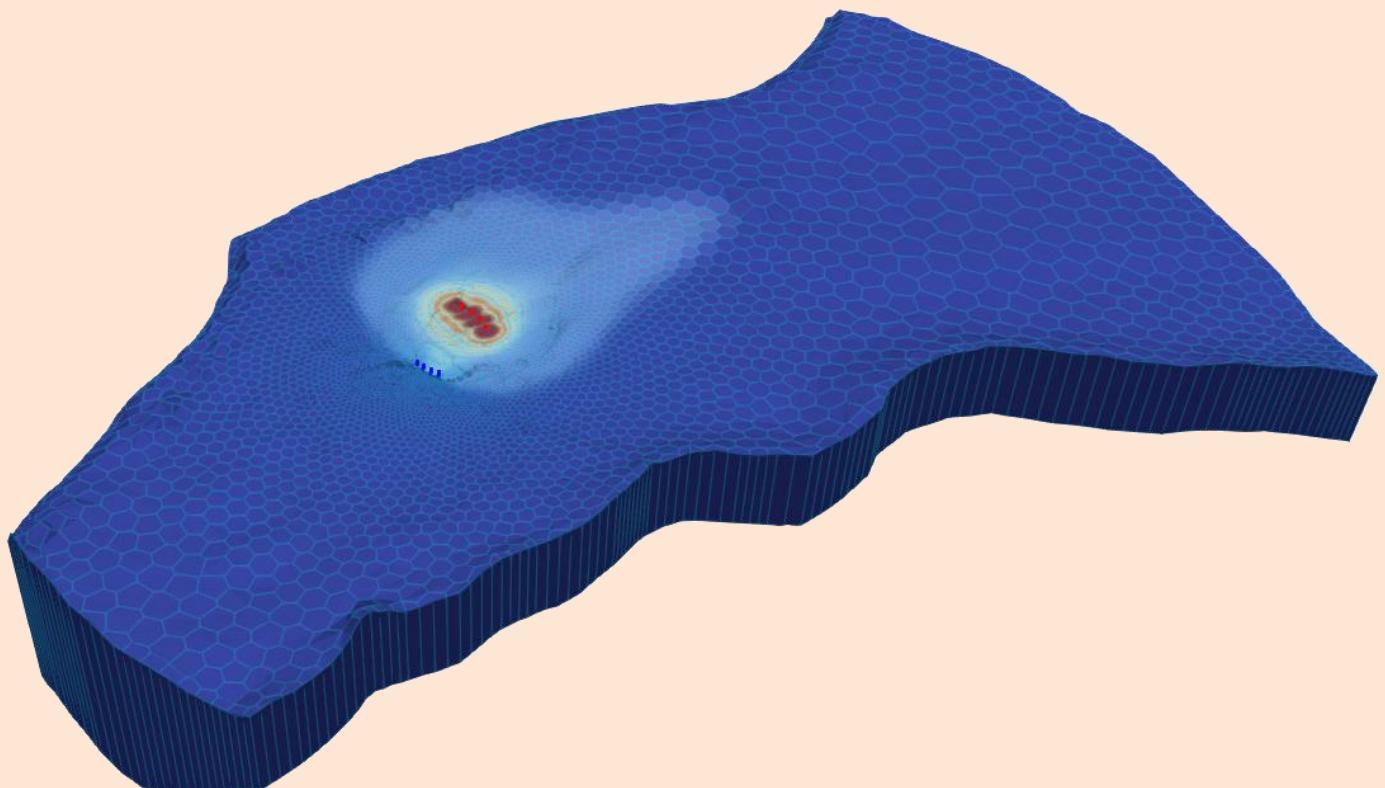


Calibration and Uncertainty Analysis for Complex Environmental Models

Second Edition



PEST: complete theory and what it means for modelling the real world

John Doherty

Calibration and Uncertainty Analysis for Complex Environmental Models

Second Edition

PEST: complete theory and what it means for modelling the real world

John Doherty

Watermark Numerical Computing,
Brisbane, Australia

First edition published in 2015 by Watermark Numerical Computing

Second edition published in 2025 by Watermark Numerical Computing

Registered office 49 Ardoyne Road

Corinda, 4075 Brisbane, Australia

www.pesthomepage.org

ISBN: 978-0-9943786-1-3

Copyright © 2025 by Watermark Numerical Computing

All rights reserved. This book or any portion thereof may not be reproduced or used without the express written permission of the publisher except for the use of brief quotations in a book review.

Preface to First Edition

PEST stands for “parameter estimation”. When it was originally written in 1994, this is all that PEST did. However over the 20 years that have elapsed since then, the capabilities of the PEST suite of software have expanded enormously. The emphasis has shifted from inversion to inversion-constrained model parameter and predictive uncertainty analysis.

The purpose of this document is to present, in one place, the theory that underpins PEST and that underpins the plethora of utility software that supports and complements PEST. In doing this, it serves the same purpose for the next generation of PEST-like software. This includes PEST++ and pyEMU. As such, it is hoped that this book provides a valuable resource for those who wish to understand inversion and inversion-constrained uncertainty analysis as it pertains to environmental models. These types of models include (but are not limited to) petroleum and geothermal reservoir models, groundwater models and surface water hydraulic and hydrologic models.

As well as presenting theory, an equally important role of this book is to draw some important conclusions from the theory. These conclusions pertain to real world model usage. Some of them question modelling practices that are commonplace today. Others suggest roles for models in environmental management beyond those which they presently play.

Although a basic knowledge of matrix algebra, and of statistical and geostatistical concepts, is required in order to follow the discussion provided herein, mathematical prerequisites are not high. Most readers will have acquired the knowledge to understand the following theory at school or through undergraduate university courses. The early part of the book reminds readers of a few basic linear analysis and statistical concepts that they may have forgotten. Refer to the many good books on these topics for further details.

Preface to Second Edition

Ten years is a long time in the modern world. Technology changes fast. So a new edition of this book is warranted.

Most of the contents of the original edition of this book, and conclusions drawn from these contents about how decision-support modelling is best undertaken, are still relevant. However some new topics require examination. Of particular importance are ensemble methods. Complementary to the importance of ensembles in modern-day numerical model history-matching and uncertainty analysis is the importance of how ensembles of heterogeneous parameter fields should be generated. On the one hand, they must be adjustable if they are to be used in data assimilation. On the other hand, they must provide realistic representations of subsurface conditions, particularly as they relate to connectedness of hydraulic properties. So generation and adjustment of nonstationary hydraulic property fields receives considerable attention in this new edition of the book.

However uncertainty quantification and uncertainty reduction do not always require the adjustment of complex parameter fields. In fact, the more complex is a parameter field, the more likely it is that its adjustment will incur predictive bias as the intricacies of hydraulic property connections are lost through the adjustment process. Data space inversion provides another possibility for harvesting information from site data and transporting this data to predictions. When implementing data space inversion, the parameters of a complex model do not require adjustment at all. Instead, data assimilation requires adjustment of the parameters

Preface

of a statistical model that links the measured past to the managed future. Depending on the circumstances, construction of this statistical model may require only a few hundred runs of a complex model. Data space inversion is therefore numerically frugal.

This new edition of the book finishes with an epilogue. This is drawn from a GMDSI monograph that was written by the author of this book. In the epilogue, predictions are made by a simple model for which history-matching poses an ill-posed inverse problem (as simple as an ill-posed inverse problem can be). These predictions are then subjected to Bayesian analysis. Based on this analysis, important general conclusions pertaining to decision-support modelling appropriateness for the making of different types of predictions are then drawn. Appropriateness is judged by how well information that is harvested from data is protected from corruption as it is transported to predictions of management interest. It is shown that opportunities for information to descend into disinformation are many; these opportunities often go unrecognized by decision-support modellers. The epilogue (and the book) finishes with a diagram that suggests optimal trajectories for decision-support modelling in different data and management contexts.

Table of Contents

1. Introduction	1
1.1 PEST and PEST++	1
1.1.1 PEST	1
1.1.2 PEST++	3
1.2 This Book	3
1.2.1 Topics	3
1.2.2 Philosophy	4
2. Vectors and Matrices	6
2.1 Introduction	6
2.2 Definitions	6
2.3 Addition and Scalar Multiplication of Matrices	7
2.4 Multiplication of Matrices	8
2.5 Partitioning of a Matrix	9
2.6 The Inverse of a Matrix	10
2.7 Determinant of a Matrix	11
2.8 The Null Space	11
2.9 The Generalized Inverse	12
2.10 Orthogonality and Projections	14
2.11 Orthogonality and Rotations	14
2.12 Eigenvectors, Eigenvalues and Positive-Definiteness	16
2.13 Singular Value Decomposition	16
2.13.1 What Singular Value Decomposition Does	16
2.13.2 Just for Completeness	21
2.14 More on Positive-Definite Matrices	21
2.15 Some Important Matrix Identities	23
3 Some Basic Probability Theory	25
3.1 A Random Variable	25
3.1.1 Probability	25
3.1.2 Confidence Intervals	26
3.2 Some Commonly Used Density Functions	27
3.2.1 Gaussian or Normal Distribution	27
3.2.2 The t or Student's t Distribution	28
3.2.3 The Chi-Squared Distribution	29
3.2.4 The F or Fisher Distribution	29
3.3 Random Vectors	29
3.4 Dependence and Independence	31
3.5 Covariance and Correlation	34
3.5.1 Concepts	34
3.5.2 Empirical Covariance Matrices	35
3.6 The Multinormal Distribution	36
3.7 Bayes Equation	37
3.8 Conditional Covariance Matrix	38
3.9 Propagation of Covariance	39

Table of Contents

3.10 A Covariance Matrix and Singular Value Decomposition	39
3.10.1 Kahunen-Loève Transformation	39
3.10.2 Principal Component Analysis.....	40
3.10.3 Inversion of an Empirical Covariance Matrix.....	41
3.10.4 Square Root of a Covariance Matrix.....	42
3.11 Spatial Correlation	42
3.12 Nonstationarity	44
3.13 Random Number Generation	45
3.14 Mahalanobis Distance	46
3.15 A Much Tougher Test	47
4. Hydraulic Property Field Generation.....	50
4.1 Background	50
4.2 Geostatistical Field Generation through Spatial Averaging.....	52
4.2.1 Advantages	52
4.2.2 Spatial Averaging Kernel	52
4.2.3 Nonstationary Spatial Correlation	54
4.2.4 Hierarchical Geostatistical Model.....	54
4.2.5 Implementation Details	55
5. Pilot Points and Conceptual Points.....	56
5.1 General	56
5.2 Pilot Points	57
5.3 Traditional Pilot Point Usage	58
5.3.1 Concepts	58
5.3.2 Accommodating Correlation Anisotropy	58
5.3.3 Accommodating Nonstationarity	58
5.3.4 Conceptual Points.....	59
5.4 Pilot Points and Spatial Averaging	60
5.4.1 Concepts	60
5.4.2 Pilot Point Parameters	62
5.4.3 Conceptual Points.....	62
5.4.4 Hierarchical History-Matching	62
5.4.5 Conditioning of Hydraulic Property Fields	63
5.5 Some Other Parameterization Devices.....	65
6. Situation Statement	67
6.1 Model Parameters.....	67
6.2 History-matching.....	68
6.3 Sampling the Posterior	72
6.3.1 Rejection Sampling	72
6.3.2 Markov Chain Monte Carlo	73
6.3.3 Ensemble Methods	75
6.3.4 Bypassing Parameters Altogether	77
6.3.5 Summary	78
6.4 Model Calibration	78
6.4.1 Calibration vs Bayes Theorem	78
6.4.2 What is Model Calibration?	79

Table of Contents

6.4.3 Regularization	80
6.4.4 So Why Calibrate?	81
6.5 Model Linearization	83
6.5.1 The Jacobian Matrix	83
6.5.2 Linear Model Analysis	84
7. Manual Regularization	86
7.1 Formulation of Equations	86
7.2 Well-Posed Inverse Problem	88
7.2.1 Data without Noise	88
7.2.2 Data with Noise	89
7.2.3 Prior Information	92
7.2.4 Objective Function Contours	93
7.2.5 Probability Density Contours of Posterior Parameter Error	95
7.2.6 Residuals	96
7.3 Post-Calibration Analysis	97
7.3.1 Problem Well-Posedness	97
7.3.2 Analysis of Residuals	102
7.3.3 Influence Statistics	104
7.4 Nonlinear Models	106
7.4.1 The Jacobian Matrix	106
7.4.2 Nonlinear Parameter Estimation	107
7.5 Critique of Manual Regularization	115
7.5.1 Man-Made Structural Noise	115
7.5.2 Parameter and Predictive Error Variance at the Scale that Matters	118
7.5.3 What is being Estimated?	120
7.5.4 Cooley's Strategy	122
7.5.5 The Role of Prior Information	123
7.5.6 Information Criteria Statistics	125
8. Mathematical Regularization	127
8.1 General	127
8.1.1 Formulation of Equations	127
8.1.2 Measurement Noise	127
8.1.3 Parameter and Predictive Error Variance	128
8.1.4 Overview of Mathematical Regularization	131
8.2 Singular Value Decomposition	133
8.2.1 Estimation of Parameters	133
8.2.2 An Example	134
8.2.3 Relationship between Estimated and Real Parameters	136
8.2.4 Kahanen-Loève Transformation	138
8.2.5 Predictive Error Variance	141
8.2.6 Well-Posedness and Flow of Information	145
8.2.7 Strengths and Weaknesses of Singular Value Decomposition	148
8.3 Tikhonov Regularization	150
8.3.1 Concepts	150
8.3.2 Formulation of Equations	153

Table of Contents

8.3.3 Optimization of Regularization Weight Factor	154
8.3.4 Relationship between Estimated and Real Parameters.....	156
8.4 Practical Implementation	157
8.4.1 Nonlinear Model Behaviour.....	157
8.4.2 Solution of Equations	159
8.4.3 Multiple Regularization Weight Factors	160
8.5 Learning from History-Matching.....	161
8.6 Hierarchical Inversion	165
9. Dimensional Reduction Methods.....	167
9.1 General	167
9.2 The Principle	167
9.3 SVD-Assist.....	168
9.3.1 How SVD-Assist Works	168
9.3.2 Regularization for SVD-Assisted Inversion.....	169
9.4 Ensemble Space Inversion	170
9.4.1 How ENSI Works.....	170
9.4.2 Regularization for ENSI.....	172
9.4.3 Strengths and Weaknesses	172
9.4.4 Methodology Variations.....	174
9.4.5 Post-ENSI Uncertainty Analysis.....	176
10. Linear Uncertainty and Error Analysis	177
10.1 General	177
10.2 Post-Calibration Subspace-Based Analyses.....	178
10.2.1 Parameter Identifiability.....	178
10.2.2 Other Subspace Insights	180
10.3 Post-Calibration Uncertainty Analysis.....	185
10.3.1 Formulation of Equations.....	185
10.3.2 Relative Parameter Uncertainty Variance Reduction.....	186
10.4 Model Predictions	187
10.4.1 Formulation of Equations.....	187
10.4.2 Effect of Transformation	187
10.4.3 Error and Uncertainty.....	189
10.5 Data Worth Analysis	191
10.6 Parameter Contributions to Predictive Uncertainty	192
11. Basics of Nonlinear Uncertainty Analysis.....	194
11.1 Introduction	194
11.1.1 General	194
11.1.2 Bayesian Analysis	194
11.2 Randomized Maximum Likelihood	196
11.2.1 Methodology	196
11.2.2 Respect for Bayes Equation	196
11.3 Approximate and Numerically Cheap RML	197
12 Iterative Ensemble Smoother	199
12.1 General	199
12.2 Mathematical Basics	200

Table of Contents

12.2.1 Simultaneous Parameter Adjustment	200
12.2.2 Empirical Covariance Calculation	200
12.2.3 Alternative Formulation	202
12.3 Numerical Implementation of IES	203
12.3.1 General	203
12.3.2 The Chen and Oliver Algorithm.....	203
12.3.3 Multiple Data Assimilation	205
12.3.4 Localization	206
12.3.5 Other Performance Enhancing Devices	207
12.4 Some Practicalities	209
12.4.1 How Many Realizations?	209
12.4.2 Weights and Noise	211
12.4.3 Ensemble Collapse	212
12.4.4 Integrity of Uncertainty Intervals	212
12.4.5 Selection of the Prior.....	213
12.4.6 Giving IES a Head Start.....	214
12.4.7 IES Following ENSI.....	215
12.4.8 Data Rich Environments	216
13. Other Parameter-Based Methods	217
13.1 General	217
13.2 Constrained Predictive Maximization/Minimization	217
13.2.1 Concepts	217
13.2.2 Limiting Objective Function	218
13.2.3 Predictive Noise	219
13.2.4 Implementation.....	220
13.2.5 Some Practicalities	221
13.2.6 A Heuristic Implementation	223
13.3 Simulator-Based Hypothesis-Testing.....	225
13.3.1 Calibration as Hypothesis-Testing	225
13.3.2 Simulator-Aided Decision-Making	227
13.3.3 Configuring a Model for Hypothesis-Testing	229
14. Data Space Inversion	232
14.1 Introduction	232
14.2 Linear DSI.....	233
14.2.1 General	233
14.2.2 Concepts	233
14.2.3 Comparison with other Methods	236
14.2.4 Data Worth Analysis	237
14.3 Nonlinear DSI	238
14.3.1 General	238
14.3.2 Statistical Model.....	238
14.3.3 Data Worth Analysis	241
14.3.4 Strengths and Weaknesses of DSI.....	242
14.4 DSIVC.....	245
14.4.1 Introduction	245

Table of Contents

14.4.2 Optimization	245
14.4.3 How DSIVC Works	248
14.5 Prior-Data Conflict	250
14.5.1 Background	250
14.5.2 The Alfonzo and Oliver Algorithm	251
15. Model Defects	255
15.1 Introduction	255
15.2 An Example	255
15.3 Mathematical Formulation of the Problem	260
15.3.1 A Defective Model	260
15.3.2 Parameters and Residuals	262
15.4 Defence against Structural Noise	265
15.4.1 Theory	265
15.4.2 Practice	266
15.5 Model Defects: a Pictorial Representation	267
15.6 Defect-Induced and Defect-Immune Predictive Error	274
15.6.1 Linear Analysis	274
15.6.2 Nonlinear Analysis	279
16 Other Issues	283
16.1 Introduction	283
16.2 Objective Function Formulation	283
16.2.1 General	283
16.2.2 Some Issues	284
16.3 Differences	286
16.4 Derivatives	287
16.5 Conclusions	289
17. References	291
Epilogue: The Journey of Information	297
E1. Introduction	297
E2. A Model	298
E2.1 Model Equations	298
E2.2 Model Predictions	299
E2.3 Prior Parameter Probability Distribution	300
E3. History-Matching the Model	300
E3.1 Conditioning	300
E3.2 Prediction #1	302
E3.3 Prediction #2	305
E4. Reflections on the Journey	307
E4.1 General	307
E4.2 Information Splitting	307
E4.3 Prior Parameter Probability Distribution	309
E4.4 Modelling as Information Partitioning	310
E4.5 Parameterisation Limits	310
E4.6 Numerical Burden Reduction	311
E4.7 Ameliorating Predictive Vulnerabilities	311

Table of Contents

E5. Modelling Appropriateness	314
E5.1 General.....	314
E5.2 The Variables	314
E5.3 Parameter-Based History-Matching	315
E5.4 Two Options	316
E5.5 Data Space Inversion	317

1. Introduction

1.1 PEST and PEST++

1.1.1 PEST

PEST stands for “parameter estimation”.

PEST first appeared in 1994. Upon its release, the feature that most distinguished PEST from earlier generations of parameter estimation software was its ability to communicate with a model through the model’s own input and output files. Its interface with a model is therefore “nonintrusive”. The consequences of this feature are more profound than simply convenience. In fact, this specification is essential for environmental model parameter estimation and uncertainty analysis. It allows “the model” to be a batch or script file which runs one or many simulators, together with one or many simulator preprocessors and postprocessors. An inversion or uncertainty analysis problem can therefore be tailored to suit the individual requirements of any environmental modelling and data context.

Since its original release, functionality offered by PEST has been expanded to include linear and nonlinear parameter and predictive uncertainty analysis. Model run parallelization was implemented using BEOPEST at first, and then PEST_HP. PEST’s inversion engine has become more powerful over time. Meanwhile, innumerable “dirty tricks” have been added to PEST’s inversion repertoire in order to ensure good performance in calibration settings that are characterized by simulator misbehaviour and highly nonlinear relationships between parameters and model outputs.

Over this same period, an ever-increasing number of utility programs were written to complement and enhance the use of PEST. Some of these utilities expedite PEST setup, while some of them complement PEST deployment in other ways. Some of the programs that are supplied with PEST perform tasks that are complementary to those performed by PEST while using the same PEST input datasets.

Tasks that are undertaken by the PEST support suite include (but are not limited to)

- building and checking of PEST input datasets;
- PEST pre- and postprocessing;
- model pre- and postprocessing;
- linear parameter and predictive uncertainty and error analysis;
- establishing the integrity of finite difference derivatives;
- matrix manipulation;
- global optimization;
- random number generation;
- ensemble-based uncertainty analysis;
- ensemble space inversion;
- data space inversion;
- analysis of prior-data conflict.

Over this same time period, another comprehensive utility suite was developed in order to expedite use of PEST in conjunction with groundwater models, particular members of the MODFLOW family (including MODFLOW-USG and MODFLOW 6). These are collectively known as the “PEST Groundwater Utility Suite”. Tasks that are performed by this suite actually

go well beyond those which assist model calibration and uncertainty analysis. Many of them were written to assist in plotting and visualization of model grids, properties and results. Tasks performed by members of the Groundwater Utility Suite include, but are not limited to

- querying of binary MODFLOW dependent-variable and budget output files;
- spatial and temporal interpolation of model outputs to sites and times of field measurements;
- temporal and spatial differencing of model outputs;
- MODFLOW array manipulation;
- pilot points parameterization;
- construction of covariance matrices based on stationary and nonstationary variograms;
- random hydraulic property field generation using stationary and nonstationary geostatistics;
- PEST input dataset preparation;
- model-to-GIS data transfer;
- model-to-SURFER data transfer;
- preparation of VTK files for model visualization;
- expeditious use of particles in history-matching and uncertainty analysis.

The PLPROC utility was written to expedite high-end inversion and uncertainty analysis, including for simulators that employ unstructured grids. Complex model preprocessing tasks are assigned to PLPROC through user-prepared scripts. During implementation of parameter estimation or uncertainty analysis, PLPROC is generally invoked by a model batch file prior to invoking the simulator. This batch file is run repeatedly by PEST as it calibrates a model. Tasks performed by PLPROC include, but are not limited to

- two- and three-dimensional pilot points parameterization, where spatial interpolation to a model grid is carried out using inverse-power-of-distance, kriging or radial basis functions;
- anisotropy-aware spatial interpolation;
- two- and three-dimensional random hydraulic property field generation using stationary or non-stationary geostatistics;
- manipulation of model input datasets using arithmetic operations of arbitrary complexity;
- superimposition of moveable, parameterizable structural features on a model grid;
- superimposition of moveable, parameterizable, sinuous, alluvial channels on a model grid;
- construction of covariance matrices based on stationary and nonstationary variograms;
- inter-layer hydraulic property interpolation.

Other programs available from the PEST web site perform tasks that include (but are not limited to)

- generalized model postprocessing and complementary PEST input dataset construction;
- surface water model postprocessing and complementary PEST input dataset construction;
- manipulation of time series used by MODFLOW 6;
- recharge simulation and irrigation demand evaluation using a lumped parameter model.

PEST, together with all of its support software, can be downloaded from the PEST web site at <https://pesthomepage.org/>

1.1.2 PEST++

PEST++ was released in 2012. Like PEST, it communicates with a model through the model's own input and output files. It does this using the same non-intrusive protocols as those that are used by PEST.

Initially, PEST++ offered similar regularized inversion capabilities to those offered by PEST. However, its main purpose was to provide a platform for collaborative programming that would enable implementation of a suite of model-value-adding tasks beyond those that were offered by PEST at the time. Its coding is object-oriented; it is written in C++. (PEST is written in FORTRAN.)

The PEST++ suite has expanded considerably since its inception. At the time of writing, the best-known member of the PEST++ suite is the PESTPP-IES ensemble smoother. This was released in 2018. PESTPP-IES implements ensemble-based parameter and predictive uncertainty analysis. Its algorithm includes a plethora of options that ensure robust performance when working with difficult numerical models and/or highly nonlinear inverse problems.

Other members of the PEST++ suite perform tasks that include (but are not limited to)

- regularized inversion;
- global sensitivity analysis;
- global optimization;
- generalized parallel model run supervision;
- linear parameter and predictive uncertainty analysis;
- optimization using chance constraints;
- nonlinear optimization under uncertainty;
- surrogate model development.

PEST++ optionally uses the same input files as PEST. PEST and PEST++ datasets are therefore interchangeable. However PEST++ also provides access to a new generation of input files that are not used by PEST.

PEST++ programs are available from the PEST++ GitHub site at <https://github.com/usgs/pestpp>

Support for PEST++ usage in the groundwater modelling context is provided by PyEMU. PyEMU provides similar functionality to that provided by PEST-complimentary utility suites. However this functionality is accessed through Python libraries. Assistance for MODFLOW inversion and uncertainty analysis is enabled by a close association between PyEMU and FloPy. PyEMU can be downloaded from its GitHub site at <https://github.com/pypest/pyemu>

1.2 This Book

1.2.1 Topics

This book is not a software manual. However it is meant to complement use of PEST, PEST++ and their utility support suites by explaining some of the theory on which these software packages are based. Particular attention is paid to the following topics:

- solution of well-posed inverse problems;
- solution of ill-posed inverse problems using regularized inversion;
- linear and nonlinear parameter and predictive uncertainty analysis;
- dimensional reduction methods such as ensemble space inversion;

- data space inversion;
- analysis of prior-data conflict;
- random number generation;
- stationary and nonstationary geostatistics.

By understanding the theory behind these topics, a modeller is more likely to use software of the PEST/PEST++ suites in ways that capitalize on their strengths, while understanding their weaknesses. This is important for their successful deployment, and for success of decision-support modelling whose implementation relies on these software suites.

1.2.2 Philosophy

This book is largely theoretical. However, it attempts to draw conclusions from theory that are salient to the way in which decision-support environmental modelling in general, and decision-support groundwater modelling in particular, should be implemented in real-world settings. Practical repercussions of theory that is developed within these pages are discussed extensively. They are also distilled into a narrative in the epilogue of this book where decision-support modelling is viewed as a “journey” taken by information from data to decisions. The epilogue suggests which decision-support contexts are most appropriate for deployment of the many methods that are discussed in this book based on the strengths and weaknesses of each. The final figure of the book, extracted from the epilogue, is presented below.

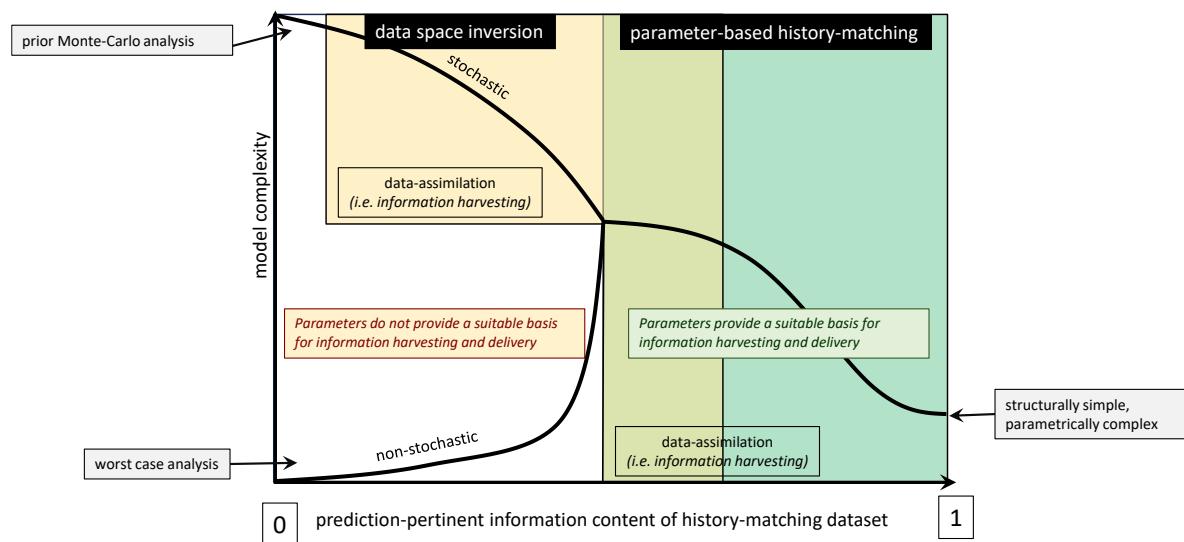


Figure 1.1. The final figure of this book. See the epilogue.

It is the opinion of the author that the role of environmental modelling in providing support for environmental decision-making is often misconstrued. Because of this, modelling is often asked to satisfy inappropriate metrics. At the same time, it often fails to implement the scientific method. Let’s take a look at why this is so.

Decision-making requires information. Information pertaining to the behaviour of complex environmental systems is diverse and multi-faceted. Ideally, it includes some insights into what may happen if a certain course of management action is pursued. It should also include insights into the repercussions of information insufficiency. This is fundamental to risk analysis. Risk analysis is fundamental to decision-making, for it is incumbent on a decision-maker to address the question “what can go wrong?”.

Numerical simulation is replete with imperfections. It should not be conceptualized as digital

replication of impossibly complex subsurface processes, and of the properties on which these processes depend. However, if properly designed, it can be used (in conjunction with software such as PEST/PEST++) to quantify and reduce the uncertainties of decision-critical predictions. At the same time, if properly deployed, it can address the question “what can go wrong?”.

Paradoxically, the more complex is a model, the more difficulty that it may experience in serving these important decision-support roles. This is because complexity attempts to express detail. However detail cannot be known. Hence detail must be expressed probabilistically. The greater the level of detail that is represented in a model, the more difficult it becomes to provide a useable prior probability distribution for this detail, and to parameterize this detail. Nevertheless, it is through adjustment of its parameters that a model is able to absorb information from field measurements of system behaviour.

Every decision-support modelling context is different. In all cases, however, the focus of modelling must be on harvesting information that may reduce the uncertainties of predictions on which management will be based. As figure 1.1 suggests, the best way to do this depends on where, and how much, decision-pertinent information is available. Where decision-pertinent information is plentiful, and a prediction is therefore data-driven, harvesting of prediction-pertinent information is easily achieved using simulator-based “machine learning”. In order to accomplish this task effectively, a model does not need to embody a “picture-perfect” replica of reality; it simply needs to fit a calibration dataset well. In contrast, where little prediction-pertinent information is available, the question that decision-support modelling must address most urgently is “what can go wrong?”. Perhaps this question is most effectively answered using worst case scenario analysis based on a simplified model.

It is the cases in between these two extremes that are the most difficult to handle. At the same time, they are the cases that are most commonly encountered in everyday modelling practice. These are the cases where uncertainties of decision-critical model prediction must be quantified as they are reduced as much as available data allow. This is not as easy as it sounds. There are many ways in which simulator-based uncertainty reduction and quantification can inculcate predictive bias. As is described in this book, there are also many safeguards that can be put in place to protect decision-critical predictions from history-matching incurred bias.

So this book is not about seeking perfection in environmental simulation, for this an impossible goal. Instead, this book describes how to create a safe passage for flow of information from data to predictions, and how to assess the predictive consequences of information insufficiency. Numerical simulation can enable both of these important decision-support processes, but only if undertaken with a thorough knowledge of theory that is presented herein, and of repercussions of this theory for everyday decision-support modelling practice.

2. Vectors and Matrices

2.1 Introduction

This chapter and the following chapter provide a brief overview of mathematical concepts on which theory covered in later chapters of this book is based. For some, these two chapters may serve as reminders of long-forgotten wisdom that was acquired at school and has since been lost in the mists of time. There are some topics, however, that a reader may not have previously encountered, for example the concepts of null space and the generalized inverse of a matrix. It is not essential that matters such as these be understood in depth. It is only important that their existence be acknowledged, and that patterns of thought that are so elegantly expressed through mathematical symbolism be retained. It is their meanings that are important, and not their details.

The reader is therefore urged to review this and the following chapters with a light heart, to be grateful for the gentle reminders of childhood days that they contain, and to wonder at the glimpses of mathematical beauty which they reveal. The going is not meant to be tough. It is meant to be nothing more than a series of rapid visitations of concepts that are important to the field of inverse problem-solving and post-calibration uncertainty analysis.

2.2 Definitions

A scalar is a single real number, such as the value of a parameter or a number that constitutes a predictive model outcome. In the present document it is represented by a lower-case, italicized character, such as x .

A vector is a collection of scalars arranged in a column. The individual scalars are its elements. They are indexed according to their place in the column. In the present text (and in most texts) a vector is denoted using a bold, lower-case letter. The vector \mathbf{x} with three elements can be written as

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \quad (2.2.1)$$

No-one who is reading this text has not heard of vectors. When taught at school vectors were often pictured as having 2 or 3 elements, these representing coordinates in two- or three-dimensional space. The vector “points” from the coordinate system origin (which has coordinates of zero) to the point in space which is designated by its elements. A vector thus has direction as well as size. Its size is its length.

The norm of a vector is another name for its size. The norm of vector \mathbf{x} is symbolized as $\|\mathbf{x}\|$. Its square is computed as

$$\|\mathbf{x}\|^2 = \sum_{i=1}^n x_i^2 \quad (2.2.2)$$

where n is the number of elements in the vector.

A vector whose elements are all zero is denoted as $\mathbf{0}$.

A matrix is a two-dimensional array of scalars. Matrices are represented using a bold capitalized letter such as \mathbf{A} . The elements of a matrix are denoted by their row and column indices. Thus the scalar a_{ij} denotes the element which occupies the i 'th row and j 'th column of the matrix \mathbf{A} . A 2×3 matrix (i.e. a matrix with two rows and three columns) named \mathbf{A} can be

written as

$$\mathbf{A} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{bmatrix} \quad (2.2.3)$$

An n -dimensional vector can be considered to be an $n \times 1$ matrix. A scalar can be considered to be a one-dimensional vector and a 1×1 matrix.

The transpose of a matrix is obtained by interchanging its rows and columns. The transpose of a matrix is indicated by a “t” superscript herein. For \mathbf{x} and \mathbf{A} defined above, \mathbf{x}^t and \mathbf{A}^t are given by

$$\mathbf{x}^t = [x_1 \ x_2 \ x_3] \quad (2.2.4)$$

$$\mathbf{A}^t = \begin{bmatrix} a_{11} & a_{21} \\ a_{12} & a_{22} \\ a_{13} & a_{23} \end{bmatrix} \quad (2.2.5)$$

Note that in the second of the above equations the subscripts on the a ’s denote the original row/column positions of these scalars in the \mathbf{A} matrix prior to transposition.

Transposing a matrix twice yields the original matrix. Thus

$$[\mathbf{A}^t]^t = \mathbf{A} \quad (2.2.6)$$

A matrix is symmetric if

$$\mathbf{A}^t = \mathbf{A} \quad (2.2.7)$$

It is obvious that a symmetric matrix must be a square matrix; as such it has an equal number of rows and columns.

The diagonal elements of a matrix are those whose row and column indices are equal. A diagonal matrix is one whose off-diagonal elements are zero. The identity matrix is a square, diagonal matrix whose diagonal elements are all 1 and whose off-diagonal elements are all 0. It is represented by \mathbf{I} . The 3×3 identity matrix is

$$\mathbf{I} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (2.2.8)$$

The trace of a matrix \mathbf{A} , denoted as $\text{tr}(\mathbf{A})$, is the sum of its diagonal elements.

A triangular matrix is one for which all elements above the diagonal are zero (this is a lower triangular matrix), or for which all elements below the diagonal are zero (this is an upper triangular matrix).

2.3 Addition and Scalar Multiplication of Matrices

Let s be a scalar. For the matrix \mathbf{A} defined in (2.2.3), the matrix $s\mathbf{A}$ is

$$s\mathbf{A} = \mathbf{A}s = \begin{bmatrix} sa_{11} & sa_{12} & sa_{13} \\ sa_{21} & sa_{22} & sa_{23} \end{bmatrix} \quad (2.3.1)$$

Matrices are added by adding their respective elements. Suppose that \mathbf{A} and \mathbf{B} are each 2×3 matrices. Then the matrix \mathbf{C} is their sum if

$$C = \begin{bmatrix} c_{11} & c_{12} & c_{13} \\ c_{21} & c_{22} & c_{23} \end{bmatrix} = \begin{bmatrix} a_{11} + b_{11} & a_{12} + b_{12} & a_{13} + b_{13} \\ a_{21} + b_{21} & a_{22} + b_{22} & a_{23} + b_{23} \end{bmatrix} \quad (2.3.2)$$

Obviously, associative and commutative rules apply, so that

$$\mathbf{A} + (\mathbf{B} + \mathbf{C}) = (\mathbf{A} + \mathbf{B}) + \mathbf{C} \quad (2.3.3)$$

and

$$\mathbf{A} + \mathbf{B} = \mathbf{B} + \mathbf{A} \quad (2.3.4)$$

2.4 Multiplication of Matrices

Let \mathbf{x} be the three-dimensional vector defined in equation (2.2.1) and \mathbf{A} be the 2×3 matrix defined in (2.2.3). Let the two-dimensional vector \mathbf{y} be the outcome of multiplication of the vector \mathbf{x} by the matrix \mathbf{A} . This is written as

$$\mathbf{y} = \mathbf{Ax} \quad (2.4.1)$$

The vector \mathbf{y} has as many rows as there are rows of \mathbf{A} . Multiplication of \mathbf{x} by \mathbf{A} can only occur if \mathbf{A} has as many columns as there are rows in \mathbf{x} . The elements of \mathbf{y} , namely y_1 and y_2 , are calculated as follows

$$y_i = \sum_{j=1}^3 a_{ij}x_j \quad (2.4.2)$$

where i can be 1 or 2 in this equation.

Matrix multiplication can be viewed in two different ways. As written in (2.4.2), the i 'th element of \mathbf{y} is calculated by travelling along the i 'th row of the matrix \mathbf{A} , and multiplying every element encountered in this journey by the element encountered in a simultaneous journey down the single column of \mathbf{x} .

There is another way to look at matrix multiplication. Let us write the matrix \mathbf{A} as

$$\mathbf{A} = [\mathbf{a}_1 \quad \mathbf{a}_2 \quad \mathbf{a}_3] \quad (2.4.3)$$

The vectors \mathbf{a}_1 , \mathbf{a}_2 and \mathbf{a}_3 are so-called “column vectors” of the matrix \mathbf{A} defined as

$$\mathbf{a}_1 = \begin{bmatrix} a_{11} \\ a_{21} \end{bmatrix}; \mathbf{a}_2 = \begin{bmatrix} a_{12} \\ a_{22} \end{bmatrix}; \mathbf{a}_3 = \begin{bmatrix} a_{13} \\ a_{23} \end{bmatrix} \quad (2.4.4)$$

Then multiplication of the vector \mathbf{x} by the matrix \mathbf{A} can be represented in the following way

$$\mathbf{y} = \mathbf{Ax} = x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + x_3\mathbf{a}_3 \quad (2.4.5)$$

The scalar product of two vectors can be represented as a matrix multiplication. Suppose that s (a scalar) is the scalar product of the vectors \mathbf{x} and \mathbf{y} . Then

$$s = \mathbf{y}^t \mathbf{x} = \sum_{i=1}^n y_i x_i \quad (2.4.6)$$

where n is the number of elements in each of \mathbf{x} and \mathbf{y} .

The norm of a vector (see equation 2.2.2) can be calculated as the square root of the scalar product of a vector with itself. That is

$$\|\mathbf{x}\| = (\mathbf{x}^t \mathbf{x})^{1/2} \quad (2.4.7)$$

Matrix multiplication readily demonstrates that

$$\mathbf{x} = \mathbf{Ix} \quad (2.4.8)$$

That is, multiplication of a vector by the identity matrix leaves the vector unchanged.

Matrix-by-matrix multiplication is defined as a simple extension of vector-by-matrix multiplication. Suppose that the matrix \mathbf{A} is multiplied by the matrix \mathbf{B} . This leads to another matrix \mathbf{C} . It is written as

$$\mathbf{C} = \mathbf{BA} \quad (2.4.9)$$

For the matrix product to exist, the number of rows of \mathbf{A} must be the same as the number of columns of \mathbf{B} . The number of rows of the product matrix \mathbf{C} is the same as that of \mathbf{B} while the number of columns of the product matrix \mathbf{C} is the same as that of \mathbf{A} . This can be written as

$$\mathbf{C}_{n \times m} = \mathbf{B}_{n \times r} \mathbf{A}_{r \times m} \quad (2.4.10)$$

where the number of rows and columns in each matrix is recorded under the matrix. Each element c_{ij} of \mathbf{C} is calculated as

$$c_{ij} = \sum_{k=1}^r b_{ik} a_{kj} \quad (2.4.11)$$

Where more than two matrices are multiplied, this is done in pairwise sequence. It is easily shown that matrix multiplication is associative. Thus

$$\mathbf{D} = \mathbf{CBA} = (\mathbf{CB})\mathbf{A} = \mathbf{C}(\mathbf{BA}) \quad (2.4.12)$$

where the brackets in the above equation indicate the order in which the operations are carried out. Note, however, that matrix multiplication is *not* commutative. In fact, unless matrices are square, the order of matrix multiplication cannot be reversed. Suppose that equation (2.4.9) holds true. Even if \mathbf{A} and \mathbf{B} are square, it does *not* therefore follow that

$$\mathbf{C} = \mathbf{AB} \quad (2.4.13)$$

The following relationships are easy to prove.

$$(\mathbf{AB})^t = \mathbf{B}^t \mathbf{A}^t \quad (2.4.14)$$

$$\mathbf{A} = \mathbf{AI} = \mathbf{IA} \quad (2.4.15)$$

In the second of the above formulas the number of rows/columns of \mathbf{I} may differ, depending on whether it pre- or post-multiplies \mathbf{A} . Suppose that \mathbf{A} is an $m \times n$ matrix. Then the \mathbf{I} matrix that pre-multiplies it is necessarily $m \times m$ and the \mathbf{I} matrix that post-multiplies it is necessarily $n \times n$. The former is thus the m -dimensional \mathbf{I} matrix whereas the latter is the n -dimensional \mathbf{I} matrix.

2.5 Partitioning of a Matrix

It is often convenient to partition a matrix into submatrices. For example the matrix \mathbf{A} may be represented as follows.

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \quad (2.5.1)$$

For such partitioning to occur, \mathbf{A}_{11} and \mathbf{A}_{12} on the one hand, and \mathbf{A}_{21} and \mathbf{A}_{22} on the other hand, must have the same number of rows. Meanwhile \mathbf{A}_{11} and \mathbf{A}_{21} on the one hand, and \mathbf{A}_{12} and \mathbf{A}_{22} on the other hand, must have the same number of columns.

Similarly, the matrix \mathbf{B} may be partitioned as

$$\mathbf{B} = \begin{bmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} \\ \mathbf{B}_{21} & \mathbf{B}_{22} \end{bmatrix} \quad (2.5.2)$$

If rows and columns of respective submatrices are compatible, then the matrix product \mathbf{AB} can

be computed as

$$\mathbf{AB} = \begin{bmatrix} \mathbf{A}_{11}\mathbf{B}_{11} + \mathbf{A}_{12}\mathbf{B}_{21} & \mathbf{A}_{11}\mathbf{B}_{12} + \mathbf{A}_{12}\mathbf{B}_{22} \\ \mathbf{A}_{21}\mathbf{B}_{11} + \mathbf{A}_{22}\mathbf{B}_{21} & \mathbf{A}_{21}\mathbf{B}_{12} + \mathbf{A}_{22}\mathbf{B}_{22} \end{bmatrix} \quad (2.5.3)$$

Matrices can be partitioned to suit the occasion. Instead of partitioning a matrix four ways (i.e. by row and by column) as expressed by equation (2.5.1), a matrix can be partitioned by column only. Thus, for example

$$\mathbf{A} = [\mathbf{A}_{11} \quad \mathbf{A}_{12}]; \quad \mathbf{B} = [\mathbf{B}_{11} \quad \mathbf{B}_{12}] \quad (2.5.4)$$

so that

$$\mathbf{AB}^t = \mathbf{A}_{11}\mathbf{B}^t_{11} + \mathbf{A}_{12}\mathbf{B}^t_{12} \quad (2.5.5)$$

2.6 The Inverse of a Matrix

Suppose that the following holds

$$\mathbf{y} = \mathbf{Ax} \quad (2.6.1)$$

Is it possible to compute \mathbf{x} from \mathbf{y} ? If \mathbf{A} is a square matrix then this may indeed be possible. Suppose that \mathbf{A} is an $n \times n$ matrix. Then calculation of \mathbf{x} from \mathbf{y} is a problem that requires solution of n linear equations (these being defined by the rows of \mathbf{A}) for n unknowns (the number of elements of \mathbf{x}). It is possible to solve these equations for \mathbf{x} provided that no row of \mathbf{A} duplicates the information contained in one or a number of the other rows of \mathbf{A} . If such duplication does indeed occur, then \mathbf{A} is said to be “rank-deficient”; this means that at least one of its rows can be expressed as a linear combination of its other rows. (The same applies to its columns.) However if the rank of \mathbf{A} is equal to the number of its rows (which is equal to the number of its columns, n in this case), then calculation of \mathbf{x} from \mathbf{y} is possible. Details of how this is best achieved can be found elsewhere. The fact that it is possible infers that a matrix exists from which \mathbf{x} can be calculated from \mathbf{y} . This matrix is referred to as the inverse of the \mathbf{A} matrix; it is denoted as \mathbf{A}^{-1} . Hence

$$\mathbf{x} = \mathbf{A}^{-1}\mathbf{y} \quad (2.6.2)$$

What if \mathbf{A} has more rows than columns? Suppose that \mathbf{A} is an $n \times m$ matrix and that n exceeds m ; can \mathbf{x} be calculated from \mathbf{y} under these circumstances? This is where things start to get complicated, for the fact that \mathbf{A} has more rows than columns infers that there is a potential for duplication of information in the \mathbf{A} matrix and \mathbf{y} vector. If this is the case, then it is also possible that these two pieces of information conflict with each other. Nevertheless, estimation of \mathbf{x} from \mathbf{y} may be required. In seeking to solve this problem we enter into the realm of parameter estimation. The problem of estimating \mathbf{x} from \mathbf{y} will be discussed extensively later in this text.

Returning to square matrices, the strict mathematical definition of the inverse of matrix \mathbf{A} is

$$\mathbf{AA}^{-1} = \mathbf{A}^{-1}\mathbf{A} = \mathbf{I} \quad (2.6.3)$$

Using this definition of the inverse of \mathbf{A} , premultiplication of both sides of (2.6.1) by \mathbf{A}^{-1} leads immediately to (2.6.2).

It is important to note that a matrix must be square to have an inverse. But it is equally important to note that not all square matrices have inverses. A square matrix does not have an inverse if it is rank-deficient. A square matrix which has an inverse is often known as a “regular” matrix. One which does not have an inverse is denoted as “singular”.

It is easily shown that if \mathbf{A} and \mathbf{B} are regular matrices then

$$(\mathbf{AB})^{-1} = \mathbf{B}^{-1}\mathbf{A}^{-1} \quad (2.6.4)$$

$$(\mathbf{A}^{-1})^t = (\mathbf{A}^t)^{-1} \quad (2.6.5)$$

$$(\mathbf{A}^{-1})^{-1} = \mathbf{A} \quad (2.6.6)$$

It is also easy to show that the inverse of a symmetric matrix is itself symmetric.

2.7 Determinant of a Matrix

The determinant of a square matrix is a scalar. It is a quantity which is well-suited to characterizing the rank of a matrix. It features in many matrix equations, including that for the multidimensional Gaussian probability distribution. It is only mentioned here. For a more detailed treatment of matrix determinants see any good book on linear algebra.

The determinant $|\mathbf{A}|$ of a one-dimensional square matrix \mathbf{A} is calculated as

$$|\mathbf{A}| = a_{11} \quad (2.7.1)$$

If \mathbf{A} is a 2×2 matrix, its determinant is calculated as

$$|\mathbf{A}| = a_{11}a_{22} - a_{12}a_{21} \quad (2.7.2)$$

Formulas for the determinant become rapidly more complex as matrix dimensions rise. Useful properties of the determinant include the following; see, for example, Koch (1999) or Mikhail (1976) for further details.

- The determinant of a triangular matrix (including a diagonal matrix) is the product of its diagonal elements.
- For square matrices \mathbf{A} and \mathbf{B}

$$|\mathbf{AB}| = |\mathbf{A}| |\mathbf{B}| \quad (2.7.3)$$

- The determinant of a matrix is zero if one of its columns is a zero vector.
- The determinant of a matrix with two rows or columns that are identical is zero.
- The determinant of a matrix with one row or column that is a multiple of another row or column is zero.
- $|\mathbf{A}| = |\mathbf{A}^t|$ (2.7.4)
- If two rows or two columns of a matrix are interchanged then $|\mathbf{A}|$ changes sign.
- If a row or a column of a matrix is multiplied by a constant, the value of the determinant is multiplied by that constant.
- The value of a determinant does not change if a multiple of a row or column is added to another row or column.
- A matrix is invertible only if its determinant is non-zero.

2.8 The Null Space

Let \mathbf{A} be an arbitrary matrix. It can be square, have more rows than columns, or have more columns than rows. If a non-zero vector \mathbf{z} can be found such that the following equality holds, then \mathbf{A} has a null space.

$$\mathbf{0} = \mathbf{Az} \quad (2.8.1)$$

The concept of the null space is important for what lies ahead. Suppose that \mathbf{y} is calculated from \mathbf{x} as

$$\mathbf{y} = \mathbf{Ax} \quad (2.8.2)$$

If the above two equations are added, we obtain

$$\mathbf{y} = \mathbf{A}(\mathbf{x} + \mathbf{z}) \quad (2.8.3)$$

Now suppose that we know \mathbf{y} and wish to back-calculate \mathbf{x} . It is obvious from (2.8.2) and (2.8.3) that inference of \mathbf{x} from \mathbf{y} is nonunique if the matrix \mathbf{A} has a null space.

Any matrix that has more columns than rows has a null space. This makes sense if back-calculation of \mathbf{x} from \mathbf{y} is thought of as an attempt to solve a set of simultaneous equations for the elements of \mathbf{x} ; when the matrix that connects these two vectors has more columns than rows there are fewer equations than unknowns.

A square matrix has a null space if it is singular, that is, if it does not have an inverse. As stated above, this arises because of duplication of information in the rows (or columns) of \mathbf{A} . A matrix with more rows than columns may or may not possess a null space.

2.9 The Generalized Inverse

Consider equation (2.8.2). Looking ahead somewhat, let us suppose that \mathbf{A} in this equation represents the action of a model on its parameters, and that these parameters are represented by the vector \mathbf{x} . Let the vector \mathbf{y} represent a set of observations from which we wish to infer values for the parameters \mathbf{x} . \mathbf{A} can be of any shape. If there are more observations than parameters \mathbf{A} will be long in the vertical direction; that is, it will have more rows than columns. If there are more parameters than observations \mathbf{A} will be long in the horizontal direction; that is, it will have more columns than rows. (Conceptually this is always the case if full account is taken of the inherent heterogeneity of natural systems.)

If \mathbf{A} has a null space, inference of \mathbf{x} from \mathbf{y} is nonunique. This means that many \mathbf{x} 's are compatible with a single \mathbf{y} . A matrix from which an \mathbf{x} which satisfies (2.8.2) can be calculated from \mathbf{y} is referred to as a generalized inverse of \mathbf{A} . It is usually denoted using a “-” superscript instead of a “-1” superscript, the latter being reserved for the strict inverse of a square matrix. Thus

$$\underline{\mathbf{x}} = \mathbf{A}^{\text{-}} \mathbf{y} \quad (2.9.1)$$

In equation (2.9.1) the underscore under \mathbf{x} is used to denote the fact that $\underline{\mathbf{x}}$ is an estimate of the \mathbf{x} of equation (2.8.2) (and a nonunique estimate at that). Where \mathbf{A} has a null space, $\mathbf{A}^{\text{-}}$ is therefore nonunique. Substitution of (2.8.2) into (2.9.1) gives

$$\underline{\mathbf{x}} = \mathbf{A}^{\text{-}} \mathbf{A} \mathbf{x} \quad (2.9.2)$$

If both sides of equation (2.9.2) are multiplied by \mathbf{A} , we obtain

$$\mathbf{A} \underline{\mathbf{x}} = \mathbf{A} \mathbf{A}^{\text{-}} \mathbf{A} \mathbf{x} \quad (2.9.3)$$

But, by definition of $\underline{\mathbf{x}}$

$$\mathbf{A} \underline{\mathbf{x}} = \mathbf{A} \mathbf{x} \quad (2.9.4)$$

This leads to the mathematical definition of the generalized inverse $\mathbf{A}^{\text{-}}$ of the matrix \mathbf{A} as

$$\mathbf{A} = \mathbf{A} \mathbf{A}^{\text{-}} \mathbf{A} \quad (2.9.5)$$

Contrast this definition with that of the strict inverse of a square, invertible matrix \mathbf{A} given by equation (2.6.3). As stated above, in contrast to \mathbf{A}^{-1} , $\mathbf{A}^{\text{-}}$ is nonunique. However where \mathbf{A} is square and invertible, then $\mathbf{A}^{\text{-}}$ and \mathbf{A}^{-1} are the same. Under these conditions $\mathbf{A}^{\text{-}}$ is unique.

Consider that \mathbf{A} has a null space and that $\mathbf{A}^{\text{-}}$ is therefore nonunique. There is one special $\mathbf{A}^{\text{-}}$, commonly denoted as \mathbf{A}^+ , which has some special properties – including that of uniqueness. \mathbf{A}^+ is variously known as the pseudoinverse of the matrix \mathbf{A} or as the Moore-Penrose inverse

of \mathbf{A} . When used to calculate $\underline{\mathbf{x}}$ through equation (2.9.1), it yields the $\underline{\mathbf{x}}$ of minimum norm for which \mathbf{x} and \mathbf{y} are related through (2.8.2). In practice \mathbf{A}^+ is obtainable through singular value decomposition of the matrix \mathbf{A} (a subject which we will visit shortly). It can be shown that the following properties characterize \mathbf{A}^+ at the same time as they make it unique; see Koch (1999) for more details. Note that the first of the following equations is simply (2.9.5) rewritten for the pseudoinverse.

$$\mathbf{A}\mathbf{A}^+\mathbf{A} = \mathbf{A} \quad (2.9.6a)$$

$$\mathbf{A}^+\mathbf{A}\mathbf{A}^+ = \mathbf{A}^+ \quad (2.9.6b)$$

$$(\mathbf{A}\mathbf{A}^+)^t = \mathbf{A}\mathbf{A}^+ \quad (2.9.6c)$$

$$(\mathbf{A}^+\mathbf{A})^t = \mathbf{A}^+\mathbf{A} \quad (2.9.6d)$$

It can also be shown that (Albert, 1972)

$$\mathbf{A}^+ = \lim_{\delta \rightarrow 0} (\mathbf{A}^t \mathbf{A} + \delta^2 \mathbf{I})^{-1} \mathbf{A}^t = \lim_{\delta \rightarrow 0} \mathbf{A}^t (\mathbf{A}\mathbf{A}^t + \delta^2 \mathbf{I})^{-1} \quad (2.9.7)$$

Equation (2.9.7) is important. Looking ahead, it says that the solution of a nonunique inverse problem achieved through Tikhonov regularization is the same as that achieved through singular value decomposition as measurement noise approaches zero. As will be discussed, singular value decomposition and Tikhonov regularization comprise the two great strands of mathematical regularization that are employed to obtain sensible solutions to ill-posed inverse problems of the type that arise in environmental model calibration (and in many other fields of study).

Equation (2.9.2) depicts the relationship between estimated $\underline{\mathbf{x}}$ and true \mathbf{x} . Recall that $\underline{\mathbf{x}}$ and \mathbf{x} are indistinguishable after multiplication by \mathbf{A} because

$$\mathbf{y} = \mathbf{A}\mathbf{x} \quad (2.9.8a)$$

and

$$\mathbf{y} = \mathbf{A}\underline{\mathbf{x}} \quad (2.9.8b)$$

Where \mathbf{A} has a null space there is no reason why $\underline{\mathbf{x}}$ should equal \mathbf{x} . Hence $\mathbf{A}^t \mathbf{A}$ of (2.9.2) is not the identity matrix. Furthermore, because of the fact that \mathbf{A} has a null space, $\mathbf{A}^t \mathbf{A}$ is rank-deficient. (The identity matrix is definitely not rank-deficient). In fact $\mathbf{A}^t \mathbf{A}$ is given a special name, this being the “resolution matrix”. Thus

$$\underline{\mathbf{x}} = \mathbf{A}^t \mathbf{A} \mathbf{x} = \mathbf{R} \mathbf{x} \quad (2.9.9)$$

Equation (2.9.6d) shows that the resolution matrix is symmetric when solution of a nonunique inverse problem is achieved using the Moore-Penrose pseudoinverse.

The resolution matrix has much to say about the nature of the vector $\underline{\mathbf{x}}$ that is achieved through solution of a nonunique inverse problem, and about how any such solution should be interpreted. Recalling the definition of vector-by-matrix multiplication, equation (2.9.9) states that each element \underline{x}_i of $\underline{\mathbf{x}}$ is obtained through weighted averaging of all the elements x_i of \mathbf{x} ; the weights used in this averaging process are the elements of the i 'th row of the resolution matrix \mathbf{R} . (This is why it is called the resolution matrix.) Obviously, the closer that \mathbf{R} resembles the identity matrix \mathbf{I} , the better is the resolving power of the inversion process. \mathbf{R} equals \mathbf{I} only when \mathbf{A} is square and regular, for then \mathbf{A}^t becomes \mathbf{A}^{-1} , and $\mathbf{A}^{-1}\mathbf{A}$ becomes \mathbf{I} . If this is not the case, then the resolution matrix is of less than full rank, and hence is non-invertible. There can therefore be many different $\underline{\mathbf{x}}$'s for a given \mathbf{x} .

2.10 Orthogonality and Projections

Two vectors are orthogonal to each other if their scalar product is equal to 0. Thus the vector \mathbf{x} is orthogonal to the vector \mathbf{y} if

$$\mathbf{x}^t \mathbf{y} = 0 \quad (2.10.1)$$

A matrix \mathbf{V} is orthonormal if each of its columns comprises a unit vector, and if those unit vectors are mutually orthogonal to each other. Let \mathbf{v}_i denote the i 'th column of \mathbf{V} . Then \mathbf{V} can be written as

$$\mathbf{V} = [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \dots \quad \mathbf{v}_m] \quad (2.10.2)$$

where m is the number of columns in \mathbf{V} . If \mathbf{V} has n rows, then each \mathbf{v}_i obviously has n elements. \mathbf{V} is orthonormal if

$$\mathbf{v}_i^t \mathbf{v}_j = 0 \quad \text{for } i \neq j \quad (2.10.3a)$$

$$\mathbf{v}_i^t \mathbf{v}_j = 1 \quad \text{for } i = j \quad (2.10.3b)$$

If \mathbf{V} is orthonormal, then obviously

$$\mathbf{V}^t \mathbf{V} = \mathbf{I} \quad (2.10.4)$$

If \mathbf{V} is also square, it follows that \mathbf{V}^t is the inverse of \mathbf{V} . Thus the transpose of a square orthonormal matrix is its inverse. From (2.6.3) it then follows that

$$\mathbf{V} \mathbf{V}^t = \mathbf{V}^t \mathbf{V} = \mathbf{I} \quad (2.10.5)$$

Let us now turn our attention to projections. A square matrix \mathbf{A} is declared to be idempotent if

$$\mathbf{A} \mathbf{A} = \mathbf{A}^2 = \mathbf{A} \quad (2.10.6)$$

If the vector \mathbf{x} is multiplied by the idempotent matrix \mathbf{A} , then \mathbf{x} is said to have undergone projection onto the range space of \mathbf{A} (i.e. the space of all possible outcomes of multiplication of any vector by \mathbf{A}). What defines an operation as a projection is that multiple operations achieve nothing that a single operation has not already achieved. This follows from the idempotency of the operation.

Suppose that the matrix \mathbf{A} is idempotent; simple matrix multiplication demonstrates that $(\mathbf{I} - \mathbf{A})$ is also idempotent. Hence $\mathbf{I} - \mathbf{A}$ is also a projection operator. If operating on the vector \mathbf{x} it can be thought of as projecting \mathbf{x} in the “other direction” to the projection achieved through multiplication by \mathbf{A} . If this “other direction” is orthogonal to the original projection direction, that is if

$$(\mathbf{I} - \mathbf{A})^t \mathbf{A} = \mathbf{0} \quad (2.10.7)$$

then \mathbf{A} is called an orthogonal projection operator. It can be shown that \mathbf{A} is an orthogonal projection operator if it is both idempotent and symmetric.

Idempotent matrices have less than full rank. They are not invertible. It can be shown that if a matrix is idempotent and regular (i.e. invertible), then it can only be \mathbf{I} . See Koch (1999) for further details.

2.11 Orthogonality and Rotations

We now consider orthogonal matrices of full rank. Thus the number of columns is equal to the number of rows, and none of the columns of the matrix are repeated.

An orthogonal matrix of full rank performs either a rotation operation on a vector, or a

reflection operation (across an axis). In either case, the length of the vector is left unchanged by this operation.

Rotation is easily demonstrated in two dimensions.

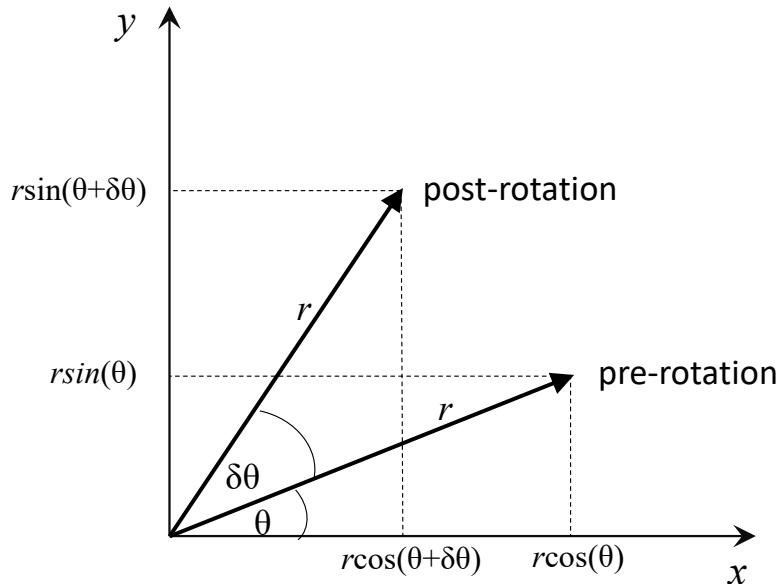


Figure 2.1. A vector or length r that undergoes rotation.

Figure 2.1 depicts a vector of length r that undergoes rotation. It is originally positioned at an angle θ above the x axis. It is rotated in the positive direction (i.e. counterclockwise) by an additional angle $\delta\theta$. Let us designate the original vector as \mathbf{w}_1 and the final vector as \mathbf{w}_2 . These vectors can be written as:

$$\mathbf{w}_1 = \begin{bmatrix} r\cos(\theta) \\ r\sin(\theta) \end{bmatrix} \quad (2.11.1a)$$

$$\mathbf{w}_2 = \begin{bmatrix} r\cos(\theta + \delta\theta) \\ r\sin(\theta + \delta\theta) \end{bmatrix} \quad (2.11.1b)$$

Now consider the matrix \mathbf{Q} , defined as follows.

$$\mathbf{Q} = \begin{bmatrix} \cos(\delta\theta) & -\sin(\delta\theta) \\ \sin(\delta\theta) & \cos(\delta\theta) \end{bmatrix} \quad (2.11.2)$$

It is readily verified that \mathbf{Q} is an orthonormal matrix. Each of its columns has a magnitude of 1.0 (square its elements and add them to verify this); the columns are perpendicular to each other (take the dot product of the two column vectors to verify this).

If we multiply \mathbf{w}_1 by the matrix \mathbf{Q} , we obtain the following vector.

$$\mathbf{Q}\mathbf{w}_1 = \begin{bmatrix} r\cos(\delta\theta)\cos(\theta) - r\sin(\delta\theta)\sin(\theta) \\ r\sin(\delta\theta)\cos(\theta) + r\cos(\delta\theta)\sin(\theta) \end{bmatrix} \quad (2.11.3)$$

Using standard trigonometric identities, the vector on the right side of equation (2.11.3) is readily verified to be \mathbf{w}_2 .

This example illustrates a more general principle. A square, full-rank orthonormal matrix with a positive determinant subjects a vector to a rotation, regardless of the number of dimensions of the vector space in which it operates. This rotation changes the direction in which the vector

points. However it does not change its magnitude.

The orthonormal matrix \mathbf{Q}^t performs a rotation by the same angle, but in the opposite direction, to the matrix \mathbf{Q} . Thus it recovers the original vector. This makes sense when it is recalled that the inverse of a square, full-rank orthonormal matrix is its transpose.

It can be shown that the determinant of an orthonormal matrix can only be +1 or -1. In the former case it performs a rotation. In the latter case it performs a reflection, possibly accompanied by a rotation. In either case, the original vector is recovered through multiplication by the matrix transpose.

2.12 Eigenvectors, Eigenvalues and Positive-Definiteness

Let \mathbf{A} be an $n \times n$ square matrix. If it is of full rank, then it has n non-zero eigenvalues; more generally \mathbf{A} has m non-zero eigenvalues, where m is the rank of \mathbf{A} . By definition, the eigenvalues λ_i of a matrix \mathbf{A} can be obtained through solution of the following equation.

$$(\mathbf{A} - \lambda_i \mathbf{I}) \mathbf{x}_i = \mathbf{0} \quad (2.12.1a)$$

That is

$$\mathbf{A} \mathbf{x}_i = \lambda_i \mathbf{x}_i \quad (2.12.1b)$$

For each eigenvalue λ_i there is a corresponding eigenvector \mathbf{x}_i .

The above definition of eigenvectors and eigenvalues seems a little abstract. Obviously, a vector undergoes no rotation when multiplied by the matrix of which it is an eigenvector. One could be forgiven for failing to see the purpose of eigenvectors and eigenvalues, if their usefulness is perceived only in terms of their ability to solve the above equation. Let the reader be assured, however, that eigenvectors and eigenvalues are some of the most important concepts in linear algebra.

A symmetric $n \times n$ matrix is positive-definite if

$$\mathbf{x}^t \mathbf{A} \mathbf{x} > 0 \text{ for all } \mathbf{x} \neq \mathbf{0} \quad (2.12.2a)$$

It is positive semi-definite if

$$\mathbf{x}^t \mathbf{A} \mathbf{x} \geq 0 \text{ for all } \mathbf{x} \neq \mathbf{0} \quad (2.12.2b)$$

Covariance matrices (see later) are always positive-definite; they have to be positive definite for probabilities to be non-negative.

It can be shown that a matrix is positive-definite only if all of its eigenvalues are positive. It can also be shown that the eigenvectors of a positive-definite matrix are orthogonal. This is important for covariance matrices as it is the basis for principal component analysis. This is discussed in the next chapter.

2.13 Singular Value Decomposition

2.13.1 What Singular Value Decomposition Does

Let \mathbf{Z} be an $n \times m$ matrix, where $n < m$. Thus \mathbf{Z} is long in the horizontal direction. This kind of matrix may be used to simulate the action of a model on its parameters where parameters are many and observations are few. (This is the case in many environmental modelling contexts where the complexity and heterogeneity of the real world is acknowledged). Singular value decomposition of \mathbf{Z} leads to the calculation of three matrices. These are denoted as \mathbf{U} , \mathbf{S} and \mathbf{V}

in the following equation.

$$\mathbf{Z} = \mathbf{U}\mathbf{S}\mathbf{V}^t \quad (2.13.1)$$

The situation is pictured in figure 2.2.

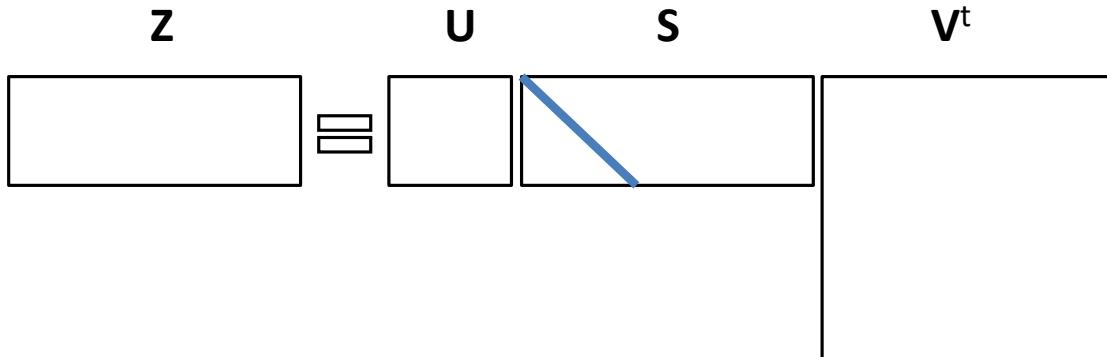


Figure 2.2. Singular value decomposition of the $n \times m$ matrix \mathbf{Z} where $n < m$.

Singular value decomposition would not be so special if the matrices that it produced were not so special. The matrix \mathbf{U} is an orthonormal matrix. Thus the vectors \mathbf{u}_i which form its columns are each of unit length and are orthogonal to each other. Collectively they encompass the range space of the matrix \mathbf{Z} . Therefore any vector \mathbf{y} that results from multiplication of any vector \mathbf{x} by the matrix \mathbf{Z} through the equation

$$\mathbf{y} = \mathbf{Z}\mathbf{x} \quad (2.13.2)$$

can be expressed as a linear combination of the \mathbf{u}_i which comprise the columns of \mathbf{U} . As such, the columns of \mathbf{U} can be considered to form a new, orthogonal, coordinate system for \mathbf{Z} output space. (\mathbf{Z} output space will become model output space later in this text when we use \mathbf{Z} to represent the action of a model on a set of parameters; the nondescript vector \mathbf{x} will then be replaced by the parameter vector \mathbf{k} .) Because \mathbf{U} is a square orthonormal matrix, it satisfies the following equation.

$$\mathbf{U}\mathbf{U}^t = \mathbf{U}^t\mathbf{U} = \mathbf{I} \quad (2.13.3)$$

where \mathbf{I} is the $n \times n$ identity matrix.

\mathbf{V} is another orthonormal matrix. It is flipped on its side in figure 2.2 which represents its transpose. The columns of \mathbf{V} span the input space of \mathbf{Z} which, in the context of the present text, is parameter space. Because it is orthonormal then

$$\mathbf{V}\mathbf{V}^t = \mathbf{V}^t\mathbf{V} = \mathbf{I} \quad (2.13.4)$$

where \mathbf{I} is the $m \times m$ identity matrix. Thus any vector \mathbf{x} on which \mathbf{Z} can operate can be expressed as a linear combination of the columns of \mathbf{V} , these comprising the m unit vectors \mathbf{v}_i which thereby form the rows of \mathbf{V}^t .

All of the elements of the matrix \mathbf{S} are zero except for those along its diagonal; its diagonal is pictured in figure 2.2. These diagonal elements are all either positive or zero. They are normally arranged from highest to lowest, starting at the top. Singular values either run out at n (the number of rows of \mathbf{Z}) or become zero before that.

Let us now partition \mathbf{S} at the point where singular values either run out or become zero. Assume that they become zero before they run out. This is illustrated in figure 2.3.

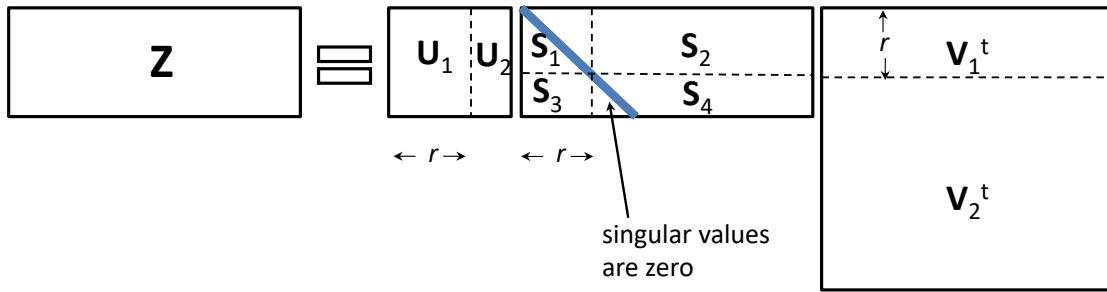


Figure 2.3. Partitioning of the U , S and V matrices according to whether singular values are zero or non-zero.

If all singular values are non-zero, then r in figure 2.2 is equal to n . Otherwise it is less than n . With S and V partitioned accordingly, equation (2.13.1) can be written as

$$Z = [U_1 \quad U_2] \begin{bmatrix} S_1 & S_2 \\ S_3 & S_4 \end{bmatrix} \begin{bmatrix} V_1^t \\ V_2^t \end{bmatrix} \quad (2.13.5)$$

$$= [U_1 S_1 + U_2 S_3 \quad U_1 S_2 + U_2 S_4] \begin{bmatrix} V_1^t \\ V_2^t \end{bmatrix} \quad (2.13.6)$$

$$= U_1 S_1 V_1^t + U_2 S_3 V_1^t + U_1 S_2 V_2^t + U_2 S_4 V_2^t \quad (2.13.7)$$

$$= U_1 S_1 V_1^t \quad (2.13.8)$$

The last of the above equations follows from its predecessor because all elements of S_2 , S_3 and S_4 are zero. We are now in a position to solve equation (2.13.2) for x . Actually we cannot solve this equation for the original x but for a vector \underline{x} which is of minimum norm and which satisfies (2.13.2). (Recall from the discussion of the Moore-Penrose pseudoinverse that singular value decomposition gives us a solution to the inverse problem which is of minimum norm). If both sides of (2.13.2) are pre-multiplied by U_1^t while the right side of (2.13.8) is substituted for Z in (2.13.2) we obtain

$$U_1^t y = U_1^t U_1 S_1 V_1^t x = S_1 V_1^t x \quad (2.13.9)$$

where equality of $U_1^t U_1$ to I follows from the orthonormality of U_1 . We now multiply each side of the above equation by S_1^{-1} . It is important to note that the inverse of S_1 cannot fail to exist because S_1 is a square diagonal matrix with non-zero elements down its diagonal; partitioning of the U , S and V matrices was designed specifically to achieve this outcome. We obtain (after recalling that the product of a matrix and its inverse is I)

$$S_1^{-1} U_1^t y = V_1^t x \quad (2.13.10)$$

Both sides of (2.13.10) are now pre-multiplied by V_1 to obtain

$$V_1 S_1^{-1} U_1^t y = V_1 V_1^t x = \underline{x} \quad (2.13.11)$$

To verify that \underline{x} satisfies equation (2.13.2) substitute it into this equation. We obtain

$$Z \underline{x} = U_1 S_1 V_1 V_1^t x = U_1 S_1 V_1^t x = Z x \quad (2.13.12)$$

where use is made of the fact that $V_1 V_1^t$ is equal to I because of orthonormality of V_1 .

Equation (2.13.11) can be re-written as

$$\underline{x} = V_1 V_1^t x \quad (2.13.13)$$

Obviously $\mathbf{V}_1\mathbf{V}_1^t$ is not equal to \mathbf{I} . $\mathbf{V}_1\mathbf{V}_1^t$ is an $m \times m$ matrix, where m is the number of elements in \mathbf{x} . (Later in this book, this will be equal to the number of parameters used by a model.) In contrast, if we reverse the order of multiplication, $\mathbf{V}_1^t\mathbf{V}_1$ is an $r \times r$ identity matrix where r is, at most, equal to the number of elements in \mathbf{y} (number of observations used in the calibration process later in this text, which is assumed, for the moment, to be less than the number of elements of \mathbf{x}). The order of matrix multiplication is very important.

Comparing (2.13.13) with (2.9.9), it is apparent that $\mathbf{V}_1\mathbf{V}_1^t$ is the resolution matrix that specifies the relationship between estimated $\underline{\mathbf{x}}$ and true \mathbf{x} when an ill-posed problem is solved using singular value decomposition. Much more will be said of this later. For now it can be stated that $\mathbf{V}_1\mathbf{V}_1^t$ is an orthogonal projection operator. The fact that it is idempotent and symmetric can be readily verified by direct multiplication. It projects \mathbf{x} from the m -dimensional space in which it resides, onto a subspace of r dimensions. This is visualized in figure 2.4, where m is 3 and r is 2.

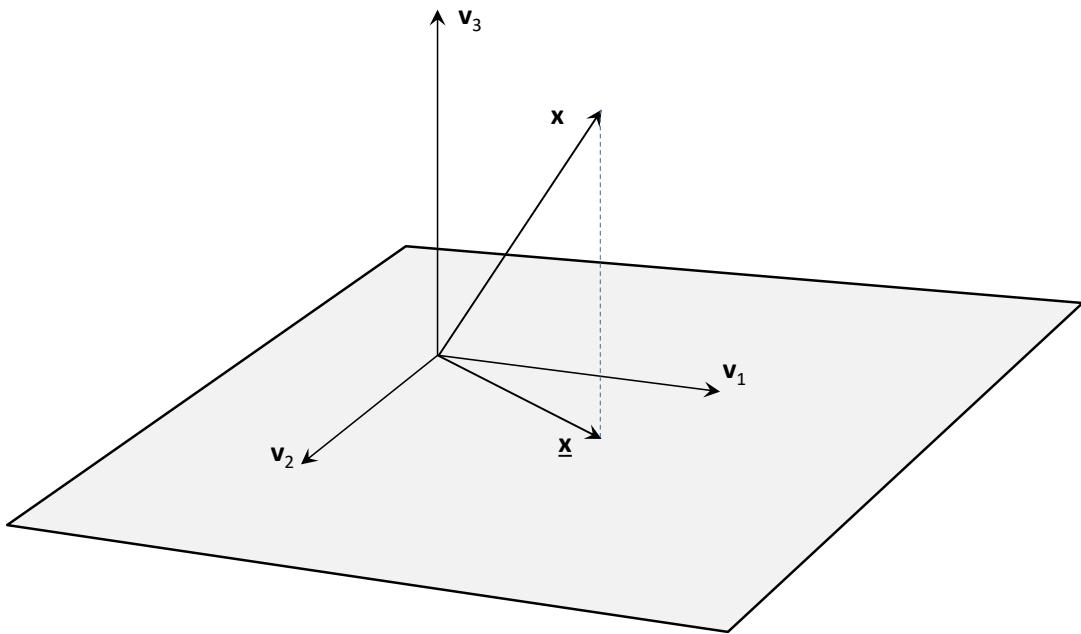


Figure 2.4. $\underline{\mathbf{x}}$ is related to \mathbf{x} through orthogonal projection onto a smaller dimensional subspace.

The subspace onto which \mathbf{x} is projected to form $\underline{\mathbf{x}}$ is spanned by the column vectors of \mathbf{V}_1 . Throughout this book this subspace is referred to as the “calibration solution space”, or simply as the “solution space”. The orthogonal complement to this subspace is the null space of \mathbf{Z} . The null space is spanned by the column vectors of \mathbf{V}_2 . Hence any vector which lies entirely in the null space of \mathbf{Z} can be formulated as a linear combination of the orthogonal unit vectors which comprise the columns of \mathbf{V}_2 . Equation (2.13.8), together with mutual orthogonality of the columns of \mathbf{V}_1 and \mathbf{V}_2 , instantly shows that multiplication of such a vector by \mathbf{Z} has an outcome of $\mathbf{0}$.

Because the column vectors of \mathbf{V}_1 and \mathbf{V}_2 collectively span the entirety of m -dimensional space, the following equation holds

$$\mathbf{V}_1\mathbf{V}_1^t + \mathbf{V}_2\mathbf{V}_2^t = \mathbf{I} \quad (2.13.14)$$

where \mathbf{I} is the $m \times m$ identity matrix. An immediate consequence of equation (2.13.14) is that if any vector is simultaneously projected onto the solution space of \mathbf{Z} and onto its null space, and

the two projected vectors are then added together, the result is the original vector. Hence from (2.13.13) \underline{x} can be considered as the original vector x with its null space component removed.

Our discussion of singular value decomposition so far has assumed that the matrix Z is elongated in the horizontal direction. The converse case is depicted in figure 2.5.

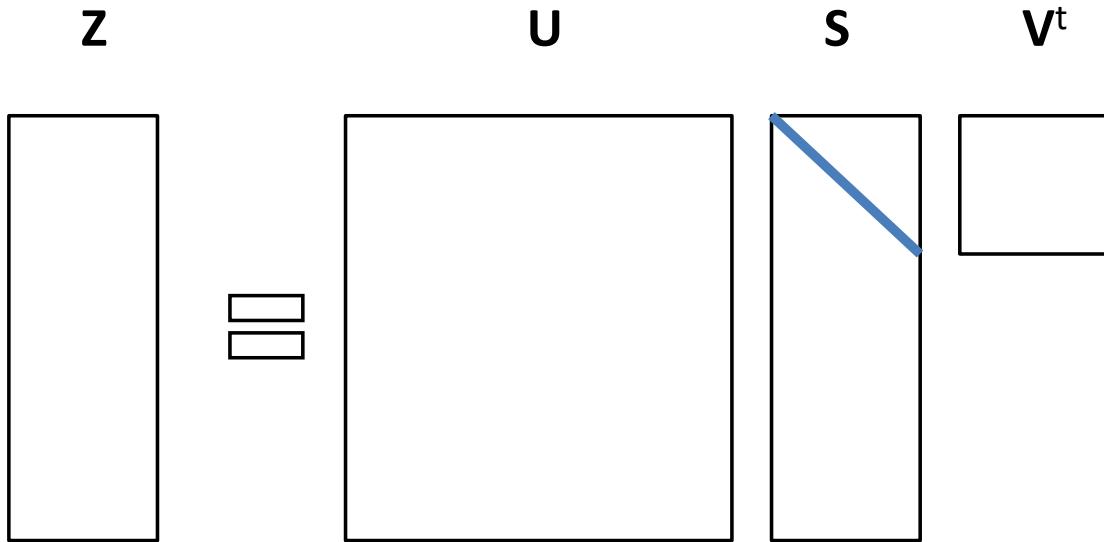


Figure 2.5. Singular value decomposition of the $n \times m$ matrix Z where $n > m$.

Where Z has more rows than columns, singular values are either all positive or become zero before the right side of the S matrix is reached. In this case, matrix partitioning that is undertaken to isolate non-zero singular values is likely to affect U more than V . See figure 2.6. If there are no zero-valued singular values, then V is not partitioned at all. In this case the inverse problem is well-posed as there is no null space.

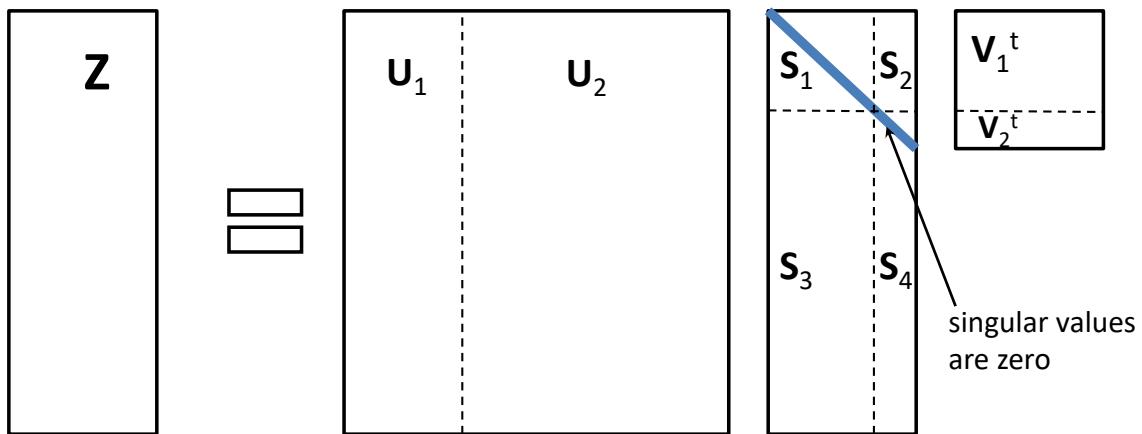


Figure 2.6. Partitioning of the U , S and V matrices where Z has more rows than columns.

So what are the implications of partitioning of U ? Where $m > n$ (i.e. where the number of columns of Z is greater than its number of rows) and where Z possesses no zero-valued singular values (because there is no room left in the S matrix to hold them), then U is not partitioned at all. However where $m < n$ (i.e. where the number of columns of Z is less than its number of rows), then U is indeed partitioned, regardless of whether any singular values are zero or not. Partitioning of U implies partitioning of the range space of the Z matrix. At best, the range space of Z can only be as high as its number of columns. Hence U_1 will possess, at most, only

as many columns as there are columns of \mathbf{Z} .

Say, for example, that \mathbf{Z} has 10 columns. (This will happen if a model has 10 parameters). Then any vector that is an outcome of multiplication by \mathbf{Z} is a linear combination of these 10 columns. Let us further suppose that \mathbf{Z} has 100 rows (which will occur if the calibration dataset is comprised of 100 observations). There are 100 degrees of freedom in the way that 100 numbers can express themselves. These are expressed by the 100 columns of \mathbf{U} . However the range space (i.e. output space) of the \mathbf{Z} matrix can only express 10 degrees of freedom. Hence \mathbf{U}_1 has at most 10 columns. The degrees of freedom in the range space of \mathbf{Z} will actually be less than this if \mathbf{Z} has a null space; under these circumstances the \mathbf{U}_1 matrix will possess fewer than 10 columns.

Suppose that we wish to solve an inverse problem based on a matrix \mathbf{Z} in which the components of \mathbf{x} are estimated from the components of an arbitrary vector \mathbf{y} . Then, to the extent that \mathbf{y} is not in the range space of \mathbf{Z} (and of \mathbf{U}_1), and is therefore not expressible as a linear combination of the columns of \mathbf{Z} (and of \mathbf{U}_1), any excess variability that it possesses beyond this can only be characterized in terms of the columns of \mathbf{U}_2 . This excess variability must simply be written off as “noise” and minimized as much as possible, as no \mathbf{x} can explain it.

Finally a geometric interpretation of SVD can be offered.

Suppose that an arbitrary vector \mathbf{x} is multiplied by a matrix \mathbf{Z} . First \mathbf{x} is multiplied by the orthonormal matrix \mathbf{V}^t . It is thus rotated and/or reflected in the input space of the matrix \mathbf{Z} in order to create a new vector of the same magnitude. This rotated/reflected vector is then shrunk or expanded along each coordinate axis. The factors by which individual components of this rotated/reflected vector are multiplied are the singular values that comprise the diagonal elements of the matrix \mathbf{S} . Another rotation/reflection is then applied through the \mathbf{U} matrix; this takes place in the range space of the matrix \mathbf{Z} . Matrix multiplication is then complete. Note, that the only part of this operation that involves expansion or contraction is embodied in singular values of the \mathbf{S} matrix.

2.13.2 Just for Completeness

You may be wondering what happens when a square orthonormal matrix is subjected to singular value decomposition. Let \mathbf{Q} be such a matrix. There is some choice here. For example, the following two equations satisfy the requirements of SVD. In the first case \mathbf{U} is equal to \mathbf{Q} and \mathbf{V} is equal to \mathbf{I} . In the second case \mathbf{U} is equal to \mathbf{I} and \mathbf{V} is equal to \mathbf{Q}^t .

$$\mathbf{Q} = \mathbf{QII}^t \tag{2.13.15a}$$

$$\mathbf{Q} = \mathbf{IIQ} \tag{2.13.15b}$$

The important thing is that the singular values are all 1.0. This is an outcome of the fact that \mathbf{Q} implements no stretching or shrinkage on any vector on which it operates.

2.14 More on Positive-Definite Matrices

If \mathbf{A} is a positive-definite matrix (such as a covariance matrix, as will be discussed shortly) then it is square and symmetric. Furthermore, its eigenvalues are all positive. If it is subjected to singular value decomposition, then the \mathbf{U} and \mathbf{V} matrices that emerge from this process are the same. Let us refer to each of these as \mathbf{E} , and to the square diagonal matrix of singular value as \mathbf{F} . Subjecting \mathbf{A} to singular value decomposition, we obtain from (2.13.1)

$$\mathbf{A} = \mathbf{EFE}^t \tag{2.14.1}$$

As usual, \mathbf{E} is an orthonormal matrix. As such, it performs a rotation in vector space, a subject to which we will return in the next chapter. (It can be shown that \mathbf{E} does not perform any reflection operations.) Furthermore, because \mathbf{E} is an orthonormal matrix, its columns are orthogonal unit vectors. Let us denote these columns as $\mathbf{e}_1, \mathbf{e}_2 \dots \mathbf{e}_n$, where \mathbf{E} is assumed to be an $n \times n$ matrix. If we now post-multiply both sides of \mathbf{A} by \mathbf{E} and invoke orthonormality of \mathbf{E} to set $\mathbf{E}^t \mathbf{E}$ equal to \mathbf{I} we obtain

$$\mathbf{AE} = \mathbf{EF} \quad (2.14.2)$$

On partitioning \mathbf{E} into its individual columns this becomes

$$\mathbf{A}[\mathbf{e}_1 \ \mathbf{e}_2 \ \dots \ \mathbf{e}_n] = [\mathbf{e}_1 \ \mathbf{e}_2 \ \dots \ \mathbf{e}_n]\mathbf{F} \quad (2.14.3)$$

Because \mathbf{F} is a diagonal matrix the above equation can be written as individual equations, each of which features an individual diagonal element of \mathbf{F} . These diagonal elements are scalars; we denote the i 'th diagonal element of \mathbf{F} as f_i . We thus obtain

$$\mathbf{A}\mathbf{e}_1 = f_1\mathbf{e}_1 \quad (2.14.4a)$$

$$\mathbf{A}\mathbf{e}_2 = f_2\mathbf{e}_2 \quad (2.14.4b)$$

.

$$\mathbf{A}\mathbf{e}_n = f_n\mathbf{e}_n \quad (2.14.4c)$$

Comparing these equations with equation (2.12.1) we see that \mathbf{e}_i are the eigenvectors of \mathbf{A} while f_i are its eigenvalues. Because of the orthonormality of \mathbf{E} , the eigenvectors of \mathbf{A} are mutually orthogonal.

Singular value decomposition can help us in other ways. It allows us to find the square root of a positive-definite matrix. We define the square root of matrix \mathbf{A} as a matrix that, if multiplied by itself, equals \mathbf{A} . Recalling that singular values are always positive, the square root of \mathbf{A} is readily written as

$$\mathbf{A}^{1/2} = \mathbf{EF}^{1/2}\mathbf{E}^t \quad (2.14.5)$$

where $\mathbf{F}^{1/2}$ is a diagonal matrix whose elements are the square roots of those of \mathbf{F} . Simple matrix multiplication shows that

$$\mathbf{F}^{1/2}\mathbf{F}^{1/2} = \mathbf{F} \quad (2.14.6)$$

Recalling the orthonormality of \mathbf{E} , from which it follows that $\mathbf{E}^t \mathbf{E}$ is equal to \mathbf{I} , it then follows from (2.14.5) that

$$\mathbf{A}^{1/2}\mathbf{A}^{1/2} = \mathbf{A} \quad (2.14.7)$$

Singular value decomposition also allows us to rapidly write the inverse of a positive definite matrix. The inverse of the matrix \mathbf{A} of equation (2.14.1) is:

$$\mathbf{A}^{-1} = \mathbf{EF}^{-1}\mathbf{E}^t \quad (2.14.8)$$

Multiplication of \mathbf{A} of equation (2.14.1) by \mathbf{A}^{-1} of equation (2.14.8) results in the following sequence of equations.

$$\mathbf{AA}^{-1} = \mathbf{E}\mathbf{F}\mathbf{E}^t\mathbf{E}\mathbf{F}^{-1}\mathbf{E}^t = \mathbf{E}\mathbf{F}\mathbf{F}^{-1}\mathbf{E}^t = \mathbf{E}\mathbf{E}^t = \mathbf{I} \quad (2.14.9a)$$

The last part of equation (2.14.9a) holds because \mathbf{E} is an orthogonal matrix of full rank. Hence $\mathbf{E}\mathbf{E}^t$ is a projection operator onto the entirety of the input space of \mathbf{E} . That is, it is the identity matrix \mathbf{I} .

Similarly, multiplication of \mathbf{A}^{-1} of equation (2.14.8) by \mathbf{A} of equation (2.14.1) results in the

following sequence of equations.

$$\mathbf{A}^{-1}\mathbf{A} = \mathbf{E}\mathbf{F}^{-1}\mathbf{E}^t\mathbf{E}\mathbf{F}\mathbf{E}^t = \mathbf{E}\mathbf{F}^{-1}\mathbf{F}\mathbf{E}^t = \mathbf{E}\mathbf{E}^t = \mathbf{I} \quad (2.14.9b)$$

2.15 Some Important Matrix Identities

We close this chapter by deriving some matrix identities. These may seem a little obscure, but they are actually used a lot when discussing highly-parameterized parameter estimation and uncertainty analysis. In fact, they are so important that one of the formulas that is derived below has its own name.

Let \mathbf{D} and \mathbf{A} be invertible (and hence square) matrices. They are not necessarily the same size. (This becomes important later in the book.)

We start our derivation with an obvious matrix identity

$$\mathbf{B}^t\mathbf{D}^{-1}\mathbf{B}\mathbf{A}\mathbf{B}^t + \mathbf{B}^t = \mathbf{B}^t\mathbf{D}^{-1}\mathbf{B}\mathbf{A}\mathbf{B}^t + \mathbf{B}^t \quad (2.15.1)$$

Then, noting that

$$\mathbf{D}^{-1}\mathbf{D} = \mathbf{I} \quad (2.15.2)$$

and that

$$\mathbf{A}^{-1}\mathbf{A} = \mathbf{I} \quad (2.15.3)$$

we can re-write (2.15.1) as

$$\mathbf{B}^t\mathbf{D}^{-1}\mathbf{B}\mathbf{A}\mathbf{B}^t + \mathbf{B}^t\mathbf{D}^{-1}\mathbf{D} = \mathbf{B}^t\mathbf{D}^{-1}\mathbf{B}\mathbf{A}\mathbf{B}^t + \mathbf{A}^{-1}\mathbf{A}\mathbf{B}^t \quad (2.15.4)$$

Thus

$$\mathbf{B}^t\mathbf{D}^{-1}(\mathbf{B}\mathbf{A}\mathbf{B}^t + \mathbf{D}) = (\mathbf{B}^t\mathbf{D}^{-1}\mathbf{B} + \mathbf{A}^{-1})\mathbf{A}\mathbf{B}^t \quad (2.15.5)$$

By premultiplying both sides of (2.15.5) by $(\mathbf{B}^t\mathbf{D}^{-1}\mathbf{B} + \mathbf{A}^{-1})^{-1}$ and postmultiplying both sides by $(\mathbf{B}\mathbf{A}\mathbf{B}^t + \mathbf{D})^{-1}$ we obtain

$$(\mathbf{B}^t\mathbf{D}^{-1}\mathbf{B} + \mathbf{A}^{-1})^{-1}\mathbf{B}^t\mathbf{D}^{-1} = \mathbf{A}\mathbf{B}^t(\mathbf{B}\mathbf{A}\mathbf{B}^t + \mathbf{D})^{-1} \quad (2.15.6)$$

This is the first of two formulas that will serve us well later in this book. For the moment notice that the large matrix that must be inverted on the left side of equation (2.15.6) has the same dimensions as \mathbf{A} , while the large matrix that must be inverted on the right side of this equation has the same dimensions as \mathbf{D} . Perhaps \mathbf{A} and \mathbf{D} are easy to invert on their own. However the same may not apply to combinations of matrices that appear inside the brackets on either side of the above equality. Numerically, the choice of which side of the equation to use will then depend on the size of these combined matrices, particularly if one of them is large.

We now derive a formula that is often referred to as the Woodbury formula, or sometimes the Sherman-Morrison-Woodbury formula. We start off with

$$\mathbf{F} = \mathbf{A} - \mathbf{A}\mathbf{B}^t(\mathbf{B}\mathbf{A}\mathbf{B}^t + \mathbf{D})^{-1}\mathbf{B}\mathbf{A} \quad (2.15.7)$$

The importance of this formula will become apparent later in this book. Obviously, \mathbf{F} is a square matrix with the same dimensions as \mathbf{A} . We will now find an alternative expression for \mathbf{F} that may be easier to compute under some circumstances.

First, re-arrange equation (2.15.7) as follows by placing both terms inside brackets and post-multiplying by \mathbf{A} .

$$\mathbf{F} = (\mathbf{I} - \mathbf{A}\mathbf{B}^t(\mathbf{B}\mathbf{A}\mathbf{B}^t + \mathbf{D})^{-1}\mathbf{B})\mathbf{A} \quad (2.15.8)$$

Before moving to the next step, we write the obvious matrix identity

$$\mathbf{I} = (\mathbf{B}^t \mathbf{D}^{-1} \mathbf{B} + \mathbf{A}^{-1})^{-1} (\mathbf{B}^t \mathbf{D}^{-1} \mathbf{B} + \mathbf{A}^{-1}) \quad (2.15.9)$$

Use this expression to replace \mathbf{I} in equation (2.15.8).

$$\mathbf{F} = ((\mathbf{B}^t \mathbf{D}^{-1} \mathbf{B} + \mathbf{A}^{-1})^{-1} (\mathbf{B}^t \mathbf{D}^{-1} \mathbf{B} + \mathbf{A}^{-1}) - \mathbf{A} \mathbf{B}^t (\mathbf{B} \mathbf{A} \mathbf{B}^t + \mathbf{D})^{-1} \mathbf{B}) \mathbf{A} \quad (2.15.10)$$

Now use equation (2.15.6) to replace the second term inside the brackets in the above equation.

$$\mathbf{F} = ((\mathbf{B}^t \mathbf{D}^{-1} \mathbf{B} + \mathbf{A}^{-1})^{-1} (\mathbf{B}^t \mathbf{D}^{-1} \mathbf{B} + \mathbf{A}^{-1}) - (\mathbf{B}^t \mathbf{D}^{-1} \mathbf{B} + \mathbf{A}^{-1})^{-1} \mathbf{B}^t \mathbf{D}^{-1} \mathbf{B}) \mathbf{A} \quad (2.15.11)$$

Now place $(\mathbf{B}^t \mathbf{D}^{-1} \mathbf{B} + \mathbf{A}^{-1})^{-1}$ outside the brackets at the front.

$$\mathbf{F} = (\mathbf{B}^t \mathbf{D}^{-1} \mathbf{B} + \mathbf{A}^{-1})^{-1} ((\mathbf{B}^t \mathbf{D}^{-1} \mathbf{B} + \mathbf{A}^{-1} - \mathbf{B}^t \mathbf{D}^{-1} \mathbf{B}) \mathbf{A}) \quad (2.15.12)$$

Therefore

$$\mathbf{F} = (\mathbf{B}^t \mathbf{D}^{-1} \mathbf{B} + \mathbf{A}^{-1})^{-1} \quad (2.15.13)$$

We have just proved the following fundamental (and extremely useful) identity.

$$\mathbf{A} - \mathbf{A} \mathbf{B}^t (\mathbf{B} \mathbf{A} \mathbf{B}^t + \mathbf{D})^{-1} \mathbf{B} \mathbf{A} = (\mathbf{B}^t \mathbf{D}^{-1} \mathbf{B} + \mathbf{A}^{-1})^{-1} \quad (2.15.14)$$

3 Some Basic Probability Theory

3.1 A Random Variable

3.1.1 Probability

Many books have been written on statistics. Only a light coverage of this broad subject area is provided herein. Only matters that are salient to the following chapters of this book are addressed.

Suppose that \underline{x} is a random number. (For the moment we will represent a random variable by underlining it with a squiggle; this notation will be dropped later for the sake of simplicity as nearly all variables that we will employ are probabilistic in nature.) The fact that \underline{x} is a random variable means that we do not know its value; it could have a range of values, with some of these values being more likely than others. If we know this range, and if we can differentiate values that \underline{x} is likely to take from those that it is not, then we know something about \underline{x} , if not its exact value. So our ignorance of \underline{x} is not complete.

In environmental modelling \underline{x} may be the value of a parameter. Suppose that someone calibrated an environmental model and informed an expert that the calibrated value of a certain parameter was a value which the expert considered unlikely. Presumably the expert would complain. He/she would inform the modeller that his/her model needed to be re-calibrated. However other values assigned to this same parameter would not assail the expert's sensibilities in this manner. This tells us something about the nature of expertise, and about the nature of knowledge to which experts lay claim. It tells us that an expert cannot say for certain what the values of environmental system properties are. (This is why they need to be informed by the calibration process after all.) However it also tells us that an expert can recognize nonsense. In other words, expertise provides certainty about the values that model parameters *cannot* take. However it does not provide certainty about their actual values.

Expert knowledge is probabilistic in nature and should therefore be expressed using probabilistic concepts. For a given random variable \underline{x} , the space of real numbers is occupied by values that \underline{x} cannot possibly have, and some that it may have. Within this latter subspace there are some values which \underline{x} is more likely to take than others. This notion is expressed mathematically by associating a probability density function with \underline{x} . This will be denoted as $f(x)$, where x (without a squiggle underline) denotes a value that \underline{x} may take. Values of x that \underline{x} is more likely to take have higher values of $f(x)$ than those which it is less likely to take.

The notion of “probability” relates to “an event”. Different realizations of \underline{x} could be generated using a random number generator that expresses $f(x)$. The generation of an individual realization of \underline{x} is an event. Conceptually, the unknown value of a real-world system property that is represented by a model's parameter is the outcome of a random number generation event. Experts do not know the outcome of each event. Ideally, however, their expertise informs them of the probability distribution associated with a parameterization event.

The probability that \underline{x} lies between two values x_1 and x_2 is denoted as $P(x_1 < \underline{x} < x_2)$. This probability is always between zero (which means “impossible”) and unity (which means “certain”). The probability that \underline{x} lies between two values x_1 and x_2 is calculated from the probability density function of \underline{x} as the integral of that function between these two values. That is

$$P(a < x < b) = \int_a^b f(x)dx \quad (3.1.1)$$

The area under the total probability density function $f(x)$ integrates to unity. This is because the value of x forthcoming from any sampling event must lie somewhere in the space of real numbers. Thus

$$\int_{-\infty}^{\infty} f(x)dx = 1 \quad (3.1.2)$$

The probability that x is less than a value x is itself a function. This is called the cumulative distribution function of x and is denoted herein as $F(x)$. Formally

$$P(x < x) = F(x) = \int_0^x f(u)du \quad (3.1.3)$$

Equation (3.1.1) can then be expanded as

$$P(a < x < b) = \int_a^b f(x)dx = F(b) - F(a) \quad (3.1.4)$$

Obviously

$$\lim_{x \rightarrow -\infty} F(x) = 0 \quad \lim_{x \rightarrow \infty} F(x) = 1 \quad (3.1.5)$$

and

$$f(x) = \frac{dF(x)}{dx} \quad (3.1.6)$$

The “expected value” of any function of x , say $g(x)$, is the weighted average of values of $g(x)$ over all values of x , with the weighting function being the probability density function. If $E[g(x)]$ is used to symbolize the expected value of $g(x)$ then

$$E[g(x)] = \int_{-\infty}^{\infty} g(x)f(x)dx \quad (3.1.7)$$

In similar fashion, the expected value $E(x)$ of x is the average of x over all possible values that x may take, with these values being weighted according to their respective probability density function values. This is also known as the average value, or mean value of x ; for the moment we use the symbol μ_x to denote the mean value of x , so that

$$E(x) = \mu_x = \int_{-\infty}^{\infty} xf(x)dx$$

The variance of x , denoted as $\text{var}(x)$ or σ_x^2 , is the expected value of the square of the difference between x and its mean. That is

$$\sigma_x^2 = E[(x - E(x))^2] = \int_{-\infty}^{\infty} (x - \mu_x)^2 f(x)dx \quad (3.1.8)$$

The square root of the variance is known as the standard deviation.

The median of x is the value of x that divides the density function in two. If this is denoted as x_m then

$$\int_{-\infty}^{x_m} f(x)dx = \int_{x_m}^{\infty} f(x)dx = \frac{1}{2} \quad (3.1.9)$$

3.1.2 Confidence Intervals

It is common to hear expressions such as “95% confidence interval” applied to a random variable. As the name implies, a “confidence interval” identifies how confident we are that a random variable lies within a certain interval. This confidence is given by the integral of the probability density function of the random variable over that interval. For a 95% confidence interval, this integral is 0.95. There is therefore only a 5% chance that the random variable lies

outside the interval.

Confidence intervals can be one-sided or two-sided. One-sided confidence intervals can be lower or upper. To add to the confusion, confidence intervals are often denoted using the symbol “ α ” where α is a low value that indicates the probability that a random variable is *not* in the interval.

Let us introduce a subscript A to x that allows us to relate the value of x_A to the area under the probability density function of \underline{x} that lies to the right of x_A . That is

$$\int_{x_A}^{\infty} f(x)dx = A \quad (3.1.10)$$

See Figure 3.1.

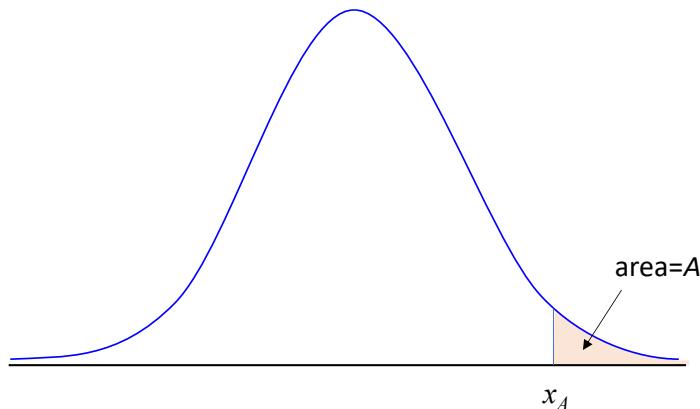


Figure 3.1. Definition of x_A .

Let us first determine the one-sided $(1-\alpha)$ confidence interval of the random variable \underline{x} . By definition, the probability that x is less than x_α is $(1-\alpha)$. The point x_α therefore marks the upper limit of the lower one-sided confidence interval of \underline{x} . That is, x_α is the upper $(1-\alpha)$ one-sided confidence limit of \underline{x} . Similarly, $x_{1-\alpha}$ is the lower $(1-\alpha)$ one-sided confidence limit of \underline{x} . It is situated at the far left of the probability density function, for the area to its right is large.

By way of example, suppose that we seek the upper and lower one sided 95% confidence limits of the random variable \underline{x} . $x_{0.05}$ is the upper one-sided 95% confidence limit of \underline{x} while $x_{0.95}$ is the lower one-sided 95% confidence limit of \underline{x} .

Let us now determine the two-sided 95% confidence interval of \underline{x} . In this case, α is still equal to 0.05, for this is the probability that \underline{x} is not within the confidence interval that we seek. However the lower and upper boundaries of the two-sided confidence interval are now $x_{0.975}$ and $x_{0.025}$, that is $x_{1-\alpha/2}$ and $x_{\alpha/2}$. There is an $\alpha/2$ chance that \underline{x} is less than $x_{1-\alpha/2}$ and an $\alpha/2$ chance that \underline{x} is greater than $x_{\alpha/2}$. Taken together, the chances of \underline{x} being outside the $(1-\alpha)$ two-sided confidence interval is therefore α , this being 5% for the present example.

Admittedly, this can become a little confusing.

3.2 Some Commonly Used Density Functions

3.2.1 Gaussian or Normal Distribution

The Gaussian or normal distribution is given by

$$f(x) = \frac{2}{\sigma_x \sqrt{2\pi}} \exp\left\{-\frac{(x-\mu_x)^2}{2\sigma_x^2}\right\} \quad (3.2.1)$$

Graphs and tables of the normal distribution function abound. As is widely known, the probability that a random normal variate lies within a standard deviation of its mean is 68.27 percent. That is

$$P[-\sigma_x < x - \mu_x < \sigma_x] = .6827 \quad (3.2.2a)$$

Probabilities related to some other intervals are as follows

$$P[-2\sigma_x < x - \mu_x < 2\sigma_x] = .9545 \quad (3.2.2b)$$

$$P[-3\sigma_x < x - \mu_x < 3\sigma_x] = .9973 \quad (3.2.2c)$$

A normal random variable is “standardized” if it is transformed such that its standard deviation is one and its mean is zero. Samples of a standardized normal random variate z are calculated from those of a normal random variable x using the equation

$$z = \frac{x - \mu_x}{\sigma_x} \quad (3.2.3)$$

The familiar “bell-shaped curve” which characterizes the normal distribution is illustrated in figure 3.2.

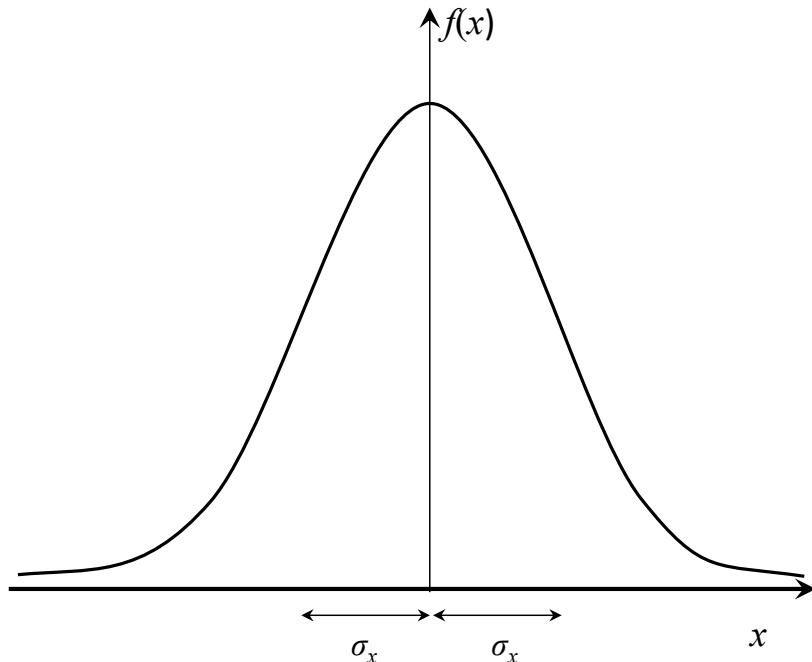


Figure 3.2. The normal probability distribution.

A number of other probability density functions feature prominently in testing and sampling. As for the normal distribution, tables and graphs of these abound. They are only mentioned briefly hereunder.

3.2.2 The t or Student's t Distribution

Let x_1, x_2, \dots, x_n be n independent (see below for a discussion of statistical independence) random variables which each possess a normal distribution with mean μ and standard deviation σ . The

random variable t defined by

$$t = \frac{\bar{x} - \mu}{s} \sqrt{n} \quad (3.2.4a)$$

where

$$\bar{x} = \frac{\sum_{i=1}^n x_i}{n} \quad \text{and} \quad s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (3.2.4b)$$

belongs to a t or Student's t distribution. This distribution describes variability in estimates of the mean of a random normal variable. A t distribution defined as above is said to possess $n-1$ degrees of freedom. Where n exceeds about 30, the Student's t distribution is almost identical to the normal distribution with mean of 0 and standard deviation $\frac{\sigma_x}{\sqrt{n}}$.

3.2.3 The Chi-Squared Distribution

Let x_1, x_2, \dots, x_n be n independent random variables, each of which is normally distributed with a mean of 0 and a standard deviation of 1. Their sum yields a random variable which belongs to a χ_n^2 distribution. This distribution depends on n which is referred to as the "degrees of freedom" of the distribution. The expected value of a χ_n^2 distribution (i.e. the mean of this distribution) is n , while its variance is $2n$. As n becomes very large, the χ^2 distribution with n degrees of freedom approaches a normal distribution with this same mean and variance.

3.2.4 The F or Fisher Distribution

Let χ_m^2 designate a random variable belonging to a chi-squared distribution with m degrees of freedom. Thus

$$\chi_m^2 = x_1^2 + x_2^2 + \dots + x_m^2 \quad (3.2.5a)$$

where x_1, x_2, \dots, x_m are independent standardized normal variates. Similarly, let χ_n^2 designate a random variable belonging to a chi-squared distribution with n degrees of freedom. Thus

$$\chi_n^2 = y_1^2 + y_2^2 + \dots + y_n^2 \quad (3.2.5b)$$

where y_1, y_2, \dots, y_n are also independent standardized normal variates. The random variable $F_{m,n}$ defined by

$$F_{m,n} = \frac{\chi_m^2}{\chi_n^2} \quad (3.2.6)$$

possess an F distribution of m, n degrees of freedom. The mean and variance of this distribution are as follows

$$\text{mean } (F_{m,n}) = \frac{n}{n-2} \quad \text{for } n > 2 \quad (3.2.7a)$$

$$\text{variance } (F_{m,n}) = \frac{2n^2(m+n-2)}{m(n-2)^2(n-4)} \quad \text{for } n > 4 \quad (3.2.7b)$$

As n approaches infinity, the F distribution with m, n degrees of freedom approaches the chi-squared distribution with m degrees of freedom.

3.3 Random Vectors

Most models have more than a single parameter. The calibration dataset against which they are calibrated is comprised of more than one observation; random noise is associated with each of these observations. We must therefore deal with collections of random numbers rather than

with individual random numbers. That is, we must deal with random vectors. A random vector is simply a vector whose elements are random numbers.

Where the variability of more than one random number must be characterized then multidimensional probabilities, and multidimensional probability density functions, must be employed.

Consider the n -dimensional random vector \mathbf{x} . Its n elements are random variates x_i . (As stated above, the squiggle under a variable that is used to denote its random status will be dropped in later notation.) Thus

$$\mathbf{x} = [x_1 \ x_2 \ \dots \ x_n]^t \quad (3.3.1)$$

An n -dimensional probability density function $f(\mathbf{x})$ can be defined that allows event probabilities to be computed for \mathbf{x} in a similar manner as for a one-dimensional probability distribution. It is defined in the following manner.

$$\begin{aligned} P(a_1 < x_1 < b_1; a_2 < x_2 < b_2; \dots; a_n < x_n < b_n) \\ = \int_{a_1}^{b_1} \int_{a_2}^{b_2} \dots \int_{a_n}^{b_n} f(x_1, x_2, \dots, x_n) dx_1 dx_2 \dots dx_n \end{aligned} \quad (3.3.2a)$$

or to put it more concisely

$$P(\mathbf{a} < \mathbf{x} < \mathbf{b}) = \int_{\mathbf{a}}^{\mathbf{b}} f(\mathbf{x}) d\mathbf{x} \quad (3.3.2b)$$

As for its one-dimensional counterpart

$$\int_{-\infty}^{\infty} f(\mathbf{x}) d\mathbf{x} = 1 \quad (3.3.3)$$

The probability that \mathbf{x} is less than \mathbf{x} is called the joint cumulative distribution function of \mathbf{x} and is denoted as $F(\mathbf{x})$. Employing vector notation, this can be expressed as

$$P(\mathbf{x} < \mathbf{x}) = F(\mathbf{x}) = \int_{\mathbf{0}}^{\mathbf{x}} f(\mathbf{u}) d\mathbf{u} \quad (3.3.4)$$

so that

$$P(\mathbf{a} < \mathbf{x} < \mathbf{b}) = \int_{\mathbf{a}}^{\mathbf{b}} f(\mathbf{x}) d\mathbf{x} = F(\mathbf{b}) - F(\mathbf{a}) \quad (3.3.5)$$

Similar limits hold for the joint cumulative distribution function as hold for the one-dimensional cumulative distribution function, namely

$$\lim_{\mathbf{x} \rightarrow \mathbf{0}} F(\mathbf{x}) = 0 \quad \lim_{\mathbf{x} \rightarrow \infty} F(\mathbf{x}) = 1 \quad (3.3.6)$$

Also

$$f(\mathbf{x}) = \frac{dF(\mathbf{x})}{d\mathbf{x}} \quad (3.3.7)$$

The expected value of a scalar function of the random vector \mathbf{x} , for example $g(\mathbf{x})$, is defined as the weighted average of that function over the n -dimensional space of \mathbf{x} ; the weighting function is the density function of \mathbf{x} . If, as above, the expected value of $g(\mathbf{x})$ is denoted as $E[g(\mathbf{x})]$ then

$$E[g(\mathbf{x})] = \int_{-\infty}^{\infty} g(\mathbf{x}) f(\mathbf{x}) d\mathbf{x} \quad (3.3.8)$$

Notice that the integration takes place over n dimensions. The mean of an element x_i of \mathbf{x} is the expected value of that element. It is given by

$$E[x_i] = \mu_i = \int_{-\infty}^{\infty} x_i f(\mathbf{x}) d\mathbf{x} \quad (3.3.9)$$

The expected values of all of the elements of \mathbf{x} can be collected into a single vector μ_x as

$$\mu_x = [\mu_1 \ \mu_2 \ \dots \ \mu_n]^t \quad (3.3.10)$$

The variance of x_i is denoted as σ^2_i and is given by

$$\sigma^2_i = E[(x_i - E(x_i))^2] = \int_{-\infty}^{\infty} (x_i - \mu_i)^2 f(\mathbf{x}) dx \quad (3.3.11)$$

Once again, the integral is over n -dimensional space.

Notice that in the above equations σ^2_{xi} has been abbreviated to σ^2_i as the “ x ” part of the subscript is redundant when x_i is part of the greater random vector \mathbf{x} . The same applies to μ_{xi} which has been abbreviated to μ_i .

The propensity for variability of a single random variable (i.e. a scalar rather than a vector variable) can be characterized by its variance. If the variable is normally distributed this is, in fact, all that is needed to characterize its variability; this is because the mean and the standard deviation alone completely define a normal distribution. (The same does not necessarily apply to other distributions.) Unfortunately, characterization of variability gets a little more complicated in multiple dimensions. This topic will be addressed shortly.

Suppose that we wish to obtain the probability density function of x_i without regard to the values that other components of the random vector \mathbf{x} may take. That is, we wish to obtain the probability density function of x_i while allowing the other components of \mathbf{x} to have any values that they please. The resulting probability density function is called the marginal probability density function of x_i . It is obtained by integrating out the other components. For illustrative purposes, suppose that n is equal to 3 and that we wish to obtain the marginal probability distribution of x_2 . Let us denote this as $f_m(x_2)$. This is obtained as

$$f_m(x_2) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x_1, x_2, x_3) dx_1 dx_3 \quad (3.3.12)$$

3.4 Dependence and Independence

Figure 3.3 shows a probability density function contour pertaining to each of two different two-variable distributions. In each case the peak of the density function lies somewhere within the contour. Outside the contour the value of the density function diminishes to zero in all directions.

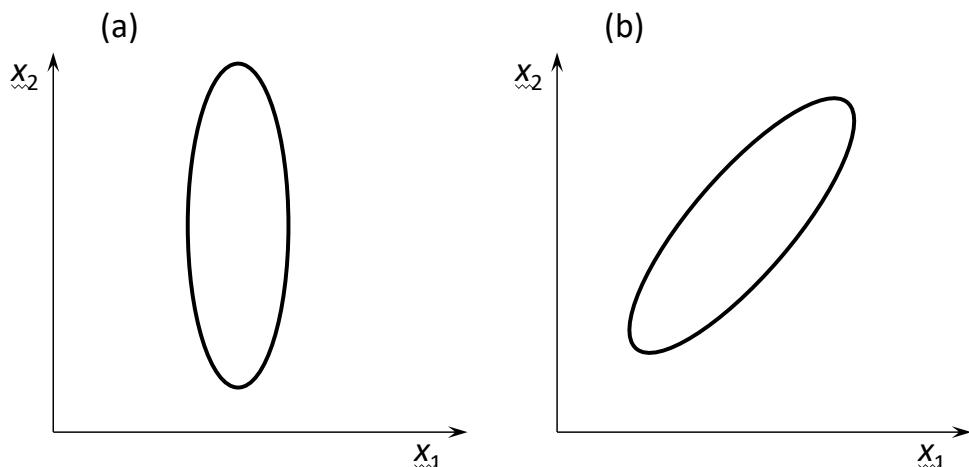


Figure 3.3. Probability density function contours where (a) random variables are independent and (b) where they are not independent.

For the distribution function depicted on the left of figure 3.3, the value taken by random variable x_2 is independent of that taken by x_1 . Thus a “statistical event” through which values are generated for both of these variables is effectively comprised of two separate events, namely generation of a value for x_1 and generation of a value for x_2 ; these events are independent of each other and can be undertaken in any order. Formally, this is expressed by the fact that the joint probability density function of x_1 and x_2 is the product of the separate, marginal probability density functions of each. For a greater number of dimensions than 2, this can be written as

$$f(x_1, x_2, \dots, x_n) = f(x_1)f(x_2)\dots f(x_n) \quad (3.4.1)$$

The same multiplicative rule applies for marginal cumulative distributions. Thus

$$F(x_1, x_2, \dots, x_n) = F(x_1)F(x_2)\dots F(x_n) \quad (3.4.2)$$

Probabilities associated with intervals are also multiplicative under these circumstances. Thus, for the two-dimensional case above, if we denote by $P(a)$ the probability that x_1 lies between a_1 and a_2 , and by $P(b)$ the probability that x_2 lies between b_1 and b_2 , the probability of joint occurrence of these events, denoted by $P(ab)$, is given by

$$P(ab) = P(a)P(b) \quad (3.4.3)$$

For the joint probability distribution whose contour is depicted on the right of figure 3.3, the values taken by x_1 and x_2 are not independent of each other. Suppose that a statistical event generates values for these two variables. If this event were to be decomposed into two parts whereby a value for x_2 is generated followed by a value for x_1 , then the probability density function used for generation of x_1 would not be its marginal distribution. This is illustrated in figures 3.4 and 3.5. Figure 3.4 shows the marginal probability density functions of the two statistically-dependent random variables. Figure 3.5 shows the so-called conditional probability density function of x_1 that must be used for generation of a value for this random variable if a value α has already been drawn for x_2 . Also shown is the conditional probability density function for x_1 if, instead of a value of α , a value of β has already been drawn for x_2 . The differences between each of these two density functions, and between the conditional and marginal density functions of x_1 are obvious.

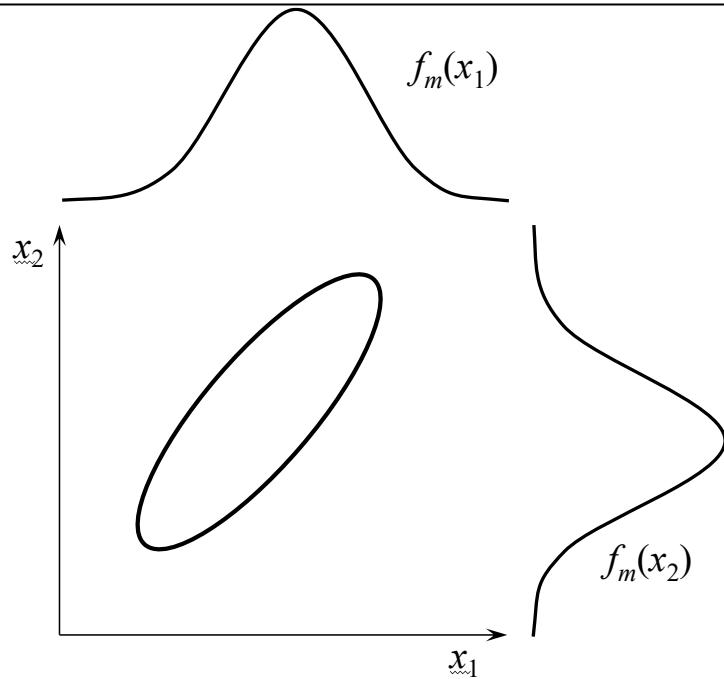


Figure 3.4. Marginal probability density functions of two random variables which are not statistically independent.

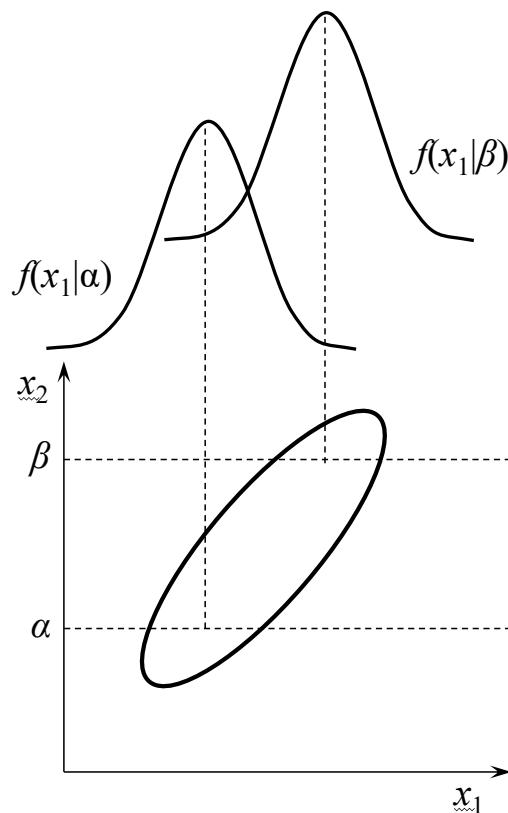


Figure 3.5. Two conditional probability density functions for random variable x_1 .

The simultaneous occurrence of event a and event b can be denoted as the event ab . (An example of such a joint event is that x_1 lies between β_1 and β_2 and that x_2 lies between α_1 and

α_2). The probability of the joint event ab is given by

$$P(ab) = P(b|a)P(a) \quad (3.4.4)$$

where $P(b|a)$ can be read as “the probability of b given a ” or “the probability of b conditional on a ”. In similar fashion, the conditional density function for x_1 conditional on x_2 is denoted as $f(x_1|x_2)$. The joint probability density function of x_1 and x_2 , i.e. $f(x_1, x_2)$ is given by

$$f(x_1, x_2) = f(x_1|x_2)f_m(x_2) \quad (3.4.5)$$

where $f_m(x_2)$ is the marginal density function of x_2 .

The above two formulas are used when decomposing a joint event into its component parts. If a value for x_2 is generated first, no restrictions on the random number generation process are operative while this value is being drawn. Hence a value for x_2 is generated using its marginal density function. However, given the statistical interdependence that exists between x_1 and x_2 , this sets a condition on the drawing of a value for x_1 . $f(x_1|x_2)$ must therefore be used when drawing a value for x_1 .

From (3.4.5) it follows that

$$f(x_1|x_2) = \frac{f(x_1, x_2)}{f_m(x_2)} \quad (3.4.6)$$

If x_1 and x_2 are statistically independent, their conditional distributions are the same as their marginal distributions. From (3.4.6) and (3.4.1)

$$f(x_1|x_2) = \frac{f(x_1, x_2)}{f_m(x_2)} = \frac{f_m(x_1)f_m(x_2)}{f_m(x_2)} = f_m(x_1) \quad (3.4.7)$$

3.5 Covariance and Correlation

3.5.1 Concepts

The variance of a single random variable is defined above. This characterizes the extent to which random samples of that variable can vary about their mean; see equation (3.1.8). The variance of a random variable which is an element of a random vector is similarly defined; see equation (3.3.11). In the latter case the integration required to calculate the expected value of the squared departure of this random vector element from its mean must be performed over multi-dimensional space, as the random variable is drawn from a multi-dimensional density function that describes both it and other components of the random vector to which it belongs. Alternatively, integration can take place over the marginal density function of the individual variable; as (3.3.12) shows, integration over the other random variables has already been performed in defining the marginal probability density function of the individual variable whose variance is sought.

Where a random variable is part of a joint probability distribution, more than just its variance may be needed to characterize its propensity to vary about its mean value, for this conveys nothing about how this propensity is linked to that of other random variables which figure in the joint distribution. Information on statistical interdependence of random variables which share a joint probability distribution can be conveyed through denoting its covariance with those other variables. Like the variance of random variable x_1 , the covariance between random variables x_1 and x_2 is expressed as an expected value; it is the expected value of the product of the departures of these two random variables from their respective means. An expected value of zero for this product indicates independence of these variables. A positive expected value indicates that the random variables tend to vary together in the same direction. A negative

expected value indicates that one random variable tends to vary in one direction from its mean as the other varies in the opposite direction.

The covariance between x_1 and x_2 is denoted as $\text{cov}(x_1, x_2)$, or more succinctly as σ_{12} . More generally, the covariance between the x_i and x_j components of the random vector \mathbf{x} is denoted as σ_{ij} . It is defined as

$$\sigma_{ij} = E[(x_i - \mu_i)(x_j - \mu_j)] = \int_{-\infty}^{\infty} (x_i - \mu_i)(x_j - \mu_j) f(x) dx \quad (3.5.1)$$

The variances of individual elements of a random vector, together with inter-element covariances between its elements, can be combined into a single matrix known as the variance-covariance matrix, or simply the covariance matrix for short. If the random vector \mathbf{x} has n dimensions, then the covariance matrix of \mathbf{x} , designated as $C(\mathbf{x})$, is a square $n \times n$ matrix defined as

$$C(\mathbf{x}) = \begin{bmatrix} \sigma_{11}^2 & \sigma_{12} & \dots & \sigma_{1n} \\ \sigma_{21} & \sigma_{22}^2 & \dots & \sigma_{2n} \\ \dots & \dots & \dots & \dots \\ \sigma_{n1} & \sigma_{n2} & \dots & \sigma_{nn}^2 \end{bmatrix} \quad (3.5.2)$$

A covariance matrix is symmetrical as $\sigma_{ij} = \sigma_{ji}$ for all i and j . It can also be shown that it must be positive-definite if it is to properly characterize the stochasticity of random variables; this ensures that no statistical events have a negative probability of occurrence.

The correlation coefficient ρ_{ij} between random vector elements x_i and x_j is calculated as

$$\rho_{ij} = \frac{\sigma_{ij}}{\sigma_i \sigma_j} = \frac{\sigma_{ij}}{\sqrt{\sigma_{ii}^2 \sigma_{jj}^2}} \quad (3.5.3)$$

The correlation coefficient varies between -1 and 1. Correlation coefficients of -1 and 1 indicate functional dependence rather than statistical inter-relatedness, so great is the propensity for the random variables to vary in the opposite or same directions respectively from their means in sympathy with each other. On the other hand, a correlation coefficient of 0 indicates statistical independence of the random variables. Correlation coefficients can be collected into a matrix; naturally enough, this is referred to as the correlation coefficient matrix. The diagonal elements of the correlation coefficient matrix are all 1.

Covariance matrixes, applied to model parameters, will figure prominently in the following chapters of this book, together with expected parameter values (i.e. parameter means). They have a significant role to play in regularized inversion (i.e. model calibration) and in uncertainty analysis. In doing this, they can be considered as an embodiment of “hydrogeological wisdom” as it pertains to a particular study site. They provide mathematical characterization of that which is known about a site, as well as that which is unknown. Knowledge is expressed through expected values, the extent of variability about these expected values (i.e. parameter variances) and the spatial connectedness of high or low values of a parameter (i.e. parameter correlation). Knowledge insufficiency is expressed by the fact that parameters require a covariance matrix, and are thus random variables.

3.5.2 Empirical Covariance Matrices

So far we have discussed what a covariance matrix means. Now we will discuss how an empirical covariance matrix can be assembled if all that we have are samples of random variables.

Suppose that we have m realizations of a random vector \underline{x} , and that \underline{x} possesses n elements. We can house all of these realizations in a matrix \mathbf{X} that has n rows and m columns. The columns of \mathbf{X} therefore comprise the m realizations of the n -dimensional random vector \underline{x} . An estimate of the mean of the i 'th element of \underline{x} is easily calculated using the standard formula:

$$\underline{x}_i = \frac{\sum_{j=1}^m x_{ij}}{m} \quad (3.5.4)$$

where x_{ij} is the element that occupies the i 'th row and j 'th column of \mathbf{X} . That is, it is the j 'th realization of the random variable x_i . This calculation can be repeated for all elements i of \underline{x} to build a vector of element means, which we denote as $\underline{\mu}$.

We now build an $n \times m$ matrix $\underline{\mathbf{X}}$ in which \underline{x} is repeated m times as each of the m columns of $\underline{\mathbf{X}}$. Use of this matrix simplifies upcoming expressions.

Empirically, the variance of the i 'th element of \underline{x} can be calculated from realizations of this element using the standard formula

$$\sigma_{ii}^{e2} = \frac{\sum_{j=1}^m (x_{ij} - \underline{x}_i)^2}{m-1} \quad (3.5.5a)$$

Similarly, the covariance between the i 'th and k 'th elements of \underline{x} can be empirically calculated as

$$\sigma_{ik}^e = \frac{\sum_{j=1}^m (x_{ij} - \underline{x}_i)(x_{kj} - \underline{x}_k)}{m-1} \quad (3.5.5b)$$

Both of the above equations can be written in concise matrix notation using the matrices \mathbf{X} and $\underline{\mathbf{X}}$ that were defined above as

$$\mathbf{C}^e(\mathbf{x}) = \frac{(\mathbf{x} - \underline{\mathbf{x}})(\mathbf{x} - \underline{\mathbf{x}})^t}{m-1} \quad (3.5.5c)$$

The greater the number of realizations that are used in computation of $\mathbf{C}^e(\mathbf{x})$ (that is, the larger is m), the better does $\mathbf{C}^e(\mathbf{x})$ approximate the true covariance matrix of \underline{x} . However for high accuracy in calculation of covariance terms for a high-dimensional \underline{x} , tens of thousands of realizations may be required.

It can be shown that an empirical covariance matrix calculated using equation (3.5.5) is positive definite if $m \geq n$, and positive semidefinite and noninvertible (i.e. singular) if $m < n$. In later chapters of this book we will encounter occasions where inversion of an empirical covariance matrix is required. A useable approximation to an inverted empirical covariance matrix can be obtained using singular value decomposition. This is discussed later in this chapter.

3.6 The Multinormal Distribution

The multinormal, or multiGaussian, probability distribution is characterized by the following density function

$$f(x_1, x_2, \dots, x_n) = f(\mathbf{x}) = \frac{e^{-\frac{1}{2}(\mathbf{x} - \mu_{\mathbf{x}})^t \mathbf{C}^{-1}(\mathbf{x})(\mathbf{x} - \mu_{\mathbf{x}})}}{(2\pi)^{\frac{n}{2}} |\mathbf{C}(\mathbf{x})|^{\frac{1}{2}}} \quad (3.6.1)$$

As can be seen from this equation, the vector of element means $\mu_{\mathbf{x}}$ and the covariance matrix $\mathbf{C}(\mathbf{x})$ are all that is required for a complete mathematical description of this distribution. For some other distributions, more specifications than this are required.

Contours of constant probability in n -dimensional parameter space are defined by the equation

$$(\mathbf{x} - \boldsymbol{\mu}_x)^t \mathbf{C}^{-1}(\mathbf{x})(\mathbf{x} - \boldsymbol{\mu}_x) = k \quad (3.6.2)$$

where k is a constant associated with the contour. In two dimensions, probability contours form ellipses; in higher dimensions they are hyper-ellipsoids. If there is no correlation between the elements of \mathbf{x} (i.e. if $\mathbf{C}(\mathbf{x})$ is a diagonal matrix), then the axes of these hyper-ellipsoids are parallel to the unit vectors $[1 \ 0 \ 0 \dots]^t$, $[0 \ 1 \ 0 \dots]^t$, etc. More generally, the axes of hyper-ellipsoids of constant probability point in the directions of the eigenvectors of $\mathbf{C}(\mathbf{x})$. The lengths of these axes are proportional to the square roots of the corresponding eigenvalues of $\mathbf{C}(\mathbf{x})$.

3.7 Bayes Equation

Equation (3.4.4) states that the probability of simultaneous occurrence of events a and b is given by the probability that event a will occur multiplied by the probability of event b conditional on event a . This equation is now repeated as (3.7.1).

$$P(ab) = P(b|a)P(a) \quad (3.7.1)$$

Obviously, the simultaneous occurrence of events a and b could also be considered as the occurrence of event b followed by the occurrence of event a . Once b has occurred, the probability of event a is then conditional on b if there is any statistical inter-relatedness between events a and b . Thus

$$P(ab) = P(a|b)P(b) \quad (3.7.2)$$

From the above two equations it immediately follows that

$$P(b|a) = \frac{P(a|b)P(b)}{P(a)} \quad (3.7.3)$$

This is Bayes equation, a subject on which much will be said in this book. The same concept can be expressed using density functions. If we drop the “m” subscript from (3.4.5) (the marginal probability density function will become the prior probability density function in later chapters), we obtain

$$f(x_1|x_2) = \frac{f(x_2|x_1)f(x_1)}{f(x_2)} \quad (3.7.4)$$

Extending the concept from random scalars to random vectors, Bayes equation in multiple dimensions becomes

$$f(\mathbf{x}_1|\mathbf{x}_2) = \frac{f(\mathbf{x}_2|\mathbf{x}_1)f(\mathbf{x}_1)}{f(\mathbf{x}_2)} \quad (3.7.5)$$

Obviously, if \mathbf{x}_1 and \mathbf{x}_2 are independent, then

$$f(\mathbf{x}_2|\mathbf{x}_1) = f(\mathbf{x}_2) \quad (3.7.6)$$

which, if substituted into (3.7.5) leads to

$$f(\mathbf{x}_1|\mathbf{x}_2) = f(\mathbf{x}_1) \quad (3.7.7)$$

However suppose that \mathbf{x}_1 and \mathbf{x}_2 are not statistically independent. Suppose also that, following a particular statistical event in which samples of \mathbf{x}_1 and \mathbf{x}_2 are generated, we acquire perfect knowledge of the values taken by the elements of \mathbf{x}_2 . $f(\mathbf{x}_2|\mathbf{x}_1)$ can then be calculated for various values of \mathbf{x}_1 and substituted into the right side of (3.7.5). The probability distribution of \mathbf{x}_1 is modified from $f(\mathbf{x}_1)$, the distribution which it would possess if sampled values of \mathbf{x}_2 are unknown, to $f(\mathbf{x}_1|\mathbf{x}_2)$. Presumably $f(\mathbf{x}_1|\mathbf{x}_2)$ is higher than $f(\mathbf{x}_1)$ for some values of \mathbf{x}_1 and lower than $f(\mathbf{x}_1)$ for other values of \mathbf{x}_1 . The probability density function of \mathbf{x}_1 is therefore modified by sampling \mathbf{x}_2 . It follows that values of \mathbf{x}_2 are informative of values of \mathbf{x}_1 . This is also apparent

from figure 3.5; in this example, sampling of \mathbf{x}_2 narrows the width of the conditional, or posterior, probability density distribution of \mathbf{x}_1 .

3.8 Conditional Covariance Matrix

Our notation will now be altered to dispense with the squiggle under a scalar or vector variable to denote its status as a random variable. This was useful up until now in order to distinguish the name of a variable from the value taken by that variable, especially in discussing probability density functions and cumulative probability density functions. As most future discussion of random variables will involve only their covariance matrices, this distinction is no longer needed. The meaning of a symbol will be clear from its context.

The acquisition of knowledge pertaining to some variables through the sampling of other variables with which they are statistically correlated is discussed above. Where the variables are collectively described by a multinormal (i.e. multiGaussian) probability density function this knowledge transfer can be quantified.

Let \mathbf{x} be a random vector. Suppose that we partition \mathbf{x} into the vectors \mathbf{x}_1 and \mathbf{x}_2 as follows

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix} \quad (3.8.1)$$

Suppose that the mean of \mathbf{x} is given by

$$\boldsymbol{\mu} = \begin{bmatrix} \boldsymbol{\mu}_1 \\ \boldsymbol{\mu}_2 \end{bmatrix} \quad (3.8.2)$$

and that the covariance matrix $\mathbf{C}(\mathbf{x})$ of \mathbf{x} can be partitioned according to the partitioning of \mathbf{x} into \mathbf{x}_1 and \mathbf{x}_2 as

$$\mathbf{C}(\mathbf{x}) = \begin{bmatrix} \mathbf{C}_{11} & \mathbf{C}_{12} \\ \mathbf{C}_{21} & \mathbf{C}_{22} \end{bmatrix} \quad (3.8.3)$$

If variability of \mathbf{x} is described by a multiGaussian distribution, then the marginal distribution of \mathbf{x}_1 is completely characterized by $\boldsymbol{\mu}_1$ and \mathbf{C}_{11} . Suppose now that we acquire perfect knowledge of values taken by elements of the \mathbf{x}_2 vector. As discussed above, this modifies the probability distribution of \mathbf{x}_1 . Let $\boldsymbol{\mu}'_1$ and \mathbf{C}'_{11} specify the mean and covariance matrix of this modified distribution. It can be shown (see, for example, Koch, 1999) that

$$\boldsymbol{\mu}'_1 = \boldsymbol{\mu}_1 + \mathbf{C}_{12}\mathbf{C}^{-1}_{22}(\mathbf{x}_2 - \boldsymbol{\mu}_2) \quad (3.8.4)$$

and

$$\mathbf{C}'_{11} = \mathbf{C}_{11} - \mathbf{C}_{12}\mathbf{C}^{-1}_{22}\mathbf{C}_{21} \quad (3.8.5)$$

Just to whet the reader's appetite, let us look ahead for a moment. Suppose that the \mathbf{x}_1 vector is comprised of model parameters. These are unknown quantities which (presumably) can be assigned a probability distribution, and hence a covariance matrix, based on expert knowledge. \mathbf{C}_{11} is therefore known. Suppose that the \mathbf{x}_2 vector represents model outputs that correspond to field measurements. These are obviously correlated with the elements of \mathbf{x}_1 because model outputs are sensitive to model parameters. The degree of correlation between each measurement-pertinent model output and each model parameter can be established empirically by undertaking many model runs based on samples of \mathbf{x}_1 . \mathbf{C}_{12} and \mathbf{C}_{21} can therefore be filled empirically, as can \mathbf{C}_{22} , based on the outcomes of these model runs. We can then fill the \mathbf{x}_2 vector of equation 3.8.4 with field measurements to calculate the post-history-matching expected values of model parameters; these are embodied in $\boldsymbol{\mu}'_1$. Meanwhile equation 3.8.5 can be used to compute the conditional (i.e. post-history-matching) uncertainties of model

parameters. This process will be discussed in much more detail later in this book.

3.9 Propagation of Covariance

Let \mathbf{x} be a random n -dimensional vector. Let \mathbf{A} be an $m \times n$ matrix. Let the m -dimensional vector \mathbf{y} be calculated from \mathbf{x} as follows

$$\mathbf{y} = \mathbf{Ax} \quad (3.9.1)$$

If \mathbf{x} is a random vector, than so too must \mathbf{y} be a random vector. Hence it possesses an $m \times m$ covariance matrix. It can be shown that

$$\mathbf{C}(\mathbf{y}) = \mathbf{AC}(\mathbf{x})\mathbf{A}^t \quad (3.9.2)$$

This relationship will be used extensively in work to follow.

3.10 A Covariance Matrix and Singular Value Decomposition

3.10.1 Kahunen-Loève Transformation

Suppose that the n -dimensional random vector \mathbf{x} has a covariance matrix $\mathbf{C}(\mathbf{x})$. For convenience (and without loss of generality), let us also suppose that the mean of \mathbf{x} is $\mathbf{0}$; elements of \mathbf{x} therefore represent perturbations from their means. As has already been discussed, $\mathbf{C}(\mathbf{x})$ is a positive-definite matrix. Hence, when subjected to singular value decomposition, an orthonormal \mathbf{E} matrix and a diagonal \mathbf{F} matrix can be found which satisfy the following equation.

$$\mathbf{C}(\mathbf{x}) = \mathbf{E}\mathbf{F}\mathbf{E}^t \quad (3.10.1)$$

Let the n -dimensional vector \mathbf{y} be calculated from \mathbf{x} using the equation

$$\mathbf{y} = \mathbf{F}^{1/2}\mathbf{E}^t\mathbf{x} \quad (3.10.2)$$

As was discussed above, the diagonal $\mathbf{F}^{1/2}$ matrix is calculated from the diagonal \mathbf{F} matrix by taking the square root of its respective elements. The diagonal $\mathbf{F}^{-1/2}$ matrix is calculated from the diagonal $\mathbf{F}^{1/2}$ matrix by individually inverting respective diagonal elements. Through use of equation (3.9.2) the covariance matrix of \mathbf{y} , that is $\mathbf{C}(\mathbf{y})$, can be calculated as

$$\mathbf{C}(\mathbf{y}) = \mathbf{F}^{-1/2}\mathbf{E}^t\mathbf{E}\mathbf{F}\mathbf{E}^t\mathbf{E}\mathbf{F}^{-1/2} = \mathbf{I} \quad (3.10.3)$$

where use is made of the fact that \mathbf{E} is an orthonormal matrix, so that

$$\mathbf{E}^t\mathbf{E} = \mathbf{I} \quad (3.10.4)$$

The covariance matrix of \mathbf{y} is thus the identity matrix.

Suppose that \mathbf{y} possesses a multiGaussian probability distribution. Then the smaller is the norm of \mathbf{y} , the higher is its likelihood. This is because the exponent in equation (3.6.1) becomes equal to minus half the square of the norm of a random vector when its covariance matrix is \mathbf{I} . An inversion methodology such as singular value decomposition which achieves a minimum norm solution to an inverse problem therefore achieves a solution of maximum likelihood to that problem if the prior covariance matrix of parameters is \mathbf{I} .

Suppose that an inverse problem is posed for \mathbf{x} , and that the prior covariance matrix of \mathbf{x} is $\mathbf{C}(\mathbf{x})$ and not \mathbf{I} . A solution of maximum likelihood for that problem can be achieved through singular value decomposition if it is first posed in terms of the transformed parameter set \mathbf{y} where the relationship between \mathbf{y} and \mathbf{x} is described by (3.10.2). The solution for \mathbf{x} can then be back-calculated from that achieved for \mathbf{y} using the relationship (easily derived from equation

3.10.2)

$$\mathbf{x} = \mathbf{E}\mathbf{F}^{1/2}\mathbf{y} \quad (3.10.5)$$

The transformation depicted in equation (3.10.2) is known as the Kahunen-Loëve transformation.

3.10.2 Principal Component Analysis

Now define the vector \mathbf{z} as

$$\mathbf{z} = \mathbf{E}^t \mathbf{x} \quad (3.10.6)$$

Because \mathbf{E}^t (like its inverse \mathbf{E}) is an orthonormal matrix, \mathbf{z} has the same magnitude as \mathbf{x} . It is rotated from \mathbf{x} to a new position in vector space through the rotation operation encapsulated in \mathbf{E}^t . \mathbf{z} is related to the vector \mathbf{y} discussed in the previous subsection (see equation 3.10.2) through the equation

$$\mathbf{y} = \mathbf{F}^{-1/2} \mathbf{z} \quad (3.10.7)$$

Recalling that $\mathbf{F}^{-1/2}$ is a diagonal matrix, \mathbf{z} points in the same direction as \mathbf{y} . The individual elements of \mathbf{z} are related to those of \mathbf{y} through an element-specific factor which is the corresponding diagonal element of $\mathbf{F}^{-1/2}$. From (3.9.2), (3.10.1) and (3.10.6) it readily follows that

$$\mathbf{C}(\mathbf{z}) = \mathbf{F} \quad (3.10.8)$$

As \mathbf{F} is a diagonal matrix, the elements of \mathbf{z} are not correlated with each other. This is despite the fact that the elements of \mathbf{x} are indeed correlated with each other if $\mathbf{C}(\mathbf{x})$ is nondiagonal.

Recall from section 2.14 that the columns of \mathbf{E} are the eigenvectors of $\mathbf{C}(\mathbf{x})$ and that the (diagonal) elements of \mathbf{F} are its eigenvalues. The columns of \mathbf{E} become the rows of \mathbf{E}^t . Each element of the vector \mathbf{z} which is calculated from \mathbf{x} using (3.10.6) is thus the scalar product of \mathbf{x} with an eigenvector of its covariance matrix $\mathbf{C}(\mathbf{x})$. From (3.10.8) the variance of each of these elements is the corresponding eigenvalue of $\mathbf{C}(\mathbf{x})$; the standard deviation of each of these elements is therefore the square root of this eigenvalue. Figure 3.6 depicts the probability density functions of the individual elements of a two-dimensional \mathbf{z} vector.

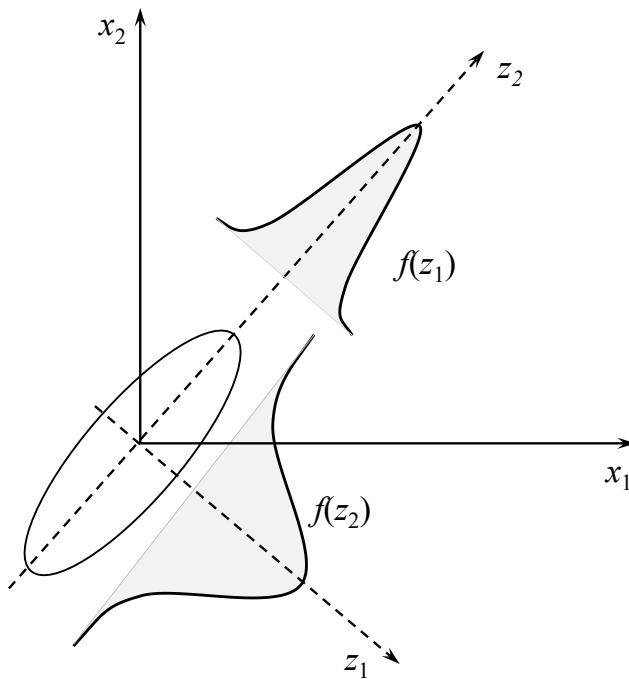


Figure 3.6. Independent probability density functions for components of \mathbf{z} obtained from \mathbf{x} using equation (3.10.6).

There is an important conclusion to be drawn from this analysis. Any random vector \mathbf{x} can be rotated to a place in vector space where its elements are mutually independent. Rotation is embodied in equation (3.10.6). Because it is the outcome of a matrix operation, any element of \mathbf{z} is a linear combination of elements of \mathbf{x} . The factors by which elements of \mathbf{x} are multiplied comprise elements of pertinent rows of the matrix \mathbf{E}^t .

In general terms, then, it follows that any set of n correlated random variables (i.e. elements of a general n -dimensional random vector) can be linearly combined in n different ways to form n independent random variables. The means of these independent random variables are calculated using the same linear combination of the means of the original variables. The standard deviations of these independent random variables are the square roots of singular values of the original covariance matrix. As has been shown in section 2.14, these are also the square roots of eigenvalues of the covariance matrix.

This is the basis for principal component analysis.

3.10.3 Inversion of an Empirical Covariance Matrix

Filling of an empirical covariance matrix is discussed in Section 3.5.2. There it is stated that occasions can arise where such a matrix requires inversion. (See, for example, the last paragraph of section 3.8; \mathbf{C}_{22} , which may be empirically evaluated, is inverted in equations 3.8.4 and 3.8.5.) However an empirical covariance matrix is singular if the number of realizations that are used in its construction is less than the number of random variables that it describes, and hence the number of rows and columns of the covariance matrix.

Let us suppose that we have filled an empirical covariance matrix \mathbf{x} to obtain $\mathbf{C}^e(\mathbf{x})$ which now requires inversion. This matrix can then be subjected to singular value decomposition; see equation (3.10.1). If the number of realizations exceeds the number of elements of \mathbf{x} , then $\mathbf{C}^{e-1}(\mathbf{x})$ can be calculated as

$$C^{e-1}(\mathbf{x}) = \mathbf{E}\mathbf{F}^{-1}\mathbf{E}^t \quad (3.10.9)$$

Pre- or postmultiplication of (3.10.1) by (3.10.9) yields the identity matrix \mathbf{I} , verifying that $C^{e-1}(\mathbf{x})$ of equation (3.10.9) is indeed the inverse of $C^e(\mathbf{x})$. Caution must be exercised in calculating this inverse, however, especially if the number of realizations does not exceed the size of \mathbf{x} by much. $C^e(\mathbf{x})$ is an approximation of the real $C(\mathbf{x})$; the same applies to $C^{e-1}(\mathbf{x})$. Furthermore, with only a small number of samples, $C^e(\mathbf{x})$ may approach singularity. In this case, errors in empirical covariance calculation may be amplified through the inversion process as reciprocals are taken of elements of \mathbf{F} . It may therefore be a good idea to replace entries of \mathbf{F}^{-1} which are particularly high (because elements of \mathbf{F} are particularly low) with values of zero before undertaking matrix multiplications embodied in equation (3.10.9). Because singular value decomposition arranges singular values in order of decreasing value in the diagonal \mathbf{F} matrix, singular values which may require this treatment occupy the trailing elements of \mathbf{F} .

A similar strategy should be followed when the number of realizations is less than the number of random variables comprising \mathbf{x} . In this case, singular value decomposition of $C^e(\mathbf{x})$ yields zero-valued trailing diagonal elements of \mathbf{F} . Corresponding elements of \mathbf{F}^{-1} should also be set to 0, together with particularly large elements of \mathbf{F}^{-1} . The inverted covariance matrix that emerges from this procedure is still rank deficient. However it may still be useable for calculations that require it.

3.10.4 Square Root of a Covariance Matrix

Using singular value decomposition, it is an easy matter to calculate the square root of a covariance matrix. From (3.10.1)

$$C^{1/2}(\mathbf{x}) = \mathbf{E}\mathbf{F}^{1/2}\mathbf{E}^t \quad (3.10.10)$$

The correctness of (3.10.10) is easily verified by multiplying $C^{1/2}(\mathbf{x})$ by itself.

$$C^{1/2}(\mathbf{x}) C^{1/2}(\mathbf{x}) = \mathbf{E}\mathbf{F}^{1/2}\mathbf{E}^t \mathbf{E}\mathbf{F}^{1/2}\mathbf{E}^t = \mathbf{E}\mathbf{F}^{1/2}\mathbf{F}^{1/2}\mathbf{E}^t = \mathbf{E}\mathbf{F}\mathbf{E}^t = C(\mathbf{x}) \quad (3.10.11)$$

$\mathbf{F}^{1/2}$ is a diagonal matrix whose elements are the square roots of those of \mathbf{F} .

3.11 Spatial Correlation

In many contexts of PEST usage, expert knowledge of parameters as it exists prior to history-matching can be expressed through positing spatial correlation between them. This is especially the case in groundwater and subsurface reservoir modelling. Spatial correlation of parameters follows from the likelihood that hydraulic properties at points which are close together are similar. However the likelihood of hydraulic property similarity is likely to decay with increasing separation of the points until, at some separation, there is likely to be no relationship at all between properties at the two points.

Suppose that a model employs m parameters. Keeping in mind that a covariance matrix is symmetric, the number of values that must be provided for the filling of a $m \times m$ covariance matrix is $m(m+1)/2$. This rises rapidly as m rises.

In contrast to this, only a few numbers are required to fill a spatial covariance matrix if the simplifying assumption of geostatistical stationarity is made. This assumption implies that system property correlation is a function only of the distance between points at which the property is specified, and not of the locations of these points. Spatial correlation is thus designated as $\rho(h)$ where h is the distance between any two points for which system property correlation must be calculated. If the variance of the system property in question is similarly uniform, then a covariance function $C(h)$ becomes immediately available. If required, such a

function can readily accommodate system property anisotropy whereby spatial correlation can decay at different rates in different orthogonal directions.

As has already been mentioned, a covariance matrix must be positive-definite. The reason for this is clear from an inspection of the exponent of equation (3.6.1). The negative of this exponent cannot become positive, for probabilities must fade as the values of random variables become very large; they cannot grow exponentially. The filling of any covariance matrix must be such as to achieve this fundamental design specification.

Certain functions which relate covariance to distance are guaranteed to achieve positive-definiteness of a spatial covariance matrix. One of these is the exponential decay function through which spatial covariance is expressed as

$$C(h) = C(0)e^{-h/a} \quad (3.11.1)$$

In equation (3.11.1), $C(0)$ is the variance of the system property under consideration. This value occupies the diagonal elements of the covariance matrix which pertains to this property; as such, it specifies the “innate variability” of this property as assessed by expert knowledge. a of equation (3.11.1) determines the rate of decay of spatial correlation with distance; the larger is a the smaller is the rate of this decay.

Spatial covariance matrices are often derived from semivariograms (often referred to as “variograms” for short). The semivariogram as a means of specifying spatial correlation has its roots in the geostatistical literature. Let us designate the value of a system property at location \mathbf{r} as $k(\mathbf{r})$. Suppose that this property has been measured at two points which are a distance h apart. The value of the semivariogram for this separation, i.e. $\gamma(h)$, is defined to be half the expected value of the squared difference between property values. That is

$$\gamma(h) = \frac{1}{2}E[(k(\mathbf{r}) - k(\mathbf{r}+\mathbf{h}))^2] \quad (3.11.2)$$

where the vector \mathbf{h} is of length h . System property anisotropy is taken into account by specifying different dependencies of γ on h in different orthogonal directions.

It is easily shown that

$$C(h) = C(0) - \gamma(h) \quad (3.11.3)$$

and that

$$\gamma(\infty) = C(0) \quad (3.11.4)$$

Figure 3.7 depicts a semivariogram and corresponding spatial covariance function. It also depicts commonly-employed descriptors of semivariogram characteristics, namely the semivariogram sill and the semivariogram range. The sill is the semivariogram value at large h which, from (3.11.4), is equal to the variance of the system property. The semivariogram range is the distance over which spatial correlation of the system property prevails.

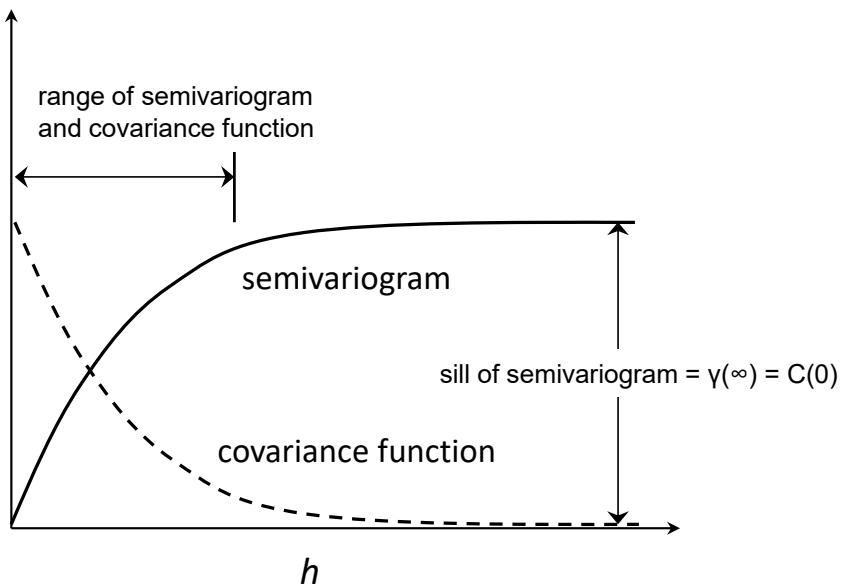


Figure 3.7. Relationship between a semivariogram and a spatial covariance function.

As already stated, mathematical functions that are used to characterize semivariograms must ensure positive-definiteness of associated covariance matrices. Four commonly-used semivariogram functions which achieve this purpose are provided below.

Spherical

$$\gamma(h) = c \cdot \left[1.5 \frac{h}{a} - 0.5 \left(\frac{h}{a} \right)^3 \right] \quad \text{if } h < a \quad (3.11.5a)$$

$$\gamma(h) = c \quad \text{if } h \geq a \quad (3.11.5b)$$

Exponential

$$\gamma(h) = c \cdot \left[1 - \exp \left(-\frac{h}{a} \right) \right] \quad (3.11.6)$$

Gaussian

$$\gamma(h) = c \cdot \left[1 - \exp \left(-\frac{h^2}{a^2} \right) \right] \quad (3.11.7)$$

Power

$$\gamma(h) = c \cdot h^a \quad (3.11.8)$$

3.12 Nonstationarity

The concept of stationarity is appealing as it simplifies construction of a covariance matrix for spatial parameterization devices such as pilot points. (These are discussed in chapter 5.). At the same time it is a little unrealistic, especially in model domains of large areal extent in which geology is likely to vary. It is also unrealistic in smaller areas whose geology may be characterized by sinuously anisotropic features such as buried alluvial channels, or is

permeated by local structural features with high degrees of connected permeability in one direction but not in the orthogonal direction. The importance of permeability connectedness in determining groundwater model outcomes is addressed by a number of authors, including Renard and Allard (2013) and Renard and Ababou (2022).

There is no reason why covariance matrices cannot embody nonstationary variograms. This allows a modeller to express the fact that the mean value of a hydraulic property, and the variability of that hydraulic property about its mean, can vary from place to place throughout a model domain. So too can the direction and magnitude of anisotropy of spatial correlation of that property. What is important, however, is that the covariance matrix remain positive definite.

Programs from the PEST support suite such as PLPROC, PPCOV_SVA and PPCOV3D_SVA allow a modeller to express spatial variability of semivariogram specifications; see Doherty, 2025a; 2025b. The spatial correlation between any two points is taken as the minimum of those attributed to the points. A temporary covariance matrix is then constructed using inter-point correlations and point-based variances. Unfortunately, a temporary covariance matrix that is built in this way is not always positive definite. This is easily rectified by first subjecting it to singular value decomposition such that

$$C_{\text{temp}}(\mathbf{k}) = \mathbf{E}\mathbf{F}\mathbf{G}^t \quad (3.12.1)$$

Parameters are represented using the symbol \mathbf{k} in equation (3.12.1). For $C_{\text{temp}}(\mathbf{k})$ matrices built in this way, the columns of \mathbf{E} and the rows of \mathbf{G} are the same except for the possibility of sign reversal. Where this exists, it is corrected, so that

$$C(\mathbf{k}) = \mathbf{E}\mathbf{F}\mathbf{E}^t \quad (3.12.2)$$

The covariance matrix is then reconstructed, as it is now positive definite. While the methodology is somewhat heuristic, it works well.

3.13 Random Number Generation

All computer programming languages provide a function for generation of random variables that possess a uniform distribution over the interval $(0,1)$; this probability distribution is symbolized using the notation $U(0,1)$. Generally, in groundwater modelling, we deal with variables that possess Gaussian or multiGaussian probability distributions. (As we shall see, the integrity of popular data assimilation and uncertainty quantification methods that work in high dimensional parameter spaces rests on an assumption of multiGaussian behaviour.)

Suppose that we have drawn two random numbers u_1 and u_2 from $U(0,1)$. The so-called ‘‘Box-Muller’’ transform can then be employed to generate two independent random numbers z_1 and z_2 belonging to $N(0,1)$ (i.e. a normal distribution with a mean of 0 and variance of 1). These are calculated from u_1 and u_2 using the following formulas.

$$z_1 = \sqrt{-\ln(u_1)} \cos(2\pi u_2) \quad (3.13.1a)$$

$$z_2 = \sqrt{-\ln(u_2)} \sin(2\pi u_1) \quad (3.13.1b)$$

For z_1 and z_2 to sample $N(\mu, \sigma^2)$ instead of $N(0,1)$, multiply them by σ and add μ .

Suppose now that we would like to generate a random vector \mathbf{w} that belongs to the multiGaussian distribution $N(\mathbf{0}, C(\mathbf{k}))$. Suppose that the vector \mathbf{w} has m elements. First, we generate realizations of m independent, standard normal variates and place them into a vector \mathbf{z} . Because each element of \mathbf{z} is sampled from $N(0,1)$, the covariance matrix of \mathbf{z} is the identity

matrix. Hence

$$\mathbf{z} \sim N(\mathbf{0}, \mathbf{I}) \quad (3.13.2)$$

Next we subject $C(\mathbf{k})$ to singular value decomposition so that

$$C(\mathbf{k}) = \mathbf{E}\mathbf{F}\mathbf{E}^t \quad (3.13.3)$$

Then \mathbf{w} can be calculated from \mathbf{z} using the equation

$$\mathbf{w} = \mathbf{E}\mathbf{F}^{1/2}\mathbf{z} \quad (3.13.4)$$

The vector \mathbf{w} has the correct covariance matrix because, from (3.9.2)

$$C(\mathbf{w}) = \mathbf{E}\mathbf{F}^{1/2}C(\mathbf{z})\mathbf{F}^{1/2}\mathbf{E}^t = \mathbf{E}\mathbf{F}^{1/2}\mathbf{I}\mathbf{F}^{1/2}\mathbf{E}^t = \mathbf{E}\mathbf{F}\mathbf{E}^t = C(\mathbf{k}) \quad (3.13.5)$$

There is another option as well. Rather than using $\mathbf{E}\mathbf{F}^{1/2}$ to transform \mathbf{z} , we could use $C^{1/2}(\mathbf{k})$. The square root of a covariance matrix is defined by equation (3.10.10). Therefore let

$$\mathbf{w} = C^{1/2}(\mathbf{k})\mathbf{z} \quad (3.13.6)$$

It is an easy matter to verify that \mathbf{w} calculated in this way has the correct covariance matrix. From (3.9.2)

$$C(\mathbf{w}) = C^{1/2}(\mathbf{k})C(\mathbf{z})C^{1/2}(\mathbf{k}) = C^{1/2}(\mathbf{k})\mathbf{I}C^{1/2}(\mathbf{k}) = C^{1/2}(\mathbf{k})C^{1/2}(\mathbf{k}) = C(\mathbf{k}) \quad (3.13.7)$$

An advantage of using (3.13.6) instead of (3.13.4) is that variates comprising the elements of \mathbf{z} are more closely related to variates comprising the elements \mathbf{k} when using (3.13.6) than when using (3.13.4). For example, if model parameters \mathbf{k} are associated with pilot points, then the elements of \mathbf{z} can be ascribed to these same points, even though they are independent normal variates instead of spatially correlated hydraulic property values. Their effects on local hydraulic properties \mathbf{k} can then be visualized; non-zero values of \mathbf{z} result in local variability of \mathbf{k} about its mean. (In the present case we assume that \mathbf{k} has a mean of zero. However its mean could just as easily be any non-zero number; it could also be spatially variable.)

So, for both of these options, the filling of the vector \mathbf{w} of correlated, multiGaussian random variables is undertaken by first generating values for a set of independent random variates whose means are zero and whose standard deviations are 1. These are often referred to as iids (for “independent and identically distributed” random variates). Generation of random parameter fields for even complex geostatistical models follows a similar trajectory. Iids are first sampled and then subjected to case-specific transformation and manipulation.

3.14 Mahalanobis Distance

Later in this book the issue of prior-data conflict is discussed. This kind of analysis attempts to answer questions such as “Does the random vector \mathbf{z} belong to a probability distribution that is characterized by mean $\underline{\mathbf{x}}$ and covariance matrix $C(\mathbf{x})$?” Let us assume a multiGaussian probability distribution for \mathbf{x} .

An answer to the above question can be sought using the concept of Mahalanobis distance. The Mahalanobis distance from an arbitrary random vector \mathbf{x} to its mean $\underline{\mathbf{x}}$ is given by

$$d_{mx} = (\mathbf{x} - \underline{\mathbf{x}})^t C^{-1}(\mathbf{x})(\mathbf{x} - \underline{\mathbf{x}}) \quad (3.14.1)$$

Notice the similarity of this definition to the term that appears in the exponent of the multiGaussian probability distribution; see equation 3.6.1.

Equation (3.14.1) can be re-written as

$$d_{mx} = (\mathbf{x} - \underline{\mathbf{x}})^t \mathbf{C}^{-1/2}(\mathbf{x}) \mathbf{C}^{-1/2}(\mathbf{x}) (\mathbf{x} - \underline{\mathbf{x}}) \quad (3.14.2)$$

That is

$$d_{mx} = \mathbf{r}^t \mathbf{r} \quad (3.14.3a)$$

where

$$\mathbf{r} = \mathbf{C}^{-1/2}(\mathbf{x}) (\mathbf{x} - \underline{\mathbf{x}}) \quad (3.14.3b)$$

From (3.9.2) the covariance matrix of \mathbf{r} can be written as

$$\mathbf{C}(\mathbf{r}) = \mathbf{C}^{-1/2}(\mathbf{x}) \mathbf{C}(\mathbf{x}) \mathbf{C}^{-1/2}(\mathbf{x}) = \mathbf{I} \quad (3.14.4)$$

The individual elements of \mathbf{r} therefore comprise independent standard normal variates. The dot product $\mathbf{r}^t \mathbf{r}$ is thus the sum of squared standard normal variates. Hence d_{mx} is a stochastic variable that belongs to a chi squared distribution of order n , where n is the number of elements in the vector \mathbf{x} . That is

$$d_{mx} \sim \chi_n^2 \quad (3.14.5)$$

It follows that to inquire whether an arbitrary random vector \mathbf{z} could be a sample of $N(\underline{\mathbf{x}}, \mathbf{C}(\mathbf{x}))$ we should calculate d_{mz} where

$$d_{mz} = (\mathbf{z} - \underline{\mathbf{x}})^t \mathbf{C}^{-1}(\mathbf{x}) (\mathbf{z} - \underline{\mathbf{x}}) \quad (3.14.6)$$

The value of d_{mz} should then be compared with upper and lower confidence limits of χ_n^2 in order to associate a level of confidence with the assertion that \mathbf{z} samples $N(\underline{\mathbf{x}}, \mathbf{C}(\mathbf{x}))$. Recall from section 3.2.3 that the mean of chi squared distribution of order n is n , while its variance is $2n$.

3.15 A Much Tougher Test

We now ask the same question. Does the random vector \mathbf{z} sample a probability distribution that is characterized by mean $\underline{\mathbf{x}}$ and covariance matrix $\mathbf{C}(\mathbf{x})$?" Let us again assume multiGaussianity of \mathbf{x} . Also, to simplify equations that follow (and without losing generality), we also assume that the mean of \mathbf{x} is zero.

As discussed above, to explore this issue using the concept of Mahalanobis distance, we could compute $\mathbf{z}^t \mathbf{C}^{-1}(\mathbf{x}) \mathbf{z}$ and then refer to a chi-squared distribution table to see where this computed value lies with respect to the range of values that a chi-squared distribution with n degrees of freedom should possess. Suppose, for example, that n is 10. After referring to a chi-squared distribution table we find that $\chi^2(10, 0.975)$ is 3.247 while $\chi^2(10, 0.025)$ is 20.483. This means that there is a 97.5% chance that a chi-squared variate with 10 degrees of freedom is above 3.247 but only a 2.5% chance that it is above 20.483. In other words, there is a 95% chance that it is between 3.247 and 20.483. Hence if the value that we compute for $\mathbf{z}^t \mathbf{C}^{-1}(\mathbf{x}) \mathbf{z}$ is outside this interval, then the hypotheses that \mathbf{z} is multiGaussian with mean $\mathbf{0}$ and covariance matrix $\mathbf{C}(\mathbf{x})$ can be rejected with a 95% level of confidence.

Actually, this is not a very strong test that $\mathbf{z} \sim N(\mathbf{0}, \mathbf{C}(\mathbf{x}))$. This can be demonstrated by example. Suppose that $\mathbf{C}(\mathbf{x})$ is \mathbf{I} . Suppose further that all of the elements of our test vector \mathbf{z} are zero except for one of them, and that this element is 4. $\mathbf{z}^t \mathbf{C}^{-1}(\mathbf{x}) \mathbf{z}$ is therefore 16. Based on the above chi-squared test, we cannot reject the hypothesis at the 95% confidence level that $\mathbf{z} \sim N(\mathbf{0}, \mathbf{I})$. However reference to a normal distribution table quickly informs us that the chances of an individual standard normal variate being 4 or above are about 0.003%.

We will now formulate a stronger test for whether $\mathbf{z} \sim N(\mathbf{0}, \mathbf{C}(\mathbf{x}))$. First we present a little theory.

Consider the scalar product between a known vector \mathbf{m} and a random vector \mathbf{z} . We continue to assume a mean of $\mathbf{0}$ for \mathbf{x} . Let us designate this scalar product as s . Then

$$s = \mathbf{m}^t \mathbf{z} \quad (3.14.7)$$

If $C(\mathbf{z})$ is the same as $C(\mathbf{x})$ then, applying the usual formula for propagation of covariance

$$\sigma_s^2 = \mathbf{m}^t C(\mathbf{x}) \mathbf{m} \quad (3.14.8)$$

Reasonableness of s calculated using (3.14.7) can be judged using σ_s^2 calculated from (3.14.8). If s lies outside its expected confidence region, then $C(\mathbf{x})$ must be rejected as representing $C(\mathbf{z})$ of which \mathbf{z} is supposedly a realization. That is, we can reject the notion that $\mathbf{z} \sim N(\mathbf{0}, C(\mathbf{x}))$.

This test can be made more convenient if the vector \mathbf{m} is normalized through division by $[\mathbf{m}^t C(\mathbf{x}) \mathbf{m}]^{1/2}$. The scalar s that we use as the basis for our test then becomes

$$s = \frac{\mathbf{m}^t \mathbf{z}}{[\mathbf{m}^t C(\mathbf{x}) \mathbf{m}]^{1/2}} \quad (3.14.9)$$

for which the variance, if $C(\mathbf{z})$ is the same as $C(\mathbf{x})$, is

$$\sigma_s^2 = \frac{\mathbf{m}^t}{[\mathbf{m}^t C(\mathbf{x}) \mathbf{m}]^{1/2}} C(\mathbf{x}) \frac{\mathbf{m}}{[\mathbf{m}^t C(\mathbf{x}) \mathbf{m}]^{1/2}} = 1 \quad (3.14.10)$$

It follows that if \mathbf{z} has a multiGaussian distribution with a mean of $\mathbf{0}$ a covariance matrix of $C(\mathbf{x})$, then the scalar $\frac{\mathbf{m}^t \mathbf{z}}{[\mathbf{m}^t C(\mathbf{x}) \mathbf{m}]^{1/2}}$ has a normal distribution with a mean of 0 and a standard deviation of 1.

To see why this test works, and to see why it is so tough, we demonstrate it for a particularly simple case. Suppose that \mathbf{z} and \mathbf{x} contain two elements. As before, we assume a mean of zero for each of these elements. They can therefore be represented by their coordinates in a two-dimensional plain. We further assume that $C(\mathbf{x}) = \mathbf{I}$. Under these conditions $\frac{\mathbf{m}}{[\mathbf{m}^t C(\mathbf{x}) \mathbf{m}]^{1/2}}$ is a unit vector that points in an arbitrary direction in this two-dimensional vector space; the direction depends on our choice of \mathbf{m} . Meanwhile $\frac{\mathbf{m}^t \mathbf{z}}{[\mathbf{m}^t C(\mathbf{x}) \mathbf{m}]^{1/2}}$ is the scalar projection of \mathbf{z} (the vector that we wish to test) onto this arbitrary unit vector. If $\mathbf{z} \sim N(\mathbf{0}, \mathbf{I})$, this scalar projection is a standard normal variate. This is clear from Figure 3.8. Regardless of the direction of a line through the origin, the projection of a 2D random vector \mathbf{z} for which $\mathbf{z} \sim N(\mathbf{0}, \mathbf{I})$ onto this line is a standard normal variate.

This strict statistical test formalizes the notion of statistical conformity in all directions of vector space. In doing so, it ensures not just that the magnitude of the tested vector is correct (as does Mahalanobis distance). It ensures that all elements of the vector adhere to their individual standard deviations. It also ensures that all linear combinations of its elements adhere to their respective one-dimensional probability distributions.

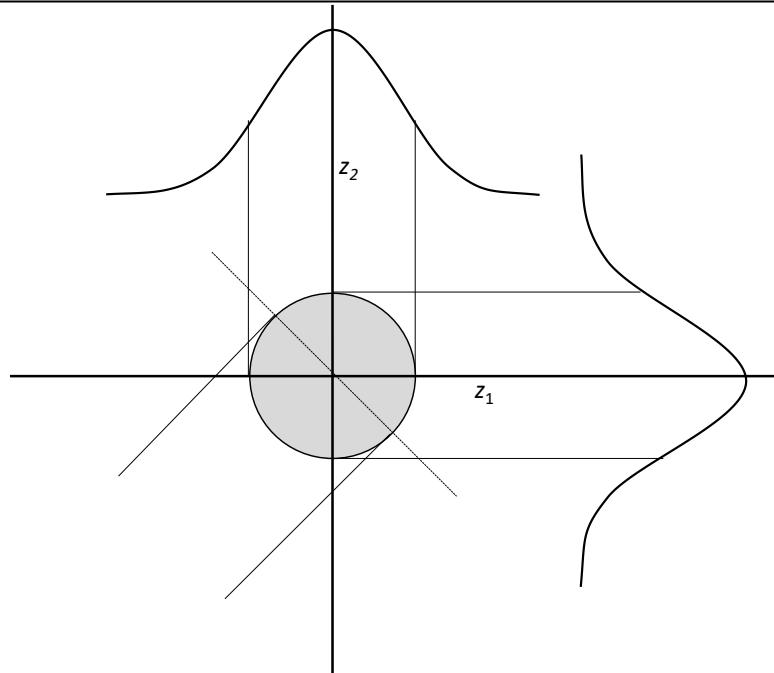


Figure 3.8. The circle has a radius of 1 standard deviation. The projection onto any line of a random 2D vector whose probability distribution is multiGaussian and for which individual elements are statistically independent is a standard normal variate.

4. Hydraulic Property Field Generation

4.1 Background

In this short chapter we pay a brief visit to a very large topic. It is that of geostatistical hydraulic property field generation. This refers to methods whereby a large two- or three-dimensional grid (possibly a model grid) is repeatedly populated with numbers that respect user-specified geostatistical and other relationships. These numbers are therefore random numbers. However they embody “randomness with rules”. These rules may be as simple as adherence to a semivariogram that prevails throughout a model domain under the assumption of geostatistical stationarity. Alternatively, grid-populated hydraulic property fields may be the outcome of detailed numerical simulation of depositional, structural and other processes that have given rise to the complex disposition of subsurface hydraulic properties at a study site; in this case strategic randomness is embodied in the simulation process in order for some aspects of the hydraulic property fields to differ from realization to realization. Or perhaps the hydraulic property field generation process may be something in between these two extremes whereby the locations and dispositions of “geobodies” whose shapes mimic those of idealized geological entities are randomly generated, and randomly emplaced, within a model domain.

Another popular methodology for geostatistical field generation is multiple point geostatistics. So-called “two-point” methods rely on a semivariogram; the correlation between hydraulic properties at two different locations is a function only of the separation between these two locations. In contrast, multiple point geostatistical methods attempt to introduce more complex patterns into emergent hydraulic property fields that cannot be described by separation-based correlation alone. These patterns are discovered by scanning a two- or three-dimensional image of the subsurface. These images may have been generated using process simulators, geobody manipulators or by other means. See, for example, Tahmasebi (2018) and Mariethoz and Caers (2015) for more information.

Random geostatistical fields may be comprised of real numbers or a small number of integers which represent different geological facies or other geological items. Fields of this type are referred to as “categorical”. (For convenience we will refer to stochastic fields that are comprised of real numbers as “continuous”.) The disposition of these integers may vary greatly from realization to realization. When supplied to a model, integers are generally replaced by the hydraulic properties of the geological entities that they represent. Alternatively, they may be populated with continuous stochastic fields whose means and semivariograms are different in different geological units.

Software for generation of continuous and categorical stochastic fields using a variety of methodologies is freely available through many stand-alone packages and Python libraries. See, for example, Deutsch and Journel (1997), Carle (1999) and Remy et al (2011).

As the focus of this book is on parameter estimation and uncertainty reduction accrued through history-matching, little attention is given to the theory and mechanics of geostatistical field generation. However we do describe one methodology that is of particular use in history-match-constrained uncertainty analysis.

Though useful (and impressive) for exploration of pre-calibration uncertainty, many methodologies of geostatistical field generation are of limited use in contexts where decision-support modelling is charged with the responsibility of harvesting information from historical measurements of system behaviour in order to reduce model parameter and predictive

uncertainty. The latter tasks require use of “parameters”. As will be discussed in the next chapter, parameters (by definition) are adjustable. This allows them to host information that constrains uncertainty, while retaining the ability to express residual uncertainty. Generation of complex, geostatistically-based hydraulic property fields is rarely done with history-match adjustment in mind.

We note, however, that use by a numerical simulator of complex geostatistical hydraulic property fields does not necessarily preclude predictive uncertainty reduction through history-matching. Later in this book we discuss data space inversion (DSI). Using only a few hundred model runs based on geostatistical hydraulic property fields of arbitrary complexity, DSI builds statistical relationships between the measured past and the managed future. These relationships are encapsulated in a statistical model whose parameters are indeed continuously adjustable. History-matching of this statistical model allows quantification and reduction of predictive uncertainty. Meanwhile, a modeller is free to endow his/her numerical model with continuous or categorical, stochastic, hydraulic property fields of arbitrary complexity if he/she feels that the integrity of decision-support numerical simulation will suffer if this is not done.

We also note that generation of geostatistical hydraulic property fields need not be a difficult undertaking if a modeller is content to represent hydraulic property heterogeneity in simple, but often effective, ways. In the next chapter we discuss pilot points. In the previous chapter we discussed covariance matrices, including covariance matrices that are based on nonstationary geostatistics. A model can be endowed with a stochastic hydraulic property field by ascribing random numbers to a set of pilot point parameters based on a covariance matrix. As is discussed in the next chapter, these hydraulic properties are then spatially interpolated to a model grid. A disadvantage of this scheme is that stochastic variability is not represented at the model cell scale; instead, it is represented at the pilot point scale (which is often sufficient). An advantage of this scheme is that random parameters that are assigned to pilot points can be easily adjusted in order for model outputs to respect history-matching constraints. This forms the basis for ensemble methods that are discussed later in this book.

We close this subsection with a philosophical observation.

Many modellers believe that a groundwater model cannot be useful unless it embodies representation of hydrogeological detail that is commensurate with their mental image of the subsurface. The author of this book contests this notion. The necessity for representation of detail in a groundwater model depends on a number of things. One of these is the sensitivity of critical model outputs (of past or future system behaviour) to this detail. Another is the extent to which predictions of future system behaviour are informed by its past behaviour. There are subtleties at play here.

Where the future behaviour of a managed system is completely informed by its behaviour under preceding conditions, predictions of future behaviour can be considered to be “data driven”. In this case, all that is required of a groundwater model is that it be a machine learning tool. Parametric detail is required only to the extent that model parameters can host prediction-pertinent information that is harvested from measurements of past system behaviour. See chapter 15 of this book for further discussion of this issue.

In contrast, where predictions are not data-driven, representation of prediction-sensitive hydraulic property detail is required for integrity of predictive uncertainty analysis. However, the details of hydraulic property detail are rarely known. It follows that representation of detail must be stochastic. Additionally, the level of detail that requires representation in a groundwater model depends on the sensitivity of management-salient model predictions to the

details of hydraulic property detail; there may therefore be sensible upper limits on the level of hydraulic property detail that requires representation in a numerical model.

In closing this subsection, we note that it may be far more difficult to derive a useable stochastic description of lithological and structural detail than to conjure up mental images of that detail. Stochastic characterization of prediction-sensitive hydraulic property detail is one of the many challenges that decision-support groundwater modelling faces.

4.2 Geostatistical Field Generation through Spatial Averaging

4.2.1 Advantages

We consider only one form of cell-by-cell stochastic hydraulic property field generation in this book. A modification of this methodology, adapted for history-matching and history-match-constrained uncertainty analysis is discussed in the next chapter.

Stochastic field generation is often referred to as “simulation” in the geostatistical literature. This is to distinguish it from other geostatistical activities such as spatial interpolation as implemented through kriging. In this text, however, we reserve the term “simulation” for numerical simulation of environmental processes such as subsurface fluid flow.

Spatial averaging as a means of geostatistical field generation is described in this book for the following reasons.

1. It is easily implemented. (It is implemented by PLPROC and by a number of programs that belong to the PEST Groundwater Utility Suite. It is also implemented by PyEMU.)
2. It can be used to populate both structured and unstructured model grids.
3. While based on cell-by-cell parameterization, it is easily modified for use with pilot points.
4. When implemented using pilot points, random parameter values that are ascribed to pilot points can be adjusted so that model outputs match field observations. Its use is therefore compatible with both model calibration and model predictive uncertainty analysis.
5. It can readily accommodate “hierarchical inversion”, in which hydraulic property parameters can be estimated and adjusted simultaneously with geostatistical hyperparameters that govern spatial patterns adopted by hydraulic properties that are employed by a model.

4.2.2 Spatial Averaging Kernel

Section 3.13 discusses generation of values of correlated random variables. One way of doing this is expressed by equation (3.13.6) which is repeated here as equation (4.2.1) with \mathbf{k}' substituted for \mathbf{w} in the previous equation.

$$\mathbf{k}' = \mathbf{C}^{1/2}(\mathbf{k})\mathbf{z} \quad (4.2.1)$$

In equation (4.2.1), the vector \mathbf{k}' is a realization of the set of random variables of type \mathbf{k} while $\mathbf{C}(\mathbf{k})$ is the covariance matrix of \mathbf{k} . The vector \mathbf{z} contains independent standard normal deviates (iids). These all have means of zero and standard deviations of one. There are as many elements in \mathbf{z} as there are in \mathbf{k} .

Equation (4.2.1) presumes that \mathbf{k} has a mean value of zero. Suppose, instead, that \mathbf{k} has a mean value of $\underline{\mathbf{k}}$. Then equation (4.2.1) becomes

$$\mathbf{k}' = \underline{\mathbf{k}} + \mathbf{C}^{1/2}(\mathbf{k})\mathbf{z} \quad (4.2.2)$$

Now recall from Chapter 2 how multiplication of a vector by a matrix works. To calculate the i 'th element of \mathbf{k}' , i.e. k'_i , one must travel along the i 'th row of matrix $C^{\frac{1}{2}}(\mathbf{k})$ while multiplying corresponding elements of \mathbf{z} . That is

$$k'_i = \underline{k}_i + \sum_{j=1}^m C_{ij}^{\frac{1}{2}}(\mathbf{k}) z_j \quad (4.2.3)$$

where m is the number of elements in both of the \mathbf{k} and \mathbf{z} vectors. Let us suppose that \mathbf{k} represents a spatially distributed hydraulic property, and that each element of \mathbf{k} corresponds to an individual model cell (or to a cell within an individual model layer if stochastic field generation is restricted to just a single model layer). The elements of \mathbf{z} can be considered to correspond to these same grid cells. Hence elements of \mathbf{z} and \mathbf{k} coincide. In the limit, where cells are very small, the summation of equation (4.2.3) becomes a spatial integral over the model domain D . Equation (4.2.3) can then be written as

$$k'(\mathbf{x}) = \underline{k}(\mathbf{x}) + \int_D f(\mathbf{x} - \mathbf{s}) z(\mathbf{s}) d\mathbf{s} \quad (4.2.4)$$

In equation (4.2.4) $k'(\mathbf{x})$ signifies the value of k' at location \mathbf{x} while $z(\mathbf{s})$ signifies the value of z at location \mathbf{s} . Where geostatistical stationarity prevails, the kernel function $f(\mathbf{x})$ is invariant in space; Oliver (1995) refers to this function as the “square root kernel” (for obvious reasons). Equation (4.2.4) then becomes a convolution integral.

Where the domain of a model is infinite, the integration limits in (4.2.4) become infinite. Under these conditions (and the continued assumption of geostatistical stationarity), it is possible to derive expressions for $f(\mathbf{x})$ that produce stochastic fields with known spatial correlation decay functions. Oliver (1995) tabulates some of these functions for spatial integration over two and three dimensions. For example, if a two-dimensional \mathbf{k} field exhibits an exponential correlation decay function described by the equation

$$c(r) = e^{-\frac{r}{a}} \quad (4.2.5)$$

then the square root kernel is

$$f(r) = K_{1/4} \left(\frac{r}{a} \right) (2\pi a r)^{-1/4} \Gamma \left(\frac{3}{4} \right)^{-1} \quad (4.2.6)$$

In these equations, r is radial distance, $K_{1/4}$ is a modified Bessel function of order $\frac{1}{4}$ and Γ is the gamma function. For the Gaussian correlation decay function

$$c(r) = e^{-\frac{r^2}{a^2}} \quad (4.2.7)$$

the square root kernel is

$$f(r) = \left(\frac{4}{a^2 \pi} \right)^{\frac{1}{2}} e^{-\frac{2r^2}{a^2}} \quad (4.2.8)$$

which is also Gaussian. Note how, for the Gaussian decay function, the correlation length of the stochastic field is $\sqrt{2}$ times larger than that of the square root kernel function. (The two functions are directly comparable in this case because both of them employ a squared exponential decay term.) The correlation function therefore has a larger lateral footprint than that of the kernel function. Oliver (1995) presents square root kernel functions pertaining to a number of other spatial correlation functions in two and three dimensions.

It is important to note that, when implementing stochastic field generation using the spatial averaging methodology, it is not essential for a modeller to employ a square root kernel function that yields a stochastic field with a known correlation decay function. The use of any

decaying kernel function yields a stochastic field that embodies decaying spatial correlation. In the end, a modeller must decide whether stochastic fields that are generated using a particular kernel function look realistic. Whether spatial correlations that they embody can be described by an analytical function is a matter of secondary importance.

4.2.3 Nonstationary Spatial Correlation

If the square root covariance matrix of equation (4.2.2) or the square root kernel function of equation (4.2.4) varies in space, then the random field that is generated from the summation or spatial averaging process is also nonstationary. Stochastic field generation using the spatial averaging methodology is therefore readily amenable to generation of nonstationary stochastic fields. See, for example, Higdon et al (1999), Fuentes (2002) and Opazo (2025) for a further discussion of these concepts. See Kitlasten et al (2025) for a real-world example.

As mentioned above, nonstationary stochastic field generation requires that a modeller employ a spatially-decaying square root kernel (we also refer to this as a “spatial averaging kernel”) whose decay rate varies in space. That is, the decay rate of the kernel (which may be anisotropic) is a function of spatial coordinates. Other geostatistical hyperparameters (such as amplitude of the spatial averaging kernel, as well as its anisotropy direction and ratio) can also be functions of spatial coordinates. Presumably the averaging kernel is described by an analytical function so that it is amenable to easy evaluation as spatial averaging proceeds over a large model grid. As stated above, choice of this function is a matter for the modeller; there is no “right” and “wrong” choice. All that matters is that stochastic hydraulic property fields that emerge from the averaging process look “realistic”.

4.2.4 Hierarchical Geostatistical Model

Suppose that we wish to estimate parameters for a groundwater model, and that these parameters pertain to a spatially-distributed hydraulic property such as hydraulic conductivity. We may wish to estimate the most likely values of this field throughout a model domain; this is what model calibration does. Alternatively, we may wish to sample the posterior probability distribution of the hydraulic property field. These are everyday decision-support modelling tasks. As has already been mentioned, and as will be discussed in detail later in this book, both of these tasks require specification of the prior probability distribution of the hydraulic property field. This is often encapsulated in a correlation decay function, or in a covariance matrix that is based on this function. Choice of the prior probability distribution has a profound effect on the spatial patterns of hydraulic property heterogeneity that emerge through model calibration. It also governs spatial patterns that are expressed by hydraulic property fields that sample the posterior probability distribution of these fields.

A major problem facing real-world decision-support modelling is that a modeller never really knows the prior probability distribution of hydraulic properties, and hence how parameter covariance should be specified when undertaking calibration and uncertainty analysis. Furthermore, it should not be forgotten that earth properties are rarely amenable to characterization using the notion of distance-dependent covariance, so the whole process is approximate at best. It seems like a good idea, therefore, to include geostatistical hyperparameters that govern local hydraulic property variances, correlation lengths and directions in the parameter set that is adjusted through history-matching. History-matching (whether Bayesian or regularized) therefore has the opportunity to estimate not just the locations of hydraulic property heterogeneity, but the patterns that this heterogeneity adopts. Inclusion of both native model parameters and geostatistical hyperparameters in a history-

matching process is referred to as a “hierarchical history-matching”.

Hydraulic property field construction through spatial averaging is well suited to this task, for it provides a natural separation between hydraulic property parameters and geostatistical hyperparameters. This enables their simultaneous but independent adjustment. As is apparent from equations (4.2.2) and (4.2.3), hydraulic properties are represented as iids, that is, as independent normal variates. Spatial integration of these independent normal variates creates a hydraulic property field. The nature of the integration kernel is determined by hyperparameters. If the latter are altered while iid parameters remain the same, patterns of hydraulic property heterogeneity change while the locations of hydraulic property heterogeneity remain unchanged. Furthermore, the prior covariance matrix of iid parameters (which is the identity matrix) remains the same, notwithstanding the fact that the correlation structure of the emergent hydraulic property field is altered. As we shall see, this makes the mechanics of hierarchical history-matching particularly easy.

In the statistical literature, use of concepts that are outlined in this section which enable joint estimation of system properties and the statistical hyperparameters that govern them, while maintaining independence of these two sets of parameters, is sometimes referred to as “noncentred parameterization”.

4.2.5 Implementation Details

Programs of the PEST Groundwater Utility Suite, as well as functions available through the PyEMU library, implement cell-by-cell stochastic hydraulic property field generation using the method of moving averages for both structured and unstructured grids. The method can also be adapted to the use of pilot points as a parameterization device. This facilitates regularized and Bayesian history-matching in contexts where geostatistical hyperparameters that govern hydraulic property variability and patterns adopted by hydraulic property heterogeneity, are spatially variable. It also facilitates hierarchical history matching where geostatistical hyperparameters that govern the spatial patterns of hydraulic property heterogeneity are estimated together with parameters that describe the locations of hydraulic property heterogeneity.

Before discussing this matter further, a discussion of pilot points is warranted. This is provided in the next chapter.

5. Pilot Points and Conceptual Points

5.1 General

This chapter introduces the notion of a model parameter. It also considers some types of parameterization devices that are commonly used in groundwater modelling, and that are supported by software that is available through the PEST and PyEMU suites. Hyperparameters are also discussed. Hyperparameters characterize prior probability distributions that are ascribed to model parameters.

Values given to parameters comprise inputs to a numerical simulator, or are used to calculate inputs to a simulator. Their defining feature is their adjustability. Ideally, for reasons that will become obvious later in this book, they are continuously adjustable. That is, a small change in the value of a parameter results in small changes to model outputs.

Parameters are adjusted during model history-matching in order to enable model outputs to replicate field measurements of system behaviour. Depending on the way in which history-matching is undertaken, they may also be asked to express as much variability as hydrogeological reasonableness allows while still ensuring that model outputs satisfy history-match constraints. Alternatively, they may be asked to suppress any heterogeneity that is not essential for meeting these constraints. In this book, we refer to the former approach to history-matching using terms such as “Bayesian history matching”; the latter approach is termed “regularized inversion” or “model calibration”.

In either case, the role of parameters can be conceptualized as that of hosting information. Some of this information is harvested from field measurements that comprise a history-matching dataset. Rarely is the totality of information that is available at a study site sufficient to ensure the provision of unique values to model parameters (unless a model parameterization scheme is designed with uniqueness in mind). Therefore parameters are adjustable both before and after history-matching. However they are generally less adjustable after history-matching than before history-matching.

It is the history-match responsiveness of parameters to measurements of system behaviour that endows them with the ability to host information that is harvested from that behaviour. At the same time, they can also be viewed as devices for expressing a deficit of information in field measurements. They fill this role by displaying an ability to adopt many different values while still allowing model outputs to fit field measurements.

It is apparent, therefore, that decision-support modelling asks much of model parameters. Hence a model parameterization scheme must be carefully designed. More than this, it can even be argued that if a model has been built to support decision-making, then parameters are the reason for the model’s existence. The making of a decision requires access to information; it also requires an assessment of the management repercussions of information insufficiency. A decision-support groundwater model must harvest information from measurements of system behaviour while expressing post-history-matching uncertainties of decision-critical model predictions. It is the model’s parameters that do all of this work. The role of a numerical simulator is therefore to host a set of parameters that can perform the tasks that decision-making requires. This view of decision-support groundwater modelling differs from more conventional views wherein simulation is considered to be of primary importance while parameters are “model accessories”, similar to fashion accessories such as handbags and necklaces.

5.2 Pilot Points

As is discussed extensively in ensuing chapters of this book, the important roles that parameters must play in decision-support modelling are best served through the deployment of many parameters. This contrasts with notions of parameter parsimony that underpinned historical approaches to groundwater modelling. However, the deployment of many parameters requires some sophistication in adjustment of those parameters in order to prevent bias-inducing parameter unruliness as an outcome of inverse problem ill-posedness. Sometimes it also requires use of inverse problem solution algorithms that avoid the requirement for an excessive number of model runs.

At the time of writing, pilot points have gained widespread acceptance as a parameterization device for groundwater models. Pilot points were introduced to groundwater modelling by Certes and de Marsily (1991). Their use was advanced by authors such as RamaRao et al (1995). The first use of pilot points in regularized inversion was documented by Doherty (2003). Their use by groundwater modellers generally was enabled by programs of the PEST Groundwater Utility Suite. Extensive and flexible use of pilot points is now enabled by the PLPROC parameter preprocessor that is provided with PEST, as well as by the PyEMU library.

The use of pilot points has many advantages. Some of these advantages are now outlined.

1. Pilot points can be deployed in large numbers. In contrast to cell-by-cell parameterization, however, their numbers are not so large that model outputs have diminished sensitivity to values that are assigned to them. This enables calculation of their sensitivities using finite differences.
2. The arrangement of pilot points within a model domain is flexible. They can be emplaced with greater density in those parts of a model domain where data are plentiful and/or where predictions are required, than in those parts of a model domain that are endowed with little data or are of little predictive interest.
3. Pilot points can be used to parameterize all or part of a model domain. Subsets of pilot points can inform zone-based sectors of a model domain independently of other zone-based sectors.
4. Pilot points can be used to parameterize one-dimensional boundary conditions, two-dimensional model layers, and/or three-dimensional spaces that encompasses many model layers.
5. Parameters associated with pilot points can be awarded covariance matrices that are thought to reflect subsurface conditions. If desired, these covariance matrices can embody non-stationarity characterization of hydraulic property variability.
6. Pilot points can host parameters of any type. They can host hydraulic properties. They can also host parameters which multiply or add to hydraulic properties. As will be discussed below, one set of pilot points (“conceptual points”) can host geostatistical hyperparameters that can modify the way in which another set of parameters, hosted by another set of pilot points, determines hydraulic properties that are ultimately used by a model.
7. Pilot points can inform the values ascribed to both structured and unstructured numerical model grids.
8. Pilot point parameters can be of two distinct types; these are described in subsequent sections of this chapter.

5.3 Traditional Pilot Point Usage

5.3.1 Concepts

Traditionally, hydraulic property values, or hydraulic property multiplier/addend values, are directly assigned to pilot points as model parameters. These values then undergo spatial interpolation to a model grid. Hydraulic properties that are interpolated to a model grid can be directly used by the model. In contrast, spatially interpolated multipliers/addends are applied to an existing model-grid-based hydraulic property field on a cell-by-cell basis; the resulting field is then used by the model.

PLPROC (supplied with PEST) supports a number of options for point-to-grid interpolation. These include kriging, radial basis functions and inverse power of distance. Care must be taken in design of an interpolation scheme to prevent the appearance of unsightly “bulls-eyes” in model hydraulic property fields in the vicinity of individual pilot points.

It is apparent that, where pilot points are used as a model parameterization device, the assignment of properties to a model domain becomes a two-step process. First values are assigned to pilot points. These values are then interpolated to the model grid. The “model” that is run by PEST or PEST++ must therefore be encapsulated in a batch file. Before invoking the simulator, the batch file must invoke one or more parameter preprocessors that implement pilot point to model grid interpolation. PLPROC has been designed to accomplish this (and many other) parameter preprocessing tasks.

5.3.2 Accommodating Correlation Anisotropy

Suppose that spatial correlation of hydraulic properties extends for a greater distance in one direction than in the perpendicular direction. (We assume two-dimensional geostatistics because they are easy to discuss; however the same considerations apply in three dimensions.) In traditional pilot point usage, this must be accommodated in two ways.

First, the covariance matrix that is assigned to parameters that are hosted by pilot points must reflect this anisotropy. As will be discussed in subsequent chapters of this book, a parameter covariance matrix serves two purposes. The first is regularization, while the second is generation of random hydraulic property values. In both cases, the geological appropriateness of a hydraulic property field that is provided to a model depends on the appropriateness of the covariance matrix that is ascribed to pilot point parameters from which this hydraulic property field is calculated.

However, model hydraulic property field appropriateness in hydrogeological circumstances where correlation length anisotropy prevails depends on more than this. It also requires that spatial interpolation from pilot points to a model grid accommodate this same correlation anisotropy. This requires “stretching” of spatial coordinates in the direction of correlation elongation, and shrinkage of spatial coordinates in the direction perpendicular to this during the spatial interpolation process. Spatial interpolation functionality provided by PLPROC accommodates anisotropy in this manner.

5.3.3 Accommodating Nonstationarity

Accommodation of nonstationarity is a little more complicated, but can still be achieved where pilot points are used in the traditional manner. In this case, the pilot point parameter covariance matrix must reflect the fact that spatial correlation lengths and directions can vary from place to place within a model domain. So too can parameter variability (i.e. the parameter variance).

As is discussed in Section 3.12, construction of a pilot point covariance matrix can accommodate geostatistical nonstationarity of this type.

Of equal importance is that spatial interpolation from pilot points to the model grid also accommodate spatially varying correlation lengths and directions. This too can be implemented using PLPROC. However, the process is a little complicated. First, spatially varying geostatistical hyperparameters such as hydraulic property variance, anisotropy ratio, correlation length and direction must be assigned to every cell of a model grid. These are then used to enable interpolation of hydraulic property values from pilot points to the model grid in ways that reflect local geostatistical conditions.

5.3.4 Conceptual Points

While the above procedure seems simple enough, it raises the question of how model cells are informed of spatially variable geostatistical properties (i.e. geostatistical hyperparameters).

One way in which to populate a model grid with geostatistical hyperparameters is to interpolate these hyperparameters from a set of points. These hyperparameters may be assigned to the set of pilot points that also host hydraulic property parameters. Alternatively, they may be assigned to another set of pilot points. The latter set of pilot points may be fewer in number than the hydraulic property pilot points. This is because hyperparameters that describe the nature of hydraulic property variability often exhibit less spatial variation than hydraulic property parameters that express this variability.

In this book, we refer to a set of pilot points that host geostatistical hyperparameters as “conceptual points”. They can be thought of as a kind of “hydrogeological mud map” of a study area through which a hydrogeologist can express what he/she knows, and does not know, about an area based on the outcomes of site characterization. A hydrogeologist may know expected values of hydraulic properties, and the penchant for hydraulic property variability about these expected values, based on the properties of materials that are presumed to occupy different parts of a model domain. As well as this, he/she may know the directions and magnitudes of spatial correlation that this variability is likely to exhibit. All of this information can be ascribed to conceptual points. See Figure 5.1.

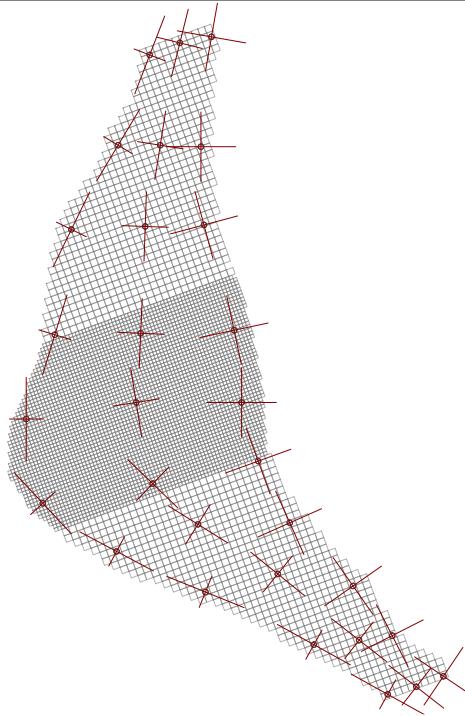


Figure 5.1. A model grid together with conceptual points. Axes of correlation anisotropy are depicted at each conceptual point. Anisotropy ratios, magnitudes and directions (as well as hydraulic property mean values and variances) are spatially interpolated from conceptual points to all cells of the model grid. These are then used to enable anisotropy-aware interpolation of local hydraulic property values from pilot points to the model grid. Placement of pilot points may be much denser than placement of conceptual points in order to facilitate expression of fine scale variability of hydraulic property values in different parts of a model domain.

As stated above, anisotropy-aware interpolation from conceptual points to a model grid can be undertaken using PLPROC. However, where correlation anisotropy is strong and variable, interpolation may need to be undertaken more than once. On the first interpolation occasion, geostatistical isotropy is assumed. Spatial interpolation of hyperparameters that embody geostatistical anisotropy from conceptual points to the model grid is therefore only approximate. On the second interpolation occasion, the previous approximate assignment of anisotropy hyperparameters to cells of the model grid supports anisotropic interpolation of geostatistical hyperparameters from conceptual points to the model grid. Perhaps the spatial interpolation process must then be implemented a third time in order for the overall interpolation process to correctly reflect the anisotropy that it attempts to interpolate.

Once all cells of the model domain have been assigned local geostatistical properties in this manner, spatially-varying anisotropic interpolation of hydraulic properties from pilot points that host these properties to the model grid can be implemented.

5.4 Pilot Points and Spatial Averaging

5.4.1 Concepts

Section 4.2 of this text describes how cells of a model grid can be collectively furnished with hydraulic property values that embody a particular, modeller-specified geostatistical characterization; optionally, this characterization can be nonstationary. Random values are first

assigned to standard normal variates (iids); these are geographically located at the centres of model cells. These iids then undergo spatial averaging using an appropriately specified kernel in order to yield hydraulic property values at these same model cells.

If desired, standard normal variates can be assigned to pilot points instead of individual model cells. The spatial integration process then becomes a summation process. Equation (4.2.2) illustrates one summation option. A more effective methodology for use with pilot points is now described.

Consider model cell i . Consider that the expected value (i.e. mean) of a distributed hydraulic property field \mathbf{k} at this cell is \underline{k}_i . This, of course, is a geostatistical hyperparameter. It may be the same for all model cells; alternatively, it may vary from cell to cell. Consider that the specified hydraulic property variance at cell i is σ_i^2 (another geostatistical hyperparameter). This hyperparameter may also be uniform or spatially variable.

Now let z_j represent the iid value ascribed to pilot point j . Let A_j be the area ascribed to pilot point j . (In PEST-support utilities which implement the methodology that is described herein, the area associated with each pilot point is assigned by the modeller.) The hydraulic property value assigned to model cell i is then calculated as

$$k_i = \underline{k}_i + \sqrt{\frac{\sigma_i^2}{\sum_{j=1}^m (f_{ij} A_j)^2} \sum_{j=1}^m f_{ij} A_j z_j} \quad (5.4.1)$$

where m is the number of pilot points. A_j values should be correct (or at least approximately so) in a relative sense. If pilot points are arranged on a regular grid, then its value is of no consequence. Only relativity of A_j correctness is necessary because the normalization term, i.e.

$\sqrt{\frac{\sigma_i^2}{\sum_{j=1}^m (f_{ij} A_j)^2}}$, ensures that the variance of k_i is indeed equal to its desired value of σ_i^2 .

In equation (5.4.1) f_{ij} is the value of the averaging kernel. For reasons discussed in section 4.2.2, this is a function of the distance between model cell i and pilot point j . f_{ij} is a decaying function of distance; calculation of distance takes correlation anisotropy into account if required. Where conditions are non-stationary, the dependence of f_{ij} on cell-to-pilot-point distance may vary from cell to cell throughout a model domain.

At the time of writing, PLPROC and programs of the PEST Groundwater Utility Suite offer a number of options for the functional dependence of f_{ij} on distance between model cell i and pilot point j . All of these options allow a modeller to specify the rate of decay of the selected function with distance. Decay rates can be exponential, Gaussian or polynomial. Similar functionality is offered by PyEMU.

This method of model grid hydraulic property assignment is similar to the “noncentred” parameterization methodology that is described in section 4.2. Parameters that ultimately determine the locations of hydraulic property heterogeneity within a model domain are iid normal variates ascribed to pilot points. Meanwhile, geostatistical hyperparameters that determine the background hydraulic property value, the permitted degree of variability about this value, and patterns that emergent heterogeneity can adopt, are the hydraulic property mean value, its variance, and variables that determine the decay rate of f_{ij} with distance. If spatial correlation is anisotropic, then two spatial f_{ij} decay rates must be supplied together with the direction of the principal axis of anisotropy. Under nonstationary conditions, some or all of these geostatistical hyperparameters may vary in space. (As was stated previously, the

description provided herein assumes a two-dimensional model domain. Similar considerations apply in three dimensions. PLPROC and PEST Groundwater Utility Suite programs offer both two- and three-dimensional versions of this methodology.)

5.4.2 Pilot Point Parameters

If random values that sample the standard normal distribution (which has a mean of zero and a standard deviation of one) are assigned to pilot point iid parameters, equation (5.4.1) can be used to equip a model domain with a spatially correlated random hydraulic property field. However this hydraulic property field will be smoother than that which would arise if iids were assigned to individual model cells. In many modelling circumstances, this is of little concern.

However use of equation (5.4.1) is not limited to generation of random hydraulic property fields. The iids which feature in this equation can be considered as model parameters. As such, they can be adjusted for the purpose of model calibration (i.e. regularized inversion). They can also be adjusted during ensemble-based Bayesian history-matching which encourages parameter diversity while ensuring that pertinent model outputs respect field measurements. Note that if all pilot point iid parameters have values of zero, then the hydraulic property field with which a model is equipped is just the mean hydraulic property field (which may, or may not, be spatially variable).

5.4.3 Conceptual Points

Where geostatistical stationarity prevails throughout a model domain, hyperparameters are spatially invariant. In contrast, under nonstationary conditions they are spatially variable. The dependence of f_{ij} of equation (5.4.1) on the distance between model cell i and pilot point j becomes cell-dependent. Therefore, before application of equation (5.4.1), geostatistical hyperparameters must be assigned to every cell of a model domain. This can be achieved through spatial interpolation from conceptual points. The means by which this can be achieved, including the means through which geostatistical hyperparameters describing spatially varying anisotropy are interpolated to cells of a model domain in a way that accommodates the anisotropy that they interpolate, is described in section 5.3.4.

5.4.4 Hierarchical History-Matching

As has already been described, hierarchical history-matching attempts to simultaneously adjust model hydraulic properties together with hyperparameters which govern the patterns that these properties adopt.

Where model hydraulic property field assignment takes place using equation (5.4.1), pilot point parameters are iids. As such, their collective prior covariance matrix (i.e. $C(\mathbf{z})$) is the identity matrix \mathbf{I} . Hence (by construction) they exhibit no spatial correlation. Meanwhile, spatial correlation that prevails in the hydraulic property field that is used by the model is determined by the f_{ij} function of equation (5.4.1) that operates on pilot point iid parameters. This function is specified through its (possibly spatially varying) hyperparameters.

Suppose that a modeller wishes to undertake hierarchical inversion. Hence values ascribed to pilot point iid parameters are estimated together with the values of at least some geostatistical hyperparameters. Let us assume that the latter are ascribed to conceptual points. When the values of one or more geostatistical hyperparameters are altered, interpolation of these altered values from conceptual points to the model grid must be undertaken afresh. Equation (5.4.1) is then implemented using the new f_{ij} function. These are not numerically difficult tasks.

However a significant conceptual/numerical advantage is accrued by separating the covariance matrix of parameters that PEST sees from the geostatistical properties of the hydraulic property field that the model sees. This separation is built into the noncentred approach to model parameterization that is now being described. The covariance matrix that PEST sees is $C(\mathbf{z})$ while the geostatistical properties of the hydraulic property field that is seen by the model is determined by the (possibly spatially-varying) f_{ij} function employed in equation (5.4.1). This is particularly advantageous when undertaking ensemble-based history-matching. As will be discussed in later chapters, ensemble-based history-matching commences with generation of random values of model parameters. These initial, random values are based on the prior probability distributions of these parameters. These initial values are then modified until model outputs match field measurements. However this is done in a way that attempts to maintain respect for their prior probability distribution. When hierarchical history-matching is undertaken, $C(\mathbf{k})$ (i.e. the covariance matrix of the parameter field that is seen by the model) is altered during the history-matching process. Hence the post-history-matching $C(\mathbf{k})$ is likely to differ from the prior $C(\mathbf{k})$. This does not matter to PEST/PEST++ as it adjusts iid model parameters that continue to respect the original $C(\mathbf{z})$ covariance matrix (i.e. the identity matrix) regardless of history-match-induced alterations to $C(\mathbf{k})$ as hyperparameters are altered.

In contrast, hierarchical ensemble-based history-matching is much more difficult to implement where pilot points are parameterized in the traditional way that is described in section 5.3. In this case the prior covariance matrix of pilot point parameters is $C(\mathbf{k})$. Alterations to $C(\mathbf{k})$ as the hierarchical inversion process progresses invalidate their prior probability distribution.

Similar considerations apply to regularized inversion. In this case the prior covariance matrix of pilot point parameters is used for regularization. When implementing noncentred pilot point parameterization, the covariance matrix that is applied to pilot point iid parameters remains the identity matrix, regardless of values adopted by geostatistical hyperparameters, and therefore regardless of the $C(\mathbf{k})$ matrix that describes the hydraulic property field that is used by the model. Traditional pilot point usage does not offer this advantage. Furthermore, because spatial interpolation from pilot points to the model grid must also respect a possibly changing $C(\mathbf{k})$ matrix, the interpolation procedure must undergo continuous updating as $C(\mathbf{k})$ is changed. In contrast, for iid pilot points parameterization that is employed in a noncentred parameterization scheme, spatial interpolation from pilot points to a model grid is replaced by spatial averaging. The averaging kernel is automatically updated as geostatistical hyperparameters are updated.

5.4.5 Conditioning of Hydraulic Property Fields

Site characterization studies often yield point measurements of system hydraulic properties. For example, estimates of hydraulic conductivity are often available from aquifer tests. It makes sense to respect these point measurements when building a model-based hydraulic property field through spatial averaging of pilot point iid parameters. “Conditioning” of the model hydraulic property field by these local hydraulic property measurements should occur whether pilot points are supplied with random iid values for the purpose of stochastic hydraulic property field generation, or whether pilot point iid parameters are undergoing history-match adjustment in order for model outputs to match a history-matching dataset.

A relatively easy way to condition stochastic or manipulated hydraulic property fields produced through spatial averaging is to assign field-measured values to local mean field values (i.e. to \underline{k}_i of equation 5.4.1) at model cells that are close to measurement points. Perhaps values of local hydraulic property variance (i.e. σ^2_i) can be reduced for these same cells.

This process can be automated using a heuristic procedure that is implemented by PEST

Groundwater Utility Suite software and by PLPROC. Before describing this methodology, we note that the kernel averaging function f_{ij} of equation (5.4.1) has a maximum value of one, and is always positive.

Let d_{ij} denote the “effective conditioning distance” from model cell i to “conditioning point” j . Conditioning points are those at which direct measurements of hydraulic property \mathbf{k} were made. d_{ij} is calculated as

$$d_{ij} = \frac{1-f_{ij}}{f_{ij}} \quad (5.4.2)$$

The conditioning distance approaches infinity where a conditioning point is far from model cell i . It approaches zero where a conditioning point is close to the centre of model cell i . Hydraulic property conditioning results in a new mean and variance for use in equation (5.4.1). We symbolize these as \underline{k}' and $\sigma_i^{2'}$. \underline{k}' is calculated as

$$\underline{k}' = \underline{k}_i + \frac{\sum_{j=1}^p (c_j - \underline{k}_i) / d_{ij}}{\sum_{j=1}^p 1/d_{ij}} \quad (5.4.3)$$

In equation (5.4.3) c_j is the measured value of the \mathbf{k} hydraulic property at conditioning point j , while p is the total number of conditioning points. Meanwhile $\sigma_i^{2'}$ is calculated as

$$\sigma_i^{2'} = \sigma_i^2 (1 - f_{ijmax}) \quad (5.4.4)$$

where f_{ijmax} is the maximum value of the weighting function f_{ij} over all conditioning points.

An alternative option for respecting direct measurements of system properties when history-matching a model is to include these measurements in a history-matching dataset. The objective function that is minimized through history-matching is therefore penalized to the extent that local estimates of hydraulic properties do not coincide with local measurements of hydraulic properties. An advantage of this option is that respect for measurements of system properties can be traded against respect for measurements of system states and fluxes (which are also featured in the objective function) if there is a conflict between the two.

The extent to which local estimates of system hydraulic properties derived through aquifer testing should be respected in hydraulic property fields that are employed by a regional groundwater model is a matter that has not received much attention in the literature. Flow fields that are simulated during aquifer test analysis are local and radial; as such, they are very different from those that are simulated by a regional model.

The resolution matrix is described in section 2.9 of this book, and will be discussed in greater detail in later chapters. Manewell et al (2023a and 2023b) evaluate this matrix where regularization is implemented by use of the Theis equation to evaluate transmissivity and storativity from pumping-induced drawdowns. They show that real-world hydraulic properties undergo complex, time-varying, spatial averaging in order to yield the single estimates of transmissivity and storativity that emerge from aquifer test interpretation. It follows that while aquifer-test-interpreted hydraulic conductivities should not be ignored by a modeller, their influence on hydraulic properties that are ascribed to a regional model should be tempered, especially where high levels of heterogeneity prevail. Manewell et al (2023b) discuss how differences in scale, observation configurations and flow fields can be accommodated when attempting to constrain regional model hydraulic conductivity fields by hydraulic conductivities that have emerged from interpretation of local aquifer tests. These authors suggest that the simplest (and possibly most effective) way to use aquifer-test-derived estimates of hydraulic conductivity in a regional groundwater model is to allow these estimates to dictate

local prior mean values of hydraulic conductivity that are employed by the model; however local prior variances about this mean should not be reduced from those that have emerged from other aspects of site characterization.

5.5 Some Other Parameterization Devices

As will be discussed in ensuing chapters, in addition to their role of hosting prediction-pertinent information that is harvested from field measurements of system behaviour, parameters often serve a regularization role. Where parameters are carelessly defined, these two roles may conflict with each other. As is discussed later, such conflicts should be recognized and rectified; in general, regularization is best achieved through mathematical means, and not through parameterization.

Zones of assumed hydraulic property constancy have been deployed as parameterization devices for groundwater models since the inception of groundwater modelling. Zone boundaries are often based on known or inferred geological boundaries. The deficiencies of zone-based parameterization are discussed later in this book. We note, however, that known geological boundaries should not be ignored when parameterizing a groundwater model. We also note that pilot points can be grouped according to the zones in which they lie. By ensuring that spatial interpolation or spatial averaging (depending on which mode of pilot point usage is adopted) does not cross zone boundaries, a history-matching process can be encouraged to introduce sharp hydraulic property changes at these boundaries in preference to more gradual spatial changes at other locations. Meanwhile, zone-specific pilot points can be used to represent intrazonal heterogeneity.

More exotic parameterization devices may be warranted at particular sites. Recall that parameters must be continuously adjustable in order to fulfill their history-matching role as hosts for information. PLPROC supports so-called “structural overlay parameters”. These are polyline or polygonal features whose vertex locations are continuously adjustable. Hydraulic properties can vary along or within these features. These properties are transferred to an underlying model grid through superimposition. At the same time, hydraulic properties at the edges of these features are “blended” with those of the material which hosts them. The blending width is set by the modeller; see Figure 5.2. This blending operation ensures that as these features are moved, model outputs do not suffer discontinuities as feature boundaries cross model cell boundaries. See Doherty (2025a) for more details.

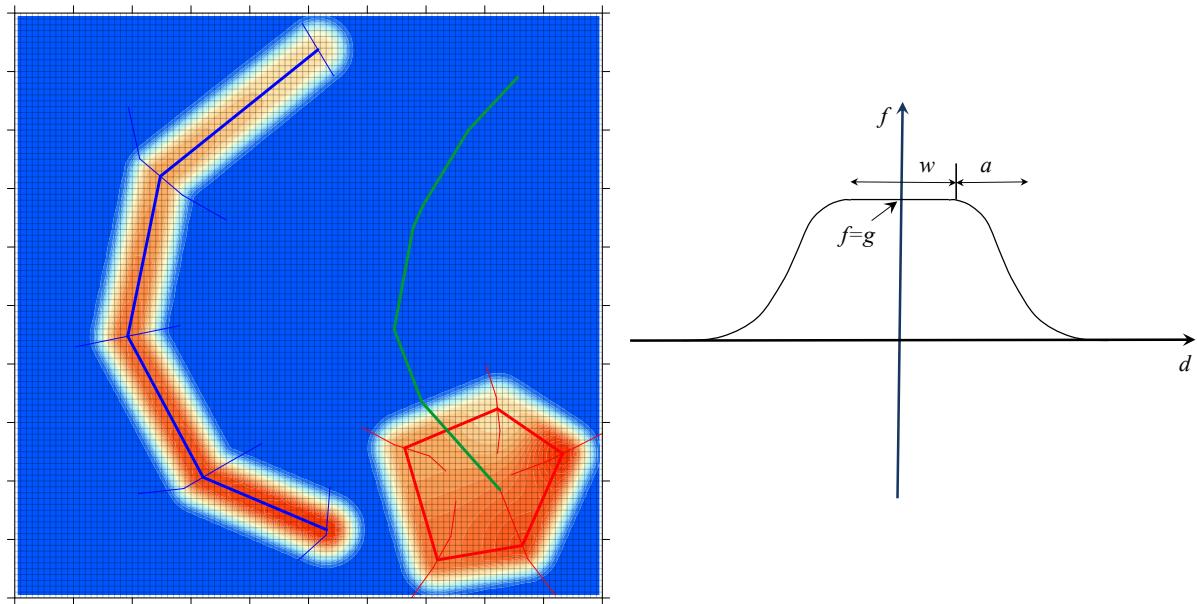


Figure 5.2. The left part of the figure depicts two examples of structural overlay parameters. One of these features is polylinear while the other is polygonal. Vertices of these features can move along “sliders” to which they are attached. The polygon can slide as a whole along another linear feature. Colouration reflects grid hydraulic property values. The right part of figure depicts the blending function through which hydraulic properties that are ascribed to a feature are mapped to an underlying model grid.

6. Situation Statement

6.1 Model Parameters

We now turn our attention from parameterization devices to the theory and practice of parameter estimation and parameter and predictive uncertainty analysis.

Let the vector \mathbf{k} denote parameters used by a model. In general, these tend to be properties of the system that a model purports to simulate, for example the permeabilities of porous media, the roughness coefficients of streams, or the coefficients appearing in semi-empirical equations that are used in place of direct simulation to calculate certain system states. However they may also be the values of stresses to which a system has been subjected, for example recharge or stream inflow. Any system properties or stresses that are incompletely known can be designated as parameters.

Because the values of model parameters are incompletely known, they are uncertain. They are also adjustable. This allows them to freely express their uncertainties. It also allows them to host information. A parameter that hosts information acquired through model history-matching possesses less variability than that which it possesses based on expert knowledge alone. Constraints on its variability are required in order for pertinent model outputs to respect field measurements of system states and fluxes. Perhaps the range of values that it can adopt is somewhat or severely limited after model history-matching. Alternatively, perhaps an individual parameter can vary just as much as ever after model history-matching, but other parameters must vary in harmony with it in order to maintain model-to-measurement fit. Either way, it is not as free as it used to be. This restriction to its freedom is how it hosts information.

The number of elements comprising the vector \mathbf{k} can be very large when simulating complex environmental systems. This is because environmental processes are complex, and the properties which govern these processes exhibit considerable spatial heterogeneity. Perhaps, as a practical measure, the number of elements comprising \mathbf{k} can be reduced where predictions of management interest are sensitive only to broadscale system behaviour and properties. In other cases this will not be possible, as decision-critical model predictions may be sensitive to fine scale spatial heterogeneity. In cases such as these, representation of system property heterogeneity at the prediction-pertinent scale becomes important despite the fact that parameters which encapsulate these properties may be neither measurable nor inferable at this scale. This philosophy differs from old-fashioned approaches to model parameterization. Specification of a model parameterization scheme should not reflect the capacity of parameters to be uniquely estimated through site characterization studies or through history-matching. Prediction-pertinent parameters require representation in a model not because they can be uniquely estimated, but precisely because they cannot be uniquely estimated. This renders their representation critical for model predictive uncertainty analysis.

The situation becomes more complex when it is considered that numerical models are imperfect simulators of the behaviour of environmental systems. A model's representation of real world processes and properties may be somewhat abstract, relying on empirical relationships rather than exact numerical simulation of some of its physical and chemical processes. Even where a model purports to be “physically-based”, environmental processes and properties can be represented at no finer scale than that of the model's grid. Hence, as will be discussed in greater detail in later chapters of this book, some of its parameters may need to adopt roles that compensate for its inadequacies as they are adjusted in order for model outputs to replicate

real-world measurements of system behaviour. For the moment, however, we will assume for convenience that models have no defects; this simplifies the following discussion.

While an environmental expert may not know the exact values of the numbers which comprise the elements of \mathbf{k} at a particular study site, his/her ignorance generally has limits. Experts rightly raise objections when confronted with models whose parameter values are inappropriate. This constitutes knowledge. Direct measurements of system properties may have been made at one or a number of locations within a model domain. This too constitutes knowledge. However these measurements may have been accompanied by error. Furthermore they may pertain to only a point or vertical line (for example a drill hole) in a much more voluminous two- or three-dimensional model domain. As such, they may only be indicative of the average system property over the model cell in which the measurement point is situated; they may be even less indicative of average system properties in neighbouring model cells.

Knowledge of the parameter set \mathbf{k} with which a model should be equipped is thus probabilistic in nature, reflecting the fact that the exact values of its elements cannot be known, but that limits can be placed on these values. It may be possible to supplement this knowledge with more qualitative, though no less important, descriptions of the patterns that are adopted by system property heterogeneity. For example an expert may know scales of spatial correlation that system properties exhibit. In the geological context, he/she may know something of the nature and disposition of structural and/or sedimentary features that are likely to give rise to anomalous system properties. He/she may therefore know something about the connectedness of hydraulic properties, particularly hydraulic conductivity. Connectedness of hydraulic conductivity is as important to movement of water and contaminants as its magnitude.

Conceptually, expert knowledge of \mathbf{k} can be expressed as a probability density function; we will refer to this probability density function as $f(\mathbf{k})$. Unfortunately, this is unlikely to be a neat and tidy analytical density function. Nor is it likely to exhibit nonstationary, or even multiGaussian behaviour. Instead, it may be possible to express $f(\mathbf{k})$ only by sampling it.

In Bayesian terminology, $f(\mathbf{k})$ is referred to as the “prior probability density function” of parameters \mathbf{k} . As such, it is the embodiment of expert knowledge as it pertains to a particular study site. Conceptually at least, a covariance matrix can be associated with $f(\mathbf{k})$. We designate this covariance matrix as $C(\mathbf{k})$.

6.2 History-matching

Suppose that we have at our disposal a set of measurements of the historical behaviour of the system which is the subject of modelling-based investigations. For a surface water system these may be streamflows. For a groundwater system they may include water levels and contaminant concentrations in wells, as well as groundwater contributions to river baseflows. System stresses at the time of these measurements may have been “natural” or anthropogenic. For a petroleum reservoir, a dataset of historical measurements may be comprised of production rates and water cuts. Being a collection of numbers, these measurements comprise a vector. We will name this vector \mathbf{h} .

Measurements of system state are accompanied by error – often referred to as “noise”. For example, samples of river flow during floods may be very noisy because of the high potential for error that accompanies the taking of such measurements. In contrast, other types of measurement can be made with a high level of instrumental precision; nevertheless their “noise content” may still be high because of their questionable representativeness of the system state that modelling is meant to portray. For example, a borehole measurement on a particular day

may not be truly representative of “steady state” or “equilibrium” groundwater conditions. Nor may it be truly representative of the average head within a thick stratigraphic layer that is represented as a single layer in a groundwater model. The subject of so-called “structural noise” will receive considerable attention in later chapters of this book. For now, discrepancies between measurements of system state and the capacity of a model to represent these states will be classified as “measurement error”, for this makes it easy to handle using the mathematics of inversion and uncertainty analysis.

Any measurement of the state of an environmental system must therefore be regarded as “contaminated” to some extent. Therefore it cannot be considered to represent the true state of the system (or the system as represented by a model) at the time at which the measurement was made. Instead, it is the state of the system plus an unknown error. For measurement i we will designate this error as ϵ_i . Collectively, the totality of measurement errors comprise the vector ϵ ; the i ’th component of ϵ is the noise associated with the i ’th component of \mathbf{h} .

Like \mathbf{k} , ϵ is a stochastic quantity. Therefore, conceptually, it can be ascribed a probability distribution $f(\epsilon)$. This probability distribution can, in turn, be ascribed a covariance matrix $C(\epsilon)$. The probability distribution associated with measurement noise is often considered to be relatively uncomplicated, unlike that associated with model parameters. In fact, the noise associated with any one measurement is often considered to be statistically independent of that associated with any other measurement. $C(\epsilon)$ thus becomes a diagonal matrix; each diagonal element of this matrix is the variance (square of the standard deviation) of measurement noise associated with the corresponding individual measurement. Furthermore, the expected value of ϵ is normally assumed to be $\mathbf{0}$. Measurements are thus assumed to be made without bias.

In common history-matching practice, a multiplier is associated with $C(\epsilon)$. This is often referred to as the “reference variance”; we designate it herein as σ_r^2 . This recognizes the fact that the level of model-to-measurement fit that can be achieved through history-matching may not be known in advance. Hence while a modeller may feel comfortable in specifying the statistics of multi-measurement error in a relative sense through use of an appropriate $C(\epsilon)$ matrix, the actual values of the elements of $C(\epsilon)$ are determined as part of the history-matching process itself while inter-element relativity is preserved. This raises the spectre, of course, that it may not be actual measurement noise that is responsible for model-to-measurement misfit. Where higher levels of misfit are encountered than was anticipated on the basis of measurement error alone, this suggests that model-to-measurement misfit may be an outcome of model structural defects. If this is the case, then a pre-specified, diagonal $C(\epsilon)$ may not be correct. This topic will be ignored for the moment, but will receive detailed treatment later in this book.

If a model is run using a particular set of parameters \mathbf{k} , then it produces a set of outputs which are associated with these parameters; we designate this set of outputs using the vector \mathbf{o} . In order to produce these outputs, the model can be considered to be an operator on these parameters. This operator will be designated as $Z[\cdot]$. The “ Z ” symbol is italicized to represent the fact that, in general, a model operator is nonlinear. It is capitalized because it operates on many parameters, these comprising the vector \mathbf{k} , and because it yields many outputs, these comprising the vector \mathbf{o} . With this notation, the action of the model over the time span for which historical measurements of system state are available can be written as

$$\mathbf{o} = Z[\mathbf{k}] \quad (6.2.1)$$

Suppose, for a moment, that we know the correct values of all model parameters and that we provide these to a model. The outputs \mathbf{o} calculated by the model under these circumstances would not reproduce the measurement dataset \mathbf{h} because, as stated above, the latter are

contaminated by measurement noise. Hence

$$\mathbf{h} = \mathbf{Z}[\mathbf{k}] + \boldsymbol{\varepsilon} \quad (6.2.2)$$

Equation (6.2.2) states that measurements comprising the history-matching dataset \mathbf{h} are scattered about the true values of system state \mathbf{o} that would be calculated by the model using the true set of parameters \mathbf{k} . This scatter is described by the probability density function $f(\boldsymbol{\varepsilon})$.

Of course, in real-world modelling contexts we do not know the true \mathbf{k} . As stated above, however, presumably, samples of \mathbf{k} can be drawn from its prior probability distribution $f(\mathbf{k})$. Suppose that we do this. For each realization of \mathbf{k} that we draw, a corresponding set of “phantom measurements” \mathbf{h} can then be obtained by first calculating \mathbf{o} using (6.2.1), then sampling $f(\boldsymbol{\varepsilon})$ to obtain a realization of $\boldsymbol{\varepsilon}$, and then adding the two together. From this it is apparent that $f(\boldsymbol{\varepsilon})$ can be considered to specify $f(\mathbf{h}|\mathbf{k})$, i.e. the probability density function of \mathbf{h} given \mathbf{k} , or the conditional probability density function of \mathbf{h} . This can be written formally as

$$f(\boldsymbol{\varepsilon}) = f(\mathbf{h}|\mathbf{k}) \quad (6.2.3)$$

Meanwhile, the marginal probability density function of \mathbf{h} , i.e. $f(\mathbf{h})$, can be sampled by repeating the following procedure.

1. Sample \mathbf{k} (i.e. generate a realization of \mathbf{k});
2. Run the model to obtain a corresponding realization of \mathbf{o} ;
3. Sample $\boldsymbol{\varepsilon}$;
4. Add it to \mathbf{o} .

If this were done many times, the complete marginal probability distribution of \mathbf{h} could be determined empirically. However, empirical estimation of $f(\mathbf{h})$ requires a great many model runs, especially if \mathbf{k} has many elements. Alternatively, if the density function $f(\mathbf{k})$ can be characterized using a mathematical expression (as can the multiGaussian probability distribution), and if the model operator $\mathbf{Z}[\cdot]$ can be similarly characterized (for example if $\mathbf{Z}[\cdot]$ is linear and can be specified as a matrix), a mathematical expression for $f(\mathbf{h})$ can be obtained. In real world modelling practice, however, models are complex and nonlinear, and no such expressions for $f(\mathbf{h})$ are obtainable.

In a real world history-matching context, a modeller has the following items at his/her disposal (or makes whatever assumptions are necessary to place these items at his/her disposal):

- a single sample of \mathbf{h} , this comprising a measurement dataset;
- knowledge of $C(\boldsymbol{\varepsilon})$ (but not necessarily of the reference variance σ_r^2);
- knowledge of $f(\mathbf{k})$ (or at least the means to generate samples of \mathbf{k}).

While \mathbf{h} and \mathbf{k} are random variables, it should be apparent from the above discussion that they are not independent random variables. They are linked by the action of the model. Knowledge of \mathbf{h} thus has a conditioning effect on $f(\mathbf{k})$. Theoretically (see section 3.7) it should therefore be possible to determine a conditional probability density function for \mathbf{k} , i.e. $f(\mathbf{k}|\mathbf{h})$, based on knowledge of \mathbf{h} . Alternatively, it should be possible to at least draw samples from this conditional distribution so that $f(\mathbf{k}|\mathbf{h})$ can be characterized empirically.

From (3.7.5)

$$f(\mathbf{k}|\mathbf{h}) = \frac{f(\mathbf{h}|\mathbf{k})f(\mathbf{k})}{f(\mathbf{h})} \quad (6.2.4)$$

This is Bayes equation as it pertains to history-matching. As stated above $f(\mathbf{k})$, the encapsulation of expert knowledge, is referred to as the prior probability density function of \mathbf{k} .

Meanwhile $f(\mathbf{k}|\mathbf{h})$ is referred to as the posterior probability density function of \mathbf{k} . $f(\mathbf{h}|\mathbf{k})$ is referred to as the likelihood function; the better is the fit between model outputs \mathbf{o} and field measurements \mathbf{h} , the higher is the value of the likelihood function. $f(\mathbf{h})$ is referred to as the model evidence.

The reduction in uncertainty of \mathbf{k} that is accrued through the conditioning process (i.e. the history-matching process) is a measure of the information content of the history-matching dataset \mathbf{h} . It is important to note, however, that \mathbf{h} is likely to be more informative of some model parameters (and combinations of parameters) than others. Hence when a history-matched model is used to make predictions of management interest, the uncertainties of some of its predictions are likely to be reduced from those that prevailed prior to history-matching, while those of others may not be reduced much at all. Moore and Doherty (2006) show an example of a “perfectly calibrated model” for which the uncertainty of one particular model output of potential management interest is barely reduced through history-matching, in spite of the conditioning exerted on some combinations of parameter by a noiseless history-matching dataset.

As stated above, the denominator of (6.2.4) is the marginal probability density function of \mathbf{h} . If this is written as a conditional (on \mathbf{k}) density function, and if \mathbf{k} is then integrated out, (6.2.4) can be re-written as

$$f(\mathbf{k}|\mathbf{h}) = \frac{f(\mathbf{h}|\mathbf{k})f(\mathbf{k})}{\int_{-\infty}^{\infty} f(\mathbf{h}|\mathbf{k})f(\mathbf{k})d\mathbf{k}} \quad (6.2.5)$$

By integrating both sides of (6.2.5) with respect to \mathbf{k} , it is apparent that the integral of the posterior density function of \mathbf{k} over all of \mathbf{k} -space is unity, as must be the case for any probability density function.

Let s , a scalar, be a prediction of management interest made by a model. Let the model operator under these predictive conditions be designated as $\mathbf{y}[\cdot]$. The bold status of “ \mathbf{y} ” indicates that, like $\mathbf{Z}[\cdot]$, it operates on the vector of model parameters \mathbf{k} ; its italicized status followed by square brackets indicates that it is a nonlinear operator and not a matrix. Use of lower case to symbolize \mathbf{y} indicates that it yields a single number, this being the prediction s . Thus

$$s = \mathbf{y}[\mathbf{k}] \quad (6.2.6)$$

Figure 6.1 conceptualizes the prior probability density function of this prediction, which we denote as $f(s)$. This density function could be sampled by first sampling \mathbf{k} based on its prior probability density function $f(\mathbf{k})$, and then running the model using each sample to obtain a corresponding sample of s . A histogram of s could then be constructed and its probability density function determined empirically. Also depicted in figure 6.1 is a schematic of the posterior probability density function of this prediction, which we denote as $f(s|\mathbf{h})$. Conceptually, this probability density function could also be constructed empirically by running the model using samples of the posterior density function of \mathbf{k} , i.e. $f(\mathbf{k}|\mathbf{h})$. The extent to which the posterior predictive uncertainty of s is smaller than its prior predictive uncertainty (i.e. the extent to which $f(s|\mathbf{h})$ is narrower than $f(s)$) reflects the information content of the history-matching dataset, the model parameters to which this information is pertinent, and the parameters to which s is most sensitive.

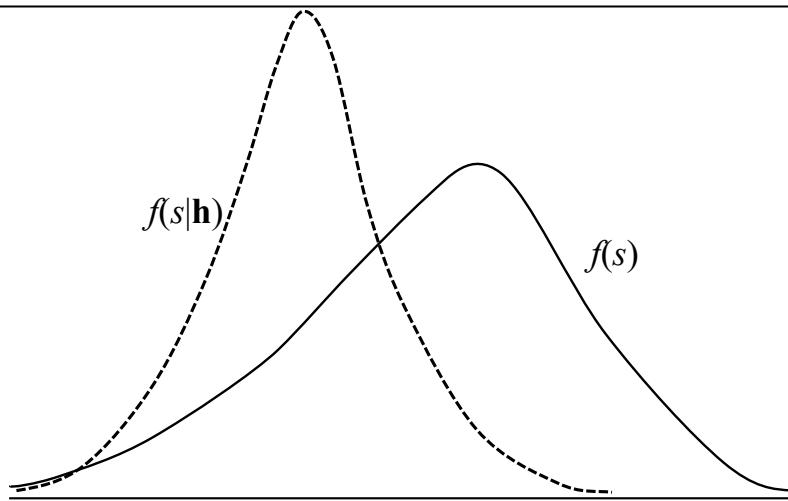


Figure 6.1. Prior and posterior predictive probability density functions. The integral under both of these is 1.0.

6.3 Sampling the Posterior

6.3.1 Rejection Sampling

The concepts expressed by Bayes equation are simple enough. Rudimentary implementation of Bayes equation is also simple enough through a process known as “rejection sampling”. This process is now described.

The prior parameter probability density function is sampled. Where the prior density function is explicit (for example if the potential for variability of parameters \mathbf{k} can be described by a multiGaussian probability distribution), the random number generation process by which samples of \mathbf{k} are obtained is relatively straightforward; this applies even under conditions of geostatistical nonstationarity that are discussed in previous chapters of this book. However where $f(\mathbf{k})$ is implicit rather than explicit and requires, for example, generation of complex, categorical geostatistical or process-driven hydraulic property fields, then the process of sampling the prior distribution of \mathbf{k} may be considerably more complex.

A model run is then undertaken based on each sampled parameter field. The fit between model outputs and field data is examined. If this fit is not within expectations based on measurement noise, then the parameter field \mathbf{k} is rejected. Otherwise it is retained as a sample of the posterior parameter probability distribution.

This process can be formalized in a way that ensures that Bayes equation is exactly obeyed. For each realization of \mathbf{k} , the exact value of the likelihood function $f(\mathbf{h}|\mathbf{k})$ can be calculated; this is calculated from model-to-measurement misfit. This is then divided by the maximum value that the likelihood function can take, that is the maximum value of $f(\mathbf{e})$. Suppose that this ratio is r . A random number is then generated from a uniform distribution that extends between zero and one. If this random number is greater than r the sample of \mathbf{k} is rejected. If it is less than r , the sample of \mathbf{k} is retained as a sample of the posterior distribution of \mathbf{k} .

For obvious reasons, this procedure is referred to as “rejection sampling”. From the above description, it is apparent that it imposes two levels of “filtering” on parameters \mathbf{k} used by a model. The first is imposed by the prior parameter probability distribution of \mathbf{k} ; sampling of $f(\mathbf{k})$, rather than any other probability distribution, results in parameter fields that are realistic as far as the study site is concerned. The second filter is comprised of the likelihood function

$f(\mathbf{h}|\mathbf{k})$. Parameter fields \mathbf{k} which give rise to poor fits with field measurements (i.e. unlikely fits) are rejected through this function.

Rejection sampling has the advantage that it is conceptually simple. Furthermore, the methodology through which it is implemented illustrates some important features of the history-matching process, namely that

- a) history-matching can be viewed as a kind of probability density function filtering operation, and that
- b) the outcome of this filtering operation is another probability density function.

The last point is important. History-matching does not promise parameter or predictive certainty. This challenges the use of terms such as “calibration” which imply a sense of parameter finality and accuracy. This matter is further discussed below.

Unfortunately, while conceptually simple, rejection sampling is impossible to implement in most practical groundwater modelling circumstances. Even where parameter numbers are very low, it is a deeply inefficient procedure, especially where a modeller is entitled to seek a good fit between model outcomes and field data. The “hit to miss ratio” (i.e. the ratio of parameter acceptances to parameter rejections) is too low to be useable, so that many model runs are squandered. Where model run times are high, this is unacceptable.

6.3.2 Markov Chain Monte Carlo

Markov chain Monte Carlo (MCMC) provides another option for Bayesian sampling of the posterior parameter probability distribution. An attractive feature of MCMC is that it is much more efficient than rejection sampling. However it is not efficient enough to be routinely useful in groundwater modelling contexts where parameter numbers may be high and model execution speeds may be low.

When implementing MCMC, the prior parameter probability distribution is generally not sampled directly. Instead, a Markov chain of parameter samples is established. In a Markov chain, the next member of the chain depends only on the previous member of the chain. The next sample in a Markov chain of parameters is generated using a proposal density function whose formulation generally depends on the value of the previous sample. The model is then run using the new parameter sample \mathbf{k} so that model outputs \mathbf{o} corresponding to this new sample can be obtained. Using $f(\mathbf{e})$, the product $f(\mathbf{h}|\mathbf{k})f(\mathbf{k})$ is then evaluated and its ratio taken with the same product calculated using the previous parameter sample. If this ratio is greater than unity, the new sample of \mathbf{k} is accepted as the next member of the Markov chain. If it is less than unity, the new sample is accepted or rejected after generating a random number from a uniform probability distribution that spans the interval between zero and one. If this random number is less than the ratio of new to old $f(\mathbf{h}|\mathbf{k})f(\mathbf{k})$, the proposed parameter set becomes the next member of the chain. Otherwise the sampled \mathbf{k} is rejected and another \mathbf{k} is tested using the same proposal density function.

This is the so-called Metropolis-Hastings algorithm. Provided the proposal density function is symmetric, that the Markov chain is irreducible (i.e. all states can be sampled from the current state) and aperiodic (i.e. the number of steps until the chain returns to its current state is not a multiple of an integer greater than 1), it can be shown that samples generated using this algorithm are actually samples of $f(\mathbf{k}|\mathbf{h})$, i.e. the posterior probability density function of \mathbf{k} . See texts such as Gelman et al (2025) for further details.

As stated above, MCMC constitutes a far more efficient means of sampling the posterior parameter probability distribution than rejection sampling. Furthermore, hyperparameters such

as the reference variance σ_r^2 can be included in the parameter conditioning process; hence the level of noise associated with a measurement dataset \mathbf{h} can be estimated at the same time as a good fit with that dataset is sought.

Nevertheless, despite the fact that powerful, parallel versions of MCMC are available to modellers through packages such as DREAM (see, for example, Vrugt, 2016), as well as through a number of Python libraries, the high model run burden of MCMC becomes a problem where model run times are long and \mathbf{k} has many elements (more than about 50). Tens of thousands to hundreds of thousands of model runs may then be required to ensure adequate sampling of the posterior parameter probability distribution. Furthermore, where parameter numbers are high, the acceptance-to-rejection ratio of the MCMC sampler may become very small, particularly where the parameter null space must be sampled for proper characterization of parameter and predictive uncertainty. This further increases the number of model runs required for sampling of the posterior parameter probability distribution. When it is recalled that the number of parameters employed by geophysical, groundwater and subsurface reservoir models on the one hand, and surface water quality and land use models on the other hand, may number in the thousands, use of MCMC in these contexts is seen to be somewhat problematic.

An obvious way to speed up the MCMC process is to reduce the number of parameters employed by a model. In some situations this may be appropriate. In other situations it may be entirely inappropriate for any of the following reasons.

- Omission of parameters from the MCMC process may artificially reduce the perceived uncertainties of some model predictions of management interest. This will be particularly the case for predictions which are sensitive to parameters which lie in the null space of the model operator \mathbf{Z} . These are the parameters that a modeller may be tempted to omit from MCMC analysis as they have little or no effect on model outcomes that correspond to the history-matching dataset \mathbf{h} .
- Where parameter simplification is not confined to the null space of \mathbf{Z} it may introduce “structural noise” to model outputs that correspond to \mathbf{h} , thereby compromising the level of fit that can be attained between these outputs and field measurements of system behaviour. The history-matching process is then further complicated by the fact that the statistics of structural noise are unknown, and that its covariance matrix is probably singular; see Doherty and Welter (2010) and chapter 15 of the present text for further details. The likelihood function then becomes more difficult to define and calculate. Choice of the likelihood function may then entail a high degree of subjectivity; see, for example, Beven (2005), Beven et al (2008) and literature on Approximate Bayesian Computation such as Vrugt and Bevin (2018), Kavetski, et al (2018) and references cited therein.
- Unless parameter simplification is “optimal” (a topic that will be addressed in chapter 15) some parameters will adopt compensatory roles as they are adjusted to ensure goodness of model-to-measurement fit during the history-matching process. Under these circumstances the history-matching process (whether it is undertaken using MCMC or by other means) can introduce substantial bias to some parameters, and possibly to some predictions. Both linear (White et al, 2014; Watson et al, 2013) and nonlinear analysis (Doherty and Christensen, 2011) demonstrate that the potential for simplification-induced predictive bias can be very large indeed. This too, is discussed in chapter 15.

Other strategies may be put in place to reduce the run time burden of MCMC analysis. The literature records use of surrogate, fast-running, models in MCMC analysis in place of

complex, slow-running, but more physically/chemically “realistic” models of environmental processes. Included in these surrogate models are stochastic surrogate-to-real-model correction terms which encapsulate their inadequacies as numerical simulator substitutes. See for example Cui and O’Sullivan (2011). While strategies such as these have met with some success, to the author’s knowledge, parameter numbers involved in surrogate model construction and consequential MCMC analysis remain small.

6.3.3 Ensemble Methods

Ensemble methods are discussed extensively later in this book. These are designed for use in high dimensional parameter spaces. Furthermore, they are able to sample posterior parameter probability distributions with high levels of model run efficiency. So they are not challenged by long model run times, unless these run times are excessively long. Their parameter adjustment algorithms are amenable to model run parallelization. This further increases their utility when deployed in conjunction with complex groundwater and other environmental models.

The literature on ensemble methods is extensive. Also, there are many variants of them, all of them having their roots in Bayes equation. In the ensemble literature, the term “data assimilation” is almost universally used to denote implementation of Bayes equation using ensemble methods.

Most ensemble methods that are used in environmental model history-matching today are an offshoot of the Kalman filter that was originally developed in the 1960s, and quickly gained utility as a navigational tool for aircraft and spacecraft because of its ability to rapidly estimate system states based on information extracted from noisy datasets. Ensembles were introduced to Kalman filtering in the 1990s. Since then, the Ensemble Kalman Filter (EnKF) has been extensively used in conjunction with large models that simulate oceanic and atmospheric processes, operation of petroleum reservoirs, and more. Through use of ensembles, the original Kalman filter has been adapted to estimate hundreds of thousands of states and parameters used by large numerical simulators.

Where a model is linear with respect to its states and parameters, and where prior parameter probability distributions and measurement noise are multiGaussian, ensemble methods comprise a direct implementation of Bayes equation. Unfortunately, these conditions are violated in most practical environmental applications. However this should not be seen as a matter of great concern, for numerical simulation and formulation of prior probability distributions for the properties of complex subsurface media are themselves accompanied by many approximations. So if quantification of posterior parameter and predictive uncertainties is indicative rather than exact, this should not be seen as a problem. In contrast, problems may arise through introduction of predictive bias in some complex modelling contexts. Hence a modeller should always be watchful for this occurrence. See Evensen et al (2022) for a comprehensive discussion of ensemble methods. See also later chapters of this book.

As stated above, ensemble Kalman methods are often used to estimate system states as well as system properties. In these applications, data assimilation is generally sequential and is repeated over many time steps, sometimes in real time. As time progresses, more and more data are assimilated so that better estimates of system states and properties emerge. It is sometimes found that estimates of system properties change over time, perhaps compensating for model inadequacies. However when predictions of system state are required for only a short time in advance, data-driven parameter compensation can increase predictive alacrity.

The term “ensemble smoother” is used to describe non-sequential, ensemble-based data assimilation that is generally applied to estimation of parameters only. The theory and use of the PESTPP-IES iterative ensemble smoother is discussed in chapter 12 of this book. See also White (2018).

The ensemble smoothing process begins by sampling the prior probability distribution of model parameters. Normally only a few hundred samples are taken, in spite of the fact that parameters may number in the tens or hundreds of thousands. Each of these samples is referred to as a “parameter realization”. The model is then run once for each realization. Model outputs that correspond to field measurements are monitored. An empirical joint covariance matrix that links parameters to these outputs is built using standard statistical formulas. Parameter estimates are then conditioned based on field measurements using concepts that are discussed in section 3.8. Algorithms that are discussed in chapter 12 thereby morph samples of the prior parameter probability distribution into samples of the posterior parameter probability distribution. This parameter adjustment process is normally iterative; data assimilation is repeated until a set of parameter realizations is derived which allow pertinent model outputs to replicate field measurements to levels that are commensurate with measurement noise. If a model prediction of management interest is made using all of these samples, the posterior probability distribution of that prediction is thereby sampled.

The attraction of ensemble methods is obvious. Their deployment has a relatively low computational cost. Ideally, outcomes of their deployment are samples of posterior parameter and predictive probability distributions. Though not strictly Bayesian where the relationships between model outputs and parameters is nonlinear, they are “Bayesian enough” to be useful for prediction uncertainty reduction and quantification. They are being increasingly seen as a vital ingredient of modern-day decision-support groundwater modelling as uncertainty assessments that they provide enable decision-maker awareness of risks that may accompany adoption of contemplated environmental management strategies. The importance of risk-assessment in decision-support modelling has been discussed by (among others) Freeze et al (1990) and Doherty and Simmons (2013).

As for all methods, strengths are accompanied by weaknesses. Sometimes attainment of a high level of model-to-measurement fit may require many iterations of the ensemble smoothing process, and hence many model runs. History-matching alacrity may be challenged in highly nonlinear contexts, and in contexts where a modeller’s notion of the prior parameter probability distribution is misaligned with hydrogeological reality. Where model-to-measurement misfit is unnecessarily high, important information may be excluded from the data assimilation process, or may be misdirected to inappropriate parameter receptacles. Because ensemble methods work in a reduced dimensional parameter space (see later in this book), the history-matching process may introduce parameter and predictive bias. At the same time, use of parameter fields which are both multiGaussian and continuously adjustable may require a somewhat abstract representation of subsurface heterogeneity.

All of this must be seen in context, however. All aspects of numerical simulation are replete with errors and approximations. The potential for direct or history-match-induced predictive bias is everywhere.

In addition to theoretical/conceptual challenges, use of ensemble methods presents some practical challenges. Modelling and model-ancillary software must enable storage and management of many realizations of model parameter fields and model outputs. Presentation and explanation of modelling results becomes more complex where a multiplicity of results is provided. Conversations with stakeholders may become longer.

6.3.4 Bypassing Parameters Altogether

A notable feature of ensemble methods is their versatility. If a covariance matrix can be built, then it can be conditioned. Posterior expected values of some of the random variables that are described by this matrix can thereby be obtained. Their posterior probability distributions can be sampled. These random variables do not need to be parameters; they can be model predictions.

Data space inversion (DSI) has already been briefly mentioned. It is mentioned again here and described in detail in chapter 14 of this book.

As for other ensemble methods, implementation of DSI begins with sampling of the prior parameter probability distribution. However parameter fields can be more complex than for other ensemble methods such as the iterative ensemble smoother (IES) because no adjustment of parameters is required. Hydraulic property fields with which a model is endowed can therefore be categorical; they do not need to be multiGaussian. In fact, as they are not actually adjusted, they do not need to be considered as model parameters.

The model is run once for each realization of hydraulic properties. (Actually, these realizations can include any model inputs that are not completely certain and hence require stochastic characterization.) The model must be run over the historical period during which measurements of system behaviour were made. It must also be run into the future so that it can make predictions pertaining to system management. Through this means, an empirical covariance matrix can be constructed that links model outputs that correspond to the measured past to those that correspond to the managed future. This matrix can then undergo conditioning based on actual field measurements. The outcomes of this conditioning process are posterior estimates of future system behaviour, together with posterior uncertainties associated with these estimates. See Sun and Durlofsky (2017), Lima et al (2020), Dellottier et al (2023), Kitlasten et al (2025) and chapter 14 of this book for more details.

The advantages of DSI are many. Its implementation generally requires only a few hundred model runs. Meanwhile, it can accommodate model hydraulic property fields of arbitrary complexity. In many modelling circumstances, its predictive integrity may be greater than that of other ensemble methods as parameter adjustment is not required. This is because relationships between model outputs that correspond to the past and those that correspond to the future are often less nonlinear than relationships between model outputs and parameters. The potential for history-match-induced predictive bias is thereby reduced.

Of course, there are disadvantages too. Successful implementation of DSI requires that the conceptual model on which the numerical model rests is reasonably good, and that inadequacies in these concepts are not exposed by field measurements. That is, it requires that so-called “prior-data conflict” be reasonably low. Fortunately, prior-data conflict is easily analysed as an adjunct to implementation of DSI; see section 14.5. Furthermore, generation of hydraulic property realizations that are used by DSI to build the past-to-future covariance matrix can readily accommodate uncertainties in the prior itself, including uncertain aspects of model structure.

A conceptual disadvantage of DSI is that while it may happily expose the range of predictive possibilities that are compatible with site concepts and site data, it cannot accompany these possibilities with corresponding parameter fields. So while DSI-based uncertainty analysis may demonstrate to modellers and decision-stakeholders that unwanted management occurrences cannot be precluded, it cannot show them hydraulic property conditions that precipitate these occurrences.

6.3.5 Summary

This brief discussion of model history-matching has demonstrated the centrality of Bayes equation to this process. It has also shown that history-matching is best understood in terms of prior and posterior probability distributions. Prior probabilities pertain to parameters. Posterior probabilities pertain to parameters and/or predictions.

A number of commonly used methods for accomplishing posterior parameter and predictive uncertainty analysis have been discussed. Their level of “Bayesian purity” varies. Those that can lay claim to Bayesian purity are generally not suitable for use in groundwater model history-matching where parameter numbers are large and model run times are long. However they are “Bayesian enough” to be useful.

We now turn our attention to the subject of model calibration. In doing so, our first task is to explain what exactly model calibration is. Our second task is to provide a justification for why it should be done at all.

6.4 Model Calibration

6.4.1 Calibration vs Bayes Theorem

Bayes equation states that parameter and predictive uncertainty prevail both before and after history-matching, and that the history-matching process may (or may not) reduce the uncertainties of management-critical predictions below their pre-history-matching levels. It would appear that Bayesian analysis is therefore mandatory.

So how does “model calibration” fit into Bayesian analysis? And what exactly is “model calibration” anyway? And why does the notion of the “calibrated model” carry such an aura of finality – as though the model construction process can be considered as complete once a model has been calibrated? Sadly, these questions can only be answered through recourse to folklore and wishful thinking; mathematics provides little assistance. We note, however, that stakeholder expectations of modelling are changing fast. The aura that surrounds the “calibrated” model is not as bright as it used to be.

As we shall now discuss, model calibration can serve a valid decision-support purpose. However this purpose, and the notion of calibration itself, must first be properly understood. Calibration must not be construed as implying certainty of parameters, nor of predictions made by a calibrated model. Nor should it diminish the fundamental truth embodied in Bayes equation that history-matching may reduce the uncertainties of some parameters but not others, and may reduce the uncertainties of some model predictions while leaving the uncertainties of other model predictions relatively unchanged.

From a Bayesian perspective, the term “model calibration” has no meaning. Nor should any conclusions regarding the predictive integrity of a model be drawn from the fact that it has been calibrated. If a model has been declared by its builder as being “well calibrated”, the only conclusion that can generally be drawn is that the model has been provided with a set of parameters for which pertinent model outputs match corresponding measurements of system behaviour reasonably well. However this tells us nothing about whether a very different set of parameters that are equally valid when judged at the bar of expert knowledge also allow the model to match measurements of system behaviour reasonably well.

In everyday usage, the term “calibration” (especially when applied to complex laboratory instruments) carries an implication of fine tuning that is necessary to achieve something approaching measurement perfection. Sometimes, it also carries an implication that small

manufacturing defects have been corrected. Application of these notions to groundwater modelling suggests that the calibration process can encourage model parameters to adopt values that compensate for model simplifications and imperfections that are essential for its construction. As will be shown in chapter 15, the matter of whether model defects can indeed be “calibrated out” is a complex one that may work for some predictions at the same time as it degrades the capacity of a calibrated model to make other predictions.

Most importantly, the notion of “calibration” implies uniqueness. However Bayes equation informs us that parameters and predictions are fundamentally stochastic quantities. Hence, at face value, the concept of uniqueness is irreconcilable with concepts that are expressed by Bayes equation. In truth however, as we shall see, calibration and Bayes equation are not incompatible. Furthermore, in highly-parameterized contexts, the same mathematics is required for both model calibration and Bayesian analysis.

Ultimately, decision-makers should be made aware of the range of predictive possibilities that are compatible with all available data. Some kind of predictive uncertainty analysis is required to answer the most important question with which decision-makers are faced. This question is “what can go wrong?” Whether an investigation into the possibility of unwanted management outcomes is conducted using strictly Bayesian principles, or whether it is based on an analysis that is illuminated by Bayesian concepts but does not rigorously implement Bayes equation, is of secondary importance to good decision-making. After all, Bayesian analysis has its problems when applied to predictions made by a groundwater model. Not least of these problems is definition of an adequate prior parameter probability distribution that respects all opportunities for hydraulic conductivity connectedness at a particular study site. Added to this are the numerous opportunities for predictive bias that accompany numerical implementation of Bayes equation in nonlinear modelling circumstances.

In the end, support for environmental management in general, and groundwater management in particular, may require a suite of analyses. There will be many occasions where inclusion of model calibration in these analyses provides insights that may be difficult to obtain in any other way.

6.4.2 What is Model Calibration?

Model calibration seeks a unique parameter field. It does this in violation of the concept that a groundwater model calibration dataset rarely possesses enough information to support unique estimation of its parameters unless parameterization is purposefully simplified to achieve this goal. However before deciding whether calibration is a worthwhile undertaking, and how it is best achieved, the metric by which model calibration should be judged must be defined.

There can be only one purpose for seeking parameter uniqueness. Calibration should pursue this purpose. This purpose is not that of ensuring that model predictions are correct, for this goal is unattainable. Instead, it should be that of ensuring that the potential for model predictive incorrectness is minimized. For a perfect simulator, this should apply to all model predictions. For an imperfect simulator, it may apply to a subset of model predictions, a matter to which we will return later in this book.

The potential for wrongness in the value of a stochastic variable is minimized if this value is central with respect to its probability distribution. This is its mean value. If the standard deviation of the random variable is large, then the potential of its mean value for wrongness may still be large. However this potential is symmetrical with respect to the mean. This is why it is minimized.

It follows that model calibration should seek a parameter field that endows a model with the ability to make predictions which are at (or close to) their posterior means. In a multiGaussian world where simulation is also perfect, there is only one parameter field that will achieve this for all model predictions. This is the parameter field that model calibration should seek. Furthermore, because only one such parameter field exists, the calibrated parameter field can claim uniqueness. The quest for uniqueness and the quest for decision-support utility therefore coincide.

As will be discussed in following chapters of this book, there are a number of ways in which a model can be calibrated. Their performances differ when judged by the metric that has just been provided. However numerical modelling of environmental processes is always approximate, and achievement of the decision-support potential of modelling always requires that many subjective judgements be made. So the goal that model calibration seeks is somewhat elusive. Nevertheless it is a goal that is worthy of pursuit. As such, it provides a basis for discussions on whether calibration of a particular model is as effective as intended, or whether it requires improvement.

6.4.3 Regularization

The term “regularization” has been used on a few occasions in this book so far, especially in association with the word “inversion”. It can now be given an exact meaning.

“Regularization” is a mathematical term. The term “regular” implies prevention of chaos. In some ways, this is what it does. It avoids parameter chaos by seeking uniqueness.

A numerical model calculates quantities of decision-support significance from properties of a natural system. As the name implies, the “inverse problem” works in the opposite direction. It attempts to infer system properties from measurements of system behaviour. As we have discussed, earth properties are generally so heterogeneous as to preclude this as a mathematical possibility if correctness is the metric by which solution of an inverse problem should be judged. Most inverse problems that are encountered in groundwater model calibration are “ill-posed” (to use a mathematical term). Therefore, while solution of an inverse problem may yield a parameter field that allows model outputs to match field data well, the parameter field that it yields may be just one of many parameter fields that do the same; there may be nothing “special” about this parameter field. And it will almost certainly be far from “correct”.

Regularized inversion seeks solution to an ill-posed inverse problem that is, in fact, special. It seeks a parameter field that supports model predictions of minimized error potential. In this book we use the term “minimum error variance parameter field” to describe this particular parameter field. This parameter field is such that all parameters, and all linear combinations of parameters (including model predictions in a linear world) exhibit this minimum error variance quality. It can be shown that there is only one parameter field that does this.

See figure 6.2 for an idealized prediction made by a calibrated model. This is virtually a repeat of figure 6.1. However a prediction \underline{s} made by the calibrated model has been added to the figure. Ideally, the value of this prediction should lie somewhere near the centre of the uncertainty band of the prediction. In this way, its potential for error is minimized. As has already been discussed, this potential may not be small; in fact it will be large if the posterior uncertainty of the prediction is large. Of course the situation becomes somewhat more complicated when the posterior predictive probability distribution is skewed because of model nonlinearity. However the principle of minimizing the error variance of decision-critical model predictions by ensuring that these predictions lie somewhere near the centres of their respective posterior

probability distributions remains intact.

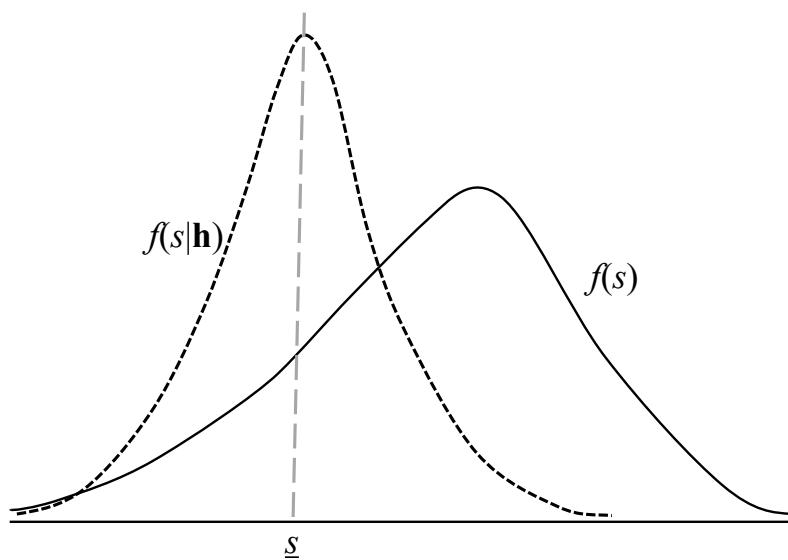


Figure 6.2. A prediction s of minimized error variance made by a calibrated model.

So model calibration and regularized inversion are the same thing. Perhaps it would have been less confusing to the groundwater industry if the term “calibration” had never been used in the first place to describe model parameter estimation. “Regularization” would have been a far better term because it is far better aligned with mathematical approaches to history-matching.

Three broad families of regularization methodologies are discussed in this book. All regularization methods that are used in practice are variants of one or more of these. In fact, most practical regularization methods that are used in groundwater model calibration are combinations of all three of them. They are:

1. manual regularization;
2. singular value decomposition;
3. Tikhonov regularization.

As we shall see in later chapters of this book, highly parameterized Bayesian methods actually use these same methods, but in different ways. This is because procedures that seek inverse problem solution diversity instead of uniqueness must not be chaotic. All of these solutions must be “realistic”. That is, as well as allowing pertinent model outputs to replicate field measurements of system behaviour, all of the parameter fields that emerge from Bayesian analysis must respect limits and correlations that are encapsulated in the prior parameter probability distribution.

6.4.4 So Why Calibrate?

The regularized parameter field that emerges from model calibration is often conceptualized as the “simplest” parameter field that is compatible with a calibration dataset. As we shall see, it can therefore be viewed as a kind of “average” parameter field. A similar parameter field (though not necessarily one that allows a model to fit field measurements as well as the calibrated parameter field) can often be obtained by averaging many samples of the posterior parameter probability distribution. Later in this book we will show that the calibrated parameter field is, in fact, a spatial average of unknown real-world hydraulic properties. We do not know

what these properties are; but the averaging function can be computed (for a linear system). The less data are available for model calibration, the more smoothing and smudging does this averaging process imply.

So in many modelling contexts, the calibrated parameter field may display very little spatial detail. However it will display as much detail as the calibration dataset can support – but no more. This can be seen as both problematic and beneficial. A modelling stakeholder should never view the calibrated parameter field as anything more than a grossly smoothed version of reality, or the “shadow” of reality cast on a small dimensional subset of parameter space (as we shall soon see). This is because calibration purposely seeks the heterogeneity that **MUST** exist to explain a calibration dataset. In contrast, Bayesian history-matching seeks the heterogeneity that **MAY** exist subject to respect for both the prior parameter probability distribution and measurements that comprise a history-matching dataset.

By acknowledging this, and by pursuing model calibration in ways that are described later in this book, a modeller may learn a lot about a study area by calibrating a model.

More often than not, a calibrated parameter field comes with some surprises. Perhaps heterogeneity has emerged with an unexpected amplitude at an unexpected place. Perhaps it adopts an unexpected shape. This challenges a modeller to review some or all of the concepts on which the current modelling process is based. It may mean that some parameters are adopting roles that compensate for defects in model construction and/or in some of its boundary conditions. The nature of aberrant heterogeneity may suggest what these defects are. After all, this heterogeneity has emerged because it **MUST** emerge in order for model outputs to fit the calibration dataset. On the other hand, the calibration process may have exposed the existence of subsurface heterogeneity of which a modeller was previously unaware. In doing so, it has “surprised” the modeller (Bredehoeft, 2005). In contrast, information (particularly subtle information) that is gleaned from a history-matching dataset may not be as clearly visible where Bayesian inversion is asked to yield hundreds of parameter fields, none of which are of minimum error variance, and hence of maximum data-compliant simplicity.

As will be extensively discussed later in this book, model calibration requires calculation of a Jacobian matrix (i.e. a sensitivity matrix). This matrix can be used for other purposes as well, including rudimentary uncertainty and data worth analysis. It can also be used to gain insights into why a model is currently incapable of fitting two different field measurements that require different alterations to the same parameters. Utility programs which implement a variety of linear analyses based on a Jacobian matrix are available through both the PEST and PyEMU suites.

Filling of a Jacobian matrix may require many model runs. For highly parameterized models, model calibration may therefore incur a higher numerical burden than ensemble-based Bayesian analysis. However, model-to-measurement fits that are attained through manipulation of a Jacobian matrix are often better than can be attained through manipulation of covariance matrices that are used by ensemble methods. Furthermore, methodologies such as Broyden Jacobian updating (see later) can improve the integrity of a Jacobian matrix “on the fly” as parameters undergo adjustment. This can greatly expedite attainment of a high level of model-to-measurement fit in nonlinear inversion contexts. Once such a fit has been attained, a solid platform has been laid for ensuing implementation of Bayesian analysis. This analysis can then begin its exploration of parameter space from a place that is (hopefully) central with respect to the posterior parameter probability distribution.

Calculation and presentation of the “calibration parameter field” can yield social benefits in

addition to mathematical and computational benefits. In most decision-support modelling circumstances, didactic outcomes of modelling are as important as numerical outcomes. Stakeholders generally want to understand why certain management choices are made. When demonstrating management-salient aspects of system behaviour in ways that can be understood by non-modellers, it may sometimes be better to employ a parameter field that possesses no more complexity than that which is required to match available data than a suite of more complex parameter fields that are used in support of more sophisticated probabilistic analysis.

In summary, model calibration can provide conceptual and computational benefits that are difficult to achieve in other ways.

We close this subsection by noting that, until recently, it was not uncommon to encounter circumstances where model calibration was impractical because of the high computational cost of filling a large Jacobian matrix using finite differences to calculate partial derivatives that are embodied in this matrix. History-matching was, therefore, unachievable through any other means than ensemble-based Bayesian analysis. This has changed with the advent of ensemble space inversion (ENSI) as a calibration option. ENSI attains the same computational speed as ensemble methods by working in a relatively low dimensional subspace of parameter space that is spanned by a set of parameter realizations that is sampled from the prior parameter probability distribution. At the same time, it attains the conceptual and practical advantages of model calibration. See chapter 9 of this book

6.5 Model Linearization

6.5.1 The Jacobian Matrix

Implementation of highly-parameterized history-matching generally requires that a model be temporarily linearized. The operation of a model on its parameters is therefore emulated by matrix multiplication. The matrix is the Jacobian matrix.

Equation (6.2.1) represents the action of a model on its parameters under calibration conditions, that is under the environmental conditions that prevailed when measurements \mathbf{h} comprising the calibration dataset were acquired. For the n individual model outputs o_i comprising the model output vector \mathbf{o} , equation (6.2.1) can be written as

$$o_1 = z_1[\mathbf{k}] \quad (6.5.1a)$$

$$o_2 = z_2[\mathbf{k}] \quad (6.5.1b)$$

...

$$o_n = z_n[\mathbf{k}] \quad (6.5.1c)$$

Using Taylor's theorem, equation (6.2.1) can be expanded as

$$\mathbf{o} = \mathbf{o}_0 + \mathbf{J}\delta\mathbf{k} + \text{other terms} \quad (6.5.2)$$

where

$$\mathbf{o}_0 = Z[\mathbf{k}_0] \quad (6.5.3)$$

$$\delta\mathbf{k} = \mathbf{k} - \mathbf{k}_0 \quad (6.5.4)$$

and \mathbf{J} is the Jacobian matrix. This is a matrix of partial derivatives of the n model outputs which correspond to measurements comprising a history-matching dataset with respect to the m (adjustable) parameters employed by a model. Specifically, the Jacobian matrix is defined as

$$\mathbf{J} = \begin{bmatrix} \frac{\partial \mathbf{z}_1}{\partial k_1} & \frac{\partial \mathbf{z}_1}{\partial k_2} & \cdots & \frac{\partial \mathbf{z}_1}{\partial k_m} \\ \frac{\partial \mathbf{z}_2}{\partial k_1} & \frac{\partial \mathbf{z}_2}{\partial k_2} & \cdots & \frac{\partial \mathbf{z}_2}{\partial k_m} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \mathbf{z}_n}{\partial k_1} & \frac{\partial \mathbf{z}_n}{\partial k_2} & \cdots & \frac{\partial \mathbf{z}_n}{\partial k_m} \end{bmatrix} \quad (6.5.5)$$

The Jacobian matrix is often referred to as the “sensitivity matrix” for obvious reasons. The m elements comprising the i ’th row of a Jacobian matrix are derivatives of the model operator as it pertains to the i ’th model output; these derivatives are taken with respect to each of the m parameters used by the model. The j ’th column of the Jacobian matrix contains derivatives with respect to the j ’th parameter of all n model outputs.

In equations (6.5.2) to (6.5.4) \mathbf{k}_0 specifies a set of “base parameters” or “reference parameters” while \mathbf{o}_0 specifies model outputs calculated using these reference parameters. Where a model is temporarily linearized as part of an iterative, nonlinear parameter estimation or uncertainty analysis procedure, \mathbf{k}_0 changes from iteration to iteration of that procedure. The Jacobian matrix is re-calculated on every occasion that \mathbf{k}_0 changes. If a model is nonlinear, the values of elements of the Jacobian matrix change with \mathbf{k}_0 . By definition, if a model were truly linear, the Jacobian matrix would not be a function of \mathbf{k}_0 .

6.5.2 Linear Model Analysis

A linear model is a model that can be represented by a matrix; the values taken by elements of this matrix are independent of values taken by model parameters. The operator $\mathbf{Z}[\cdot]$ therefore becomes a matrix operator.

Rarely is an environmental model linear. However the assumption of near-linear behaviour for small variations of parameters from a set of reference values facilitates many useful analyses. It underpins algorithms that improve parameter values in successive iterations of a nonlinear inversion or uncertainty analysis procedure. The benefits of model linearization go well beyond this, however. Where reference parameter values \mathbf{k}_0 are obtained through model calibration, and presumably therefore represent minimum error variance estimates of a real world parameter set \mathbf{k} , reasonable estimates of the uncertainties of thus-estimated parameters, and of predictions that are sensitive to them, can be readily obtained if linear model behaviour in the vicinity of the estimated \mathbf{k}_0 parameter set is assumed. In addition to this, linear analysis can teach a modeller much about the benefits (or lack thereof) that are accrued through model calibration in a particular modelling context, and of the sources of parameter and predictive uncertainty that remain after calibration has taken place. Much will be said on this subject in later chapters of this book; see, in particular, chapter 10.

To facilitate analyses that are described throughout this text, and to make the equations that emerge from these analyses somewhat simpler, equation (6.5.2) is written as

$$\mathbf{o} = \mathbf{Z}\mathbf{k} \quad (6.5.6)$$

In equation (6.5.6), \mathbf{Z} replaces the Jacobian matrix; the model is thus assumed to be truly linear instead of just locally near-linear. Furthermore \mathbf{k} in equation (6.5.6) replaces $\delta\mathbf{k}$ of (6.5.2). It is thus equal to $\mathbf{k} - \mathbf{k}_0$. In many of the equations that are presented later in this book, \mathbf{k}_0 represents the pre-calibration expected value of \mathbf{k} . That is, \mathbf{k}_0 represents the prior mean of \mathbf{k} expressed as

$$\mathbf{k}_0 = E[\mathbf{k}] \quad (6.5.7)$$

\mathbf{k} of equation (6.5.6) thus represents departures of parameters from their pre-calibration expected values. As such, if pre-calibration uncertainty of parameters is expressed by the covariance matrix $C(\mathbf{k})$, this matrix is just as applicable to the \mathbf{k} of (6.5.6) as it is to the \mathbf{k} of (6.2.1).

Because

$$\mathbf{0} = \mathbf{Z}\mathbf{0} \quad (6.5.8)$$

\mathbf{o} of equation (6.5.6) represents departures of model outputs from $\mathbf{Z}\mathbf{k}_0$. If \mathbf{h} represents the observed values of these departures, then (6.2.2) becomes

$$\mathbf{h} = \mathbf{Z}\mathbf{k} + \boldsymbol{\varepsilon} \quad (6.5.9)$$

Measurement noise is presumably independent of \mathbf{k} . Furthermore its expected value is $\mathbf{0}$. That is

$$E[\boldsymbol{\varepsilon}] = \mathbf{0} \quad (6.5.10)$$

$\boldsymbol{\varepsilon}$ of equation (6.5.10) has the same statistical properties as $\boldsymbol{\varepsilon}$ of equation (6.2.2) and is thus characterized by the same covariance matrix $C(\boldsymbol{\varepsilon})$.

For a linear model, equation (6.2.6) becomes

$$s = \mathbf{y}^t \mathbf{k} \quad (6.5.11)$$

where the element y_i of the vector \mathbf{y} expresses the sensitivity of the prediction s to parameter k_i , this comprising the i 'th element of the parameter vector \mathbf{k} . If, as above, \mathbf{k} in equation (6.5.11) is considered to represent the departure of parameter values from their prior expected values, then the value of the prediction s represented by equation (6.5.11) represents the departure of this prediction from that which the model would make if provided with prior expected parameter values.

As has already been discussed, in most cases of environmental interest, \mathbf{k} will have many elements. It will probably cite more parameters than can be uniquely estimated. The matrix \mathbf{Z} may therefore possess a null space. In some cases \mathbf{k} will have more elements than \mathbf{h} ; if this is the case, the matrix \mathbf{Z} will be long in the horizontal direction and will certainly possess a null space. Linear model analysis is well equipped to explore the repercussions of this situation.

7. Manual Regularization

7.1 Formulation of Equations

In this chapter we start by assuming linear model behaviour. Equations for solution of the inverse problem will be formulated and then solved. The means through which nonlinear problems can be solved using the same theory in slightly modified form will then be explored.

We start with equation (6.5.9), which is now repeated.

$$\mathbf{h} = \mathbf{Zk} + \boldsymbol{\varepsilon} \quad (7.1.1)$$

Manual regularization achieves calibration uniqueness through parameter simplification. Suppose that the outcome of the simplification process is a reduced set of parameters \mathbf{p} . (The number of elements comprising the vector \mathbf{p} must be less than that comprising the vector \mathbf{h} or the inverse problem cannot possibly be well-posed.) Let it further be supposed that the elements of \mathbf{p} are defined through combining some elements of the larger parameter set \mathbf{k} into single elements of \mathbf{p} and through fixing other elements of \mathbf{k} at “known” values so that they do not need to be estimated at all. As a result, \mathbf{p} has fewer elements than \mathbf{k} . Hopefully, this leads to a well-posed (often referred to as “over-determined”) inverse problem through which elements of \mathbf{p} can be uniquely estimated. The attainment of uniqueness is, of course, the goal of regularization. As we shall see, however, it will often not be possible to judge the success or otherwise of a manual regularization enterprise until solution of the newly-posed inverse problem is attempted. Furthermore, as we shall also see, it will probably not ever be possible to judge the optimality or otherwise of the manual regularization process. (Recall the metric for optimal regularization presented in the previous chapter.)

The relationship between \mathbf{k} -parameters and \mathbf{p} -parameters can be expressed through the following matrix equation:

$$\mathbf{k} = \mathbf{Lp} \quad (7.1.2)$$

In equation (7.1.2) \mathbf{L} is a matrix that calculates values for elements of the complex parameter set \mathbf{k} from those of the simple parameter set \mathbf{p} . For example, if a 3-element vector \mathbf{p} represents the values of system properties in each of three zones of assumed (for the purpose of parameter estimation) piecewise constancy into which the domain of a spatial model has been subdivided, and if \mathbf{k} (a 13-element vector) represents a more complex set of model parameters which characterize spatial heterogeneity at a finer scale, then equation (7.1.2) may be written as

$$\begin{bmatrix} k_1 \\ k_2 \\ k_3 \\ k_4 \\ k_5 \\ k_6 \\ k_7 \\ k_9 \\ k_{10} \\ k_{11} \\ k_{12} \\ k_{13} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} p_1 \\ p_2 \\ p_3 \end{bmatrix} \quad (7.1.3)$$

In this example an element of \mathbf{k} simply adopts the value of the \mathbf{p} -zone in which it lies; \mathbf{L} is a selection matrix. If (7.1.2) is substituted into (7.1.1) we obtain

$$\mathbf{h} = \mathbf{ZLp} + \boldsymbol{\varepsilon} \quad (7.1.4)$$

If \mathbf{X} is now defined as

$$\mathbf{X} = \mathbf{ZL} \quad (7.1.5)$$

then

$$\mathbf{h} = \mathbf{Xp} + \boldsymbol{\varepsilon} \quad (7.1.6)$$

We now have a modified equation that represents the action of the model under calibration conditions. However this equation uses a simplified parameter set \mathbf{p} instead of the more complex parameter set \mathbf{k} . If \mathbf{X} does not have a null space, then each member p_i of the reduced parameter set \mathbf{p} is, conceptually at least, uniquely estimable. Let us denote the estimated parameter set as $\underline{\mathbf{p}}$.

Equation (7.1.2) allows us to calculate an estimated \mathbf{k} (i.e. $\underline{\mathbf{k}}$) from an estimated \mathbf{p} (i.e. $\underline{\mathbf{p}}$). It is important to note, however, that even if $\underline{\mathbf{p}}$ is uniquely estimable from \mathbf{h} , the same does not apply to $\underline{\mathbf{k}}$. Hence the $\underline{\mathbf{k}}$ that emerges from solution of the manually-regularized inverse problem may be very different from the correct \mathbf{k} ; it will be unique, but probably incorrect. $\underline{\mathbf{k}}$ will approach the true \mathbf{k} only if there is no variability of properties within zones of assumed homogeneity that are used to define \mathbf{p} .

A modeller may undertake manual regularization when he/she is presented with an inverse problem that he/she perceives to be ill-posed. However it is important that parameterization simplification (defined by the matrix \mathbf{L}) that defines this regularization strategy be strategic. If possible, manual regularization should not compromise a model's ability to fit a calibration dataset (or should do so as little as possible). In fact, if parameter simplification induces too much model-to-measurement misfit, then it can be concluded that more parameters than are offered to the inversion process through the vector \mathbf{p} are in fact estimable.

Ideally, a parameter simplification strategy \mathbf{L} is sought for which the following equation holds:

$$\mathbf{ZLp} - \mathbf{Xp} = (\mathbf{ZL} - \mathbf{X})\mathbf{p} = \mathbf{0} \quad (7.1.7)$$

An \mathbf{L} that achieves (7.1.7) has no impact on model-to-measurement misfit.

However even if an \mathbf{L} can be found for which (7.1.7) holds true, another important question arises. While (7.1.2) can be used for computation of $\underline{\mathbf{k}}$ from $\underline{\mathbf{p}}$, what exactly is the relationship between the estimated $\underline{\mathbf{k}}$ and the true \mathbf{k} , and between the estimated $\underline{\mathbf{p}}$ and the true \mathbf{k} ? Can $\underline{\mathbf{p}}$ be construed to be an estimate of the average of \mathbf{k} within each \mathbf{p} -zone which encapsulates multiple \mathbf{k} parameters? The answer to this questions is "no", unless special measures are taken in an effort to achieve this status for $\underline{\mathbf{p}}$. These measures are described in Cooley (2004) and Cooley and Christensen (2006). They will be discussed later in this chapter; so too will the nature of the relationships between $\underline{\mathbf{k}}$ and \mathbf{k} and between $\underline{\mathbf{p}}$ and \mathbf{k} .

In most cases, an \mathbf{L} cannot be found for which (7.1.7) holds exactly. Hence the outcomes of manual regularization are visible to the inversion process as simplification-induced model-to-measurement misfit. Starting from (7.1.1)

$$\mathbf{h} = \mathbf{Zk} + \boldsymbol{\varepsilon}$$

$$\mathbf{h} = \mathbf{Zk} - \mathbf{Xp} + \mathbf{Xp} + \boldsymbol{\varepsilon}$$

$$\mathbf{h} = \mathbf{Xp} + (\mathbf{ZL} - \mathbf{X})\mathbf{p} + \boldsymbol{\varepsilon}$$

$$\mathbf{h} = \mathbf{Xp} + \boldsymbol{\eta} + \boldsymbol{\varepsilon} \quad (7.1.8)$$

$$\mathbf{h} = \mathbf{Xp} + \boldsymbol{\tau} \quad (7.1.9)$$

where

$$\boldsymbol{\eta} = (\mathbf{ZL} - \mathbf{X})\mathbf{p} \quad (7.1.10)$$

and

$$\boldsymbol{\tau} = \boldsymbol{\eta} + \boldsymbol{\varepsilon} \quad (7.1.11)$$

Equation (7.1.9) is the same as (7.1.6) except for the fact that it acknowledges the presence of simplification-induced noise $\boldsymbol{\eta}$ in addition to measurement noise $\boldsymbol{\varepsilon}$. Continuing with the above zonal example, this “structural noise” $\boldsymbol{\eta}$ will only be zero if model outputs under calibration conditions are completely insensitive to any heterogeneity which exists within each zone of assumed piecewise constancy which underpins definition of \mathbf{p} ; that is, structural noise will only be zero if equation (7.1.7) holds. Regardless of this, equation (7.1.9) defines the inverse problem that must be solved for \mathbf{p} .

The logic above is equally applicable to other forms of parameter simplification such as the fixing of some parameters at values that are based on expert knowledge and the estimation of others. In all cases, it is possible to formulate an equation which describes an inverse problem which features a simplified parameter set \mathbf{p} in place of a more complex parameter set \mathbf{k} . Ideally, solution of this problem provides a unique \mathbf{p} from which a unique \mathbf{k} can be calculated using an appropriate \mathbf{L} matrix which defines parameter simplification. However while \mathbf{p} may be unique, its relationship to real system properties \mathbf{k} remains unknown; so too does the relationship between estimates of \mathbf{k} achieved through the manual regularization process and real system properties \mathbf{k} . As will be discussed, this makes analysis of parameter and predictive error difficult, if not impossible, in the manually-regularized context.

As demonstrated above, structural noise can be introduced to the inverse problem through the parameter simplification process. The existence of this noise has a number of consequences. Firstly, it may limit the fit between model outcomes and field measurements that is achieved through the parameter estimation process. Some of the information that is hosted by the calibration dataset therefore does not get to inform model parameters. What is worse is that while we may know something about the statistics of measurement noise, generally we know nothing about the statistics of structural noise. It thus becomes very difficult to establish the potential for error in \mathbf{p} as an estimate of \mathbf{p} . This further compounds the problem of assessing post-calibration parameter and predictive uncertainty.

These problems will be ignored for a while as we focus on solution of the inverse problem that manual regularization delivers to us.

7.2 Well-Posed Inverse Problem

7.2.1 Data without Noise

Let us put aside for the moment the fact that, in the environmental modelling context, well-posed inverse problems rarely happen naturally but are the outcomes of regularization. Instead, let us suppose that the parameters of the real world are as simple as those employed by a model. Hence simplification-induced structural noise is zero. At the same time, questions pertaining to the relationship between estimated parameters and real-world parameters can be dispensed with. In doing this, critical questions related to parameter and predictive uncertainty can also be put to one side – for the moment at least.

We start with equation (7.1.9) but replace τ with ϵ . Thus

$$\mathbf{h} = \mathbf{Xp} + \boldsymbol{\epsilon} \quad (7.2.1)$$

If there is no noise associated with the measurement dataset, (7.2.1) becomes

$$\mathbf{h} = \mathbf{Xp} \quad (7.2.2)$$

While this equation is simpler than (7.2.1) it still cannot be solved for \mathbf{p} as \mathbf{X} is not a square matrix. Instead, it is a matrix with more rows than columns; as such, it does not have an inverse.

Let us now multiply both sides of (7.2.2) by \mathbf{X}^t . Then

$$\mathbf{X}^t \mathbf{h} = \mathbf{X}^t \mathbf{Xp} \quad (7.2.3)$$

$\mathbf{X}^t \mathbf{X}$ is a square matrix. If \mathbf{X} does not have a null space, $\mathbf{X}^t \mathbf{X}$ is invertible. We can thus multiply both sides of (7.2.3) by $(\mathbf{X}^t \mathbf{X})^{-1}$ to obtain \mathbf{p} .

$$(\mathbf{X}^t \mathbf{X})^{-1} \mathbf{X}^t \mathbf{h} = (\mathbf{X}^t \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Xp} = \mathbf{p} \quad (7.2.4)$$

If, instead of multiplying both sides of (7.2.2) by \mathbf{X}^t and then by $(\mathbf{X}^t \mathbf{X})^{-1}$, we had multiplied by $\mathbf{X}^t \mathbf{Q}$ and then $(\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1}$ where \mathbf{Q} is a square matrix of full rank so that $(\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1}$ exists, we would have obtained

$$(\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \mathbf{h} = \mathbf{p} \quad (7.2.5)$$

The assumption of no noise implies a perfect fit between model outputs and a calibration dataset. An underscore is not used to distinguish estimated \mathbf{p} from true \mathbf{p} as was done in the preceding section. Values for \mathbf{p} obtained through (7.2.4) and (7.2.5) are identical; both are equal to the true \mathbf{p} .

7.2.2 Data with Noise

We now acknowledge the presence of noise in the calibration dataset. It will be seen that equations derived in the previous section can still be used for estimation of \mathbf{p} , notwithstanding noise contamination of the calibration dataset. However the estimate of \mathbf{p} obtained through these equations is no longer the true \mathbf{p} ; as such it is denoted as \mathbf{p} . This estimate has a potential for error. A task that awaits us is to quantify this potential.

Before proceeding further, equation (7.2.1) will be transformed slightly to a form that ensures estimates of \mathbf{p} that are of minimum error variance. Recall from the previous chapter that this is the metric for good model calibration. Let us suppose that measurement noise can be characterized by the covariance matrix $\mathbf{C}(\boldsymbol{\epsilon})$. As such, $\mathbf{C}(\boldsymbol{\epsilon})$ is a positive-definite matrix; hence it possesses an inverse. Let a matrix \mathbf{Q} be defined as being proportional to this inverse, with the proportionality constant being denoted as σ_r^2 . Then

$$\mathbf{Q} = \sigma_r^2 \mathbf{C}^{-1}(\boldsymbol{\epsilon}) \quad (7.2.6)$$

Note that if $\mathbf{C}(\boldsymbol{\epsilon})$ is diagonal, then so too is \mathbf{Q} ; it is also invertible and hence of full rank. \mathbf{Q} is referred to as the “weight matrix”. σ_r^2 is variously referred to as the “reference variance” and “variance of unit weight”.

We now define the vector of residuals \mathbf{r} as that which lists differences between model outputs and corresponding field measurements. Thus

$$\mathbf{r} = \mathbf{h} - \mathbf{Xp} \quad (7.2.7)$$

A so-called objective function Φ can be formulated as the sum of weighted squared residuals. Thus

$$\Phi = \mathbf{r}^t \mathbf{Q} \mathbf{r} = (\mathbf{h} - \mathbf{X} \mathbf{p})^t \mathbf{Q} (\mathbf{h} - \mathbf{X} \mathbf{p}) \quad (7.2.8)$$

The lower is the objective function, the better is the fit between model outputs and field data. Notice the similarity of this definition of the objective function to the exponent appearing in the equation for the multiGaussian (i.e. multinormal) density function; see equation (3.6.1). If measurement noise is characterized by a multiGaussian distribution, then minimization of the objective function defined by (7.2.8) is the same as maximization of the likelihood of residuals which describe model-to-measurement misfit.

We now expand (7.2.8) as

$$\begin{aligned} \Phi &= \mathbf{h}^t \mathbf{Q} \mathbf{h} - \mathbf{h}^t \mathbf{Q} \mathbf{X} \mathbf{p} - \mathbf{p}^t \mathbf{X}^t \mathbf{Q} \mathbf{h} + \mathbf{p}^t \mathbf{X}^t \mathbf{X} \mathbf{p} \\ &= \mathbf{h}^t \mathbf{Q} \mathbf{h} - 2 \mathbf{p}^t \mathbf{X}^t \mathbf{Q} \mathbf{h} + \mathbf{p}^t \mathbf{X}^t \mathbf{X} \mathbf{p} \end{aligned} \quad (7.2.9)$$

To find the minimum of the objective function, we differentiate Φ by each element of \mathbf{p} and set each derivative to zero. Thus

$$\frac{d\Phi}{d\mathbf{p}} = \begin{bmatrix} \frac{\partial \Phi}{\partial p_1} \\ \frac{\partial \Phi}{\partial p_2} \\ \vdots \\ \frac{\partial \Phi}{\partial p_m} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad (7.2.10)$$

where m is the number of elements of \mathbf{p} . Obviously, the derivative with respect to \mathbf{p} of the first term of equation (7.2.9) is 0. We state without proof that

$$\frac{\partial (\mathbf{p}^t \mathbf{X}^t \mathbf{Q} \mathbf{h})}{\partial \mathbf{p}} = \mathbf{X}^t \mathbf{Q} \mathbf{h} \quad (7.2.11)$$

$$\frac{\partial (\mathbf{p}^t \mathbf{X}^t \mathbf{Q} \mathbf{X} \mathbf{p})}{\partial \mathbf{p}} = 2 \mathbf{X}^t \mathbf{Q} \mathbf{X} \quad (7.2.12)$$

Substituting these into (7.2.9) and rearranging we find

$$2 \mathbf{X}^t \mathbf{Q} \mathbf{X} \mathbf{p} = 2 \mathbf{X}^t \mathbf{Q} \mathbf{h} \quad (7.2.13)$$

From which

$$\underline{\mathbf{p}} = (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \mathbf{h} \quad (7.2.14)$$

This is the same equation as (7.2.5); it is comforting that solution of the inverse problem with noise taken into account is the same as that without noise. Use of (7.2.14) to solve the inverse problem constitutes use of the Gauss-Newton method for estimation of parameters.

We now explore the potential for error in the parameter set $\underline{\mathbf{p}}$ that is estimated using (7.2.14). If (7.2.1) is substituted into (7.2.14) we obtain

$$\underline{\mathbf{p}} = (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \mathbf{X} \mathbf{p} + (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \boldsymbol{\varepsilon} = \mathbf{p} + (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \boldsymbol{\varepsilon} \quad (7.2.15)$$

So that the error in estimated $\underline{\mathbf{p}}$ becomes

$$\underline{\mathbf{p}} - \mathbf{p} = (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \boldsymbol{\varepsilon} \quad (7.2.16)$$

If $\boldsymbol{\varepsilon}$ is zero, then so too is parameter error. Under any other circumstances we cannot know parameter error as we do not know $\boldsymbol{\varepsilon}$. However we can say something about the propensity for parameter error through calculation of its covariance matrix. Using equation (3.9.2) for propagation of covariance, we obtain for $\mathbf{C}(\underline{\mathbf{p}} - \mathbf{p})$, the covariance matrix of parameter error

$$\mathbf{C}(\underline{\mathbf{p}} - \mathbf{p}) = (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \mathbf{C}(\boldsymbol{\varepsilon}) \mathbf{Q} \mathbf{X} (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \quad (7.2.17)$$

Here we have used the fact that $\mathbf{X}^t \mathbf{Q} \mathbf{X}$ is a symmetric matrix and that its inverse is therefore symmetric; the transpose of a symmetric matrix is the original matrix. If the weight matrix \mathbf{Q} is chosen in accordance with (7.2.6), (7.2.17) becomes

$$C(\mathbf{p} - \mathbf{p}) = \sigma_r^2 (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \quad (7.2.18)$$

It can be shown (see for example Koch, 1999) that with this choice of \mathbf{Q} , \mathbf{p} satisfies our calibration metric. It is the so-called best linear unbiased estimator of \mathbf{p} ; as an estimate of \mathbf{p} it has less error variance than any other estimate. The above choice of \mathbf{Q} is critical to achieving this status as it minimizes the trace of $C(\mathbf{p} - \mathbf{p})$. This choice of \mathbf{Q} can also be shown to minimize the error variance of any linear combination of parameters calculated as

$$\underline{s} = \mathbf{w}^t \mathbf{p} \quad (7.2.19)$$

If the components of the vector \mathbf{w} comprise sensitivities of a model prediction to parameters \mathbf{p} , then equation (7.2.19) is the equation which is used to make that prediction using the calibrated model. The “real” value of the prediction is

$$s = \mathbf{w}^t \mathbf{p} \quad (7.2.20)$$

By subtracting (7.2.20) from (7.2.19), predictive error is calculated as

$$\underline{s} - s = \mathbf{w}^t (\mathbf{p} - \mathbf{p}) \quad (7.2.21)$$

From (3.9.2) the covariance matrix of predictive error, which for a scalar such as a single prediction is simply the variance of that error (recall that variance is the square of standard deviation), is then given by

$$\sigma_{\underline{s}-s}^2 = \mathbf{w}^t C(\mathbf{p} - \mathbf{p}) \mathbf{w} = \sigma_r^2 \mathbf{w}^t (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{w} \quad (7.2.22)$$

As stated above, this is minimized by choosing \mathbf{Q} according to (7.2.6).

Let Φ_0 represent the minimized value of the objective function defined according to equation (7.2.8). An unbiased estimate $\underline{\sigma}_r^2$ of the reference variance σ_r^2 can be calculated from this minimized objective function as

$$\underline{\sigma}_r^2 = \frac{\Phi_0}{n-m} \quad (7.2.23)$$

where n is the number of observations featured in the calibration dataset (i.e. the number of elements of \mathbf{h}) and m is the number of estimated parameters (i.e. number of elements of \mathbf{p}).

If the reference variance σ_r^2 does not need to be estimated from the minimized objective function because it is known in advance, and if measurement noise has a multinormal distribution, so too does $\mathbf{p} - \mathbf{p}$ (if the model is linear). The same applies to $\underline{s} - s$ for any prediction s made by the model. However if $\underline{\sigma}_r^2$ needs to be estimated using (7.2.23), then parameter and predictive error have Student’s t distributions with $n-m$ degrees of freedom. Meanwhile $\underline{\sigma}_r^2$ itself has a chi-squared distribution with $n-m$ degrees of freedom. See Koch (1999) for details.

Before completing this section, it is worth pointing out that equation (7.2.1) is often transformed in order to make some of the above equations a little less cumbersome. If all terms of (7.2.1) are pre-multiplied by $\mathbf{Q}^{1/2}$, where \mathbf{Q} is chosen according to (7.2.6), we obtain

$$\mathbf{Q}^{1/2} \mathbf{h} = \mathbf{Q}^{1/2} \mathbf{X} \mathbf{p} + \mathbf{Q}^{1/2} \boldsymbol{\xi} \quad (7.2.24)$$

This can be written as

$$\mathbf{g} = \mathbf{Y} \mathbf{p} + \boldsymbol{\xi} \quad (7.2.25)$$

where

$$\mathbf{g} = \mathbf{Q}^{\frac{1}{2}}\mathbf{h} \quad (7.2.26)$$

$$\mathbf{Y} = \mathbf{Q}^{\frac{1}{2}}\mathbf{X} \quad (7.2.27)$$

$$\boldsymbol{\xi} = \mathbf{Q}^{\frac{1}{2}}\boldsymbol{\epsilon} \quad (7.2.28)$$

Using (3.9.2) and (7.2.6) the covariance matrix $C(\boldsymbol{\xi})$ of $\boldsymbol{\xi}$ can be written as

$$C(\boldsymbol{\xi}) = \mathbf{Q}^{\frac{1}{2}}C(\boldsymbol{\epsilon})\mathbf{Q}^{\frac{1}{2}} = \sigma_r^2\mathbf{I} \quad (7.2.29)$$

\mathbf{p} can then be estimated as

$$\mathbf{p} = (\mathbf{Y}^t\mathbf{Y})^{-1}\mathbf{Y}^t\mathbf{g} \quad (7.2.30)$$

Meanwhile the covariance matrix of parameter error is

$$C(\mathbf{p} - \mathbf{p}) = \sigma_r^2(\mathbf{Y}^t\mathbf{Y})^{-1} \quad (7.2.31)$$

Substitution of (7.2.26) and (7.2.27) into (7.2.30) and (7.2.31) readily confirms that the latter two equations are equivalent to (7.2.14) and (7.2.18).

7.2.3 Prior Information

To increase the estimability of model parameters, measurements of system state can be supplemented by “observations” that pertain directly to these parameters. Such observations are often referred to as “prior information”. These supplementary observations may be comprised of direct measurements of system properties which the parameters represent. Alternatively, they may express expert knowledge of likely relationships between parameter values. Let the vector \mathbf{b} designate the values of these measurements, and let the matrix \mathbf{B} designate the relationship of these observations to system parameters. For example, if the extent of prior information that is included in the parameter estimation process is a single observation that the difference between p_i and p_j is likely to be 2.0, then \mathbf{b} will have a single element (this being 2.0), while the matrix \mathbf{B} will be a $1 \times m$ matrix, where m is the number of elements in the parameter vector \mathbf{p} . The i ’th element of the single row of \mathbf{B} will be 1 and the j ’th element of its single row will be -1.

Prior information can be included in the inversion process simply by adding elements to the \mathbf{h} vector and rows to the \mathbf{X} matrix as appropriate. However it is sometimes useful to separate terms which pertain to prior information from those which pertain to measurements of system state so that these two aspects of the calibration dataset can be given separate treatment if desired. This is especially useful where prior information is used to express Tikhonov regularization; see the next chapter.

With specific mention made of prior information, equation (7.2.1) can be written as

$$\begin{bmatrix} \mathbf{h} \\ \mathbf{b} \end{bmatrix} = \begin{bmatrix} \mathbf{X} \\ \mathbf{B} \end{bmatrix} \mathbf{p} + \begin{bmatrix} \boldsymbol{\epsilon} \\ \boldsymbol{\zeta} \end{bmatrix} \quad (7.2.32)$$

In equation (7.2.32) $\boldsymbol{\zeta}$ is the “noise” associated with \mathbf{b} . If the value of \mathbf{b} is based on expert knowledge of \mathbf{p} (from which it is calculated) then this “noise” arises from prior uncertainty of \mathbf{p} . If this uncertainty is characterized by the covariance matrix $C(\mathbf{p})$, then

$$C(\boldsymbol{\zeta}) = \mathbf{B}C(\mathbf{p})\mathbf{B}^t \quad (7.2.33)$$

Equation (7.2.14), when applied to (7.2.32), becomes

$$\underline{\mathbf{p}} = \left\{ [\mathbf{X}^t \quad \mathbf{B}^t] \begin{bmatrix} \mathbf{Q}_w & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_p \end{bmatrix} [\mathbf{X}] \right\}^{-1} [\mathbf{X}^t \quad \mathbf{B}^t] \begin{bmatrix} \mathbf{Q}_w & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_p \end{bmatrix} [\mathbf{h}] \quad (7.2.34)$$

where \mathbf{Q}_w is the weight matrix assigned to the measurement dataset, and \mathbf{Q}_p is the weight matrix assigned to prior information. Ideally, following (7.2.6), \mathbf{Q}_p should be formulated as

$$\mathbf{Q}_p = \sigma_r^2 \mathbf{C}^{-1}(\zeta) \quad (7.2.35)$$

Note that, for consistency, the same reference variance must be used in formulation of \mathbf{Q}_p as is used in formulation of \mathbf{Q}_w . Using (7.2.6) for \mathbf{Q}_w , the latter is formulated as

$$\mathbf{Q}_w = \sigma_r^2 \mathbf{C}^{-1}(\varepsilon) \quad (7.2.36)$$

After due processing, equation (7.2.36) becomes

$$\underline{\mathbf{p}} = (\mathbf{X}^t \mathbf{Q}_w \mathbf{X} + \mathbf{B}^t \mathbf{Q}_p \mathbf{B})^{-1} (\mathbf{X}^t \mathbf{Q}_w \mathbf{h} + \mathbf{B}^t \mathbf{Q}_p \mathbf{b}) \quad (7.2.37)$$

If the prior information component of the observation dataset consists simply of a series of “observations” that each parameter is ideally equal to its prior expected value, then \mathbf{B} becomes the identity matrix \mathbf{I} . Furthermore, by cancelling out the reference variance as it is featured in the \mathbf{Q}_w and \mathbf{Q}_p matrices appearing in inverted and non-inverted matrix products on the right of (7.2.37), this equation becomes

$$\underline{\mathbf{p}} = [\mathbf{X}^t \mathbf{C}^{-1}(\varepsilon) \mathbf{X} + \mathbf{C}^{-1}(\mathbf{p})]^{-1} \mathbf{X}^t \mathbf{C}^{-1}(\varepsilon) \mathbf{h} \quad (7.2.38a)$$

On the other hand, if the reference variance is not cancelled, the equation is

$$\underline{\mathbf{p}} = [\mathbf{X}^t \mathbf{Q}_w \mathbf{X} + \mathbf{Q}_p]^{-1} \mathbf{X}^t \mathbf{Q}_w \mathbf{h} \quad (7.2.38b)$$

The absence of \mathbf{b} in equations (7.2.38a) and (7.2.38b) is an outcome of the linearization procedure discussed in section 6.5.2 whereby pre-calibration (i.e. prior) expected values of all parameters are zero.

In similar fashion to (7.2.18), the post-calibration covariance matrix of parameter error under these circumstances is easily calculated as

$$\mathbf{C}(\underline{\mathbf{p}} - \mathbf{p}) = \sigma_r^2 [\mathbf{X}^t \mathbf{Q}_w \mathbf{X} + \mathbf{Q}_b]^{-1} \quad (7.2.39a)$$

which is equivalent to

$$\mathbf{C}(\underline{\mathbf{p}} - \mathbf{p}) = [\mathbf{X}^t \mathbf{C}^{-1}(\varepsilon) \mathbf{X} + \mathbf{C}^{-1}(\mathbf{p})]^{-1} \quad (7.2.39b)$$

It can be shown that where prior parameter uncertainties and measurement noise possess multiGaussian probability distributions, equations (7.2.38a) and (7.2.39b) are derivable from Bayes equation; this is further discussed in chapter 10. Hence for this over-determined (i.e. well-posed) case, the post-calibration potential for parameter error and posterior parameter uncertainty are the same.

7.2.4 Objective Function Contours

If a model is linear, and if the inverse problem is well-posed, contours of the objective function form hyper-ellipsoids if plotted in parameter space. At their centre is the unique minimum of the objective function. This is schematized for a two-dimensional inverse problem in figure 7.1.

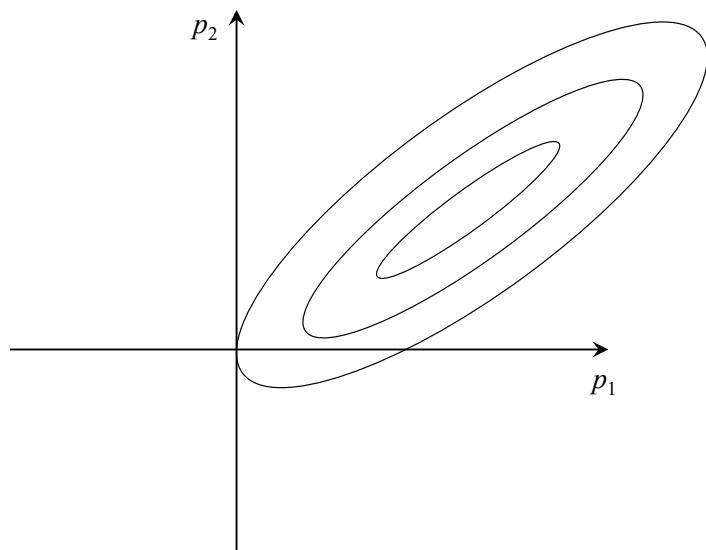


Figure 7.1. Objective function contours for a two-parameter, linear inverse problem.

If an inverse problem is nonlinear, then objective function contours do not form hyper-ellipsoids. Nevertheless, if the inverse problem is well-posed, the objective function contours close around a discrete minimum. This is illustrated for a two-parameter problem in figure 7.2.

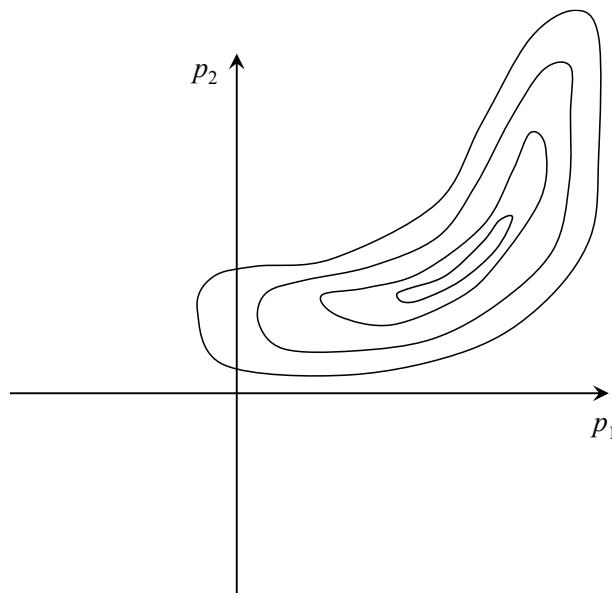


Figure 7.2. Objective function contours for a two-parameter, nonlinear, inverse problem.

Where a problem is ill-posed contours of the objective function do not close. Objective function contours for an ill-posed, nonlinear, two-parameter, inverse problem are schematized in figure 7.3.

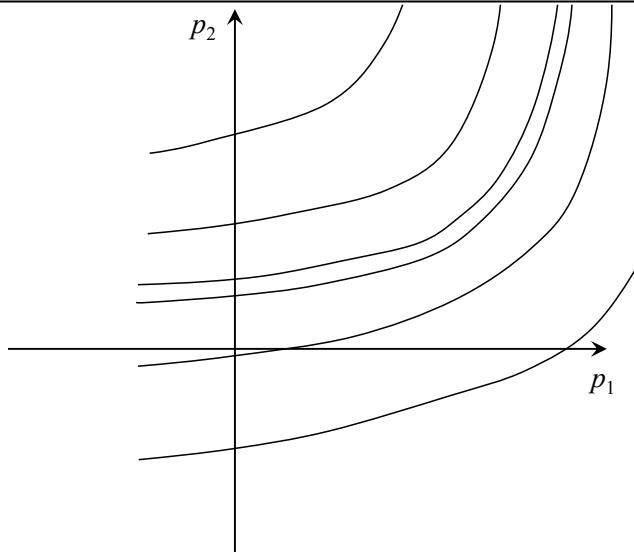


Figure 7.3. Objective function contours for a nonlinear, two-parameter, ill-posed inverse problem.

Where a problem is highly nonlinear, the objective function may contain local optima. Where it is nonlinear and ill-posed, objective function contours may not close about these local objective function minima.

7.2.5 Probability Density Contours of Posterior Parameter Error

As discussed in section 3.6, contours of the probability density function of a Gaussian distribution are elliptical in two dimensions, ellipsoidal in three dimensions and hyper-ellipsoidal in greater than three dimensions. The same applies for the Student's t distribution.

If a model is linear, and if measurement noise is normally distributed, the post-calibration probability distribution of parameter error is multinormal if the reference variance is not estimated through the calibration process, and Student's t if it is. Note that where a problem is well-posed, the prior probability distribution of model parameters is of no consequence unless preferred values for model parameters are featured as prior information. This is generally not warranted if an inverse problem is well-posed. However it is integral to Tikhonov regularization if a problem is ill-posed. For the moment, however, the discussion is focussed on well-posed inverse problems.

Following model calibration, contours of isoparameter error probability are centred on estimated parameter values. Under these circumstances, objective function contours, and contours of equal posterior parameter error probability (also known as parameter error confidence ellipsoids), coincide. The association of an objective function value with a parameter error probability value depends on the dimensionality of parameter space, and on whether a multinormal or Student's t distribution is used to characterize parameter error; see section 13.2. Where a model is mildly nonlinear, the relationship between objective function value and post-calibration parameter error probability value is still approximately maintained. However the relationship fades as a model becomes very nonlinear.

As discussed in section 3.10, eigenvectors of the covariance matrix of a set of variables whose distribution is multiGaussian point along the axes of hyper-ellipsoids of constant probability. As was also stated, the ratios of the lengths of these axes for a hyper-ellipsoid associated with any probability value are the ratios of the square roots of respective covariance matrix eigenvalues. Because of the coincidence of hyper-ellipsoids of constant probability with

objective function contours, eigenvector and eigenvalue associations of the former also characterize the geometry of the latter.

7.2.6 Residuals

The discussion so far has focused on the fate of parameters as a model is calibrated. Some features of calibrated model outputs and corresponding residuals are also of interest. These features are easier to examine if we formulate the inverse problem using equation (7.2.25) so that the covariance matrix of measurement noise is proportional to the identity matrix \mathbf{I} . Thus \mathbf{Y} replaces \mathbf{X} as the model matrix, \mathbf{g} replaces \mathbf{h} as the observation dataset, and ξ replaces ϵ as the vector of measurement noise. Transformed residuals \mathbf{f} are calculated as

$$\mathbf{f} = \mathbf{g} - \mathbf{Yp} \quad (7.2.40)$$

We denote residuals for a calibrated model as $\underline{\mathbf{f}}$ to complement the fact that they are calculated using parameters $\underline{\mathbf{p}}$. Thus

$$\underline{\mathbf{f}} = \mathbf{g} - \mathbf{Yp} \quad (7.2.41)$$

From (7.2.30) this becomes

$$\underline{\mathbf{f}} = \mathbf{g} - \mathbf{Y}(\mathbf{Y}^t \mathbf{Y})^{-1} \mathbf{Y}^t \mathbf{g} = (\mathbf{I} - \mathbf{Y}(\mathbf{Y}^t \mathbf{Y})^{-1} \mathbf{Y}^t) \mathbf{g} \quad (7.2.42)$$

Let \mathbf{e} designate outputs calculated by the model \mathbf{Y} that correspond to field measurements. Thus

$$\mathbf{e} = \mathbf{Yp} \quad (7.2.43)$$

For the calibrated model these become

$$\underline{\mathbf{e}} = \mathbf{Yp} = \mathbf{Y}(\mathbf{Y}^t \mathbf{Y})^{-1} \mathbf{Y}^t \mathbf{g} \quad (7.2.44)$$

From (7.2.44) and (7.2.42)

$$\underline{\mathbf{e}}^t \underline{\mathbf{f}} = \mathbf{g}^t \mathbf{Y}(\mathbf{Y}^t \mathbf{Y})^{-1} \mathbf{Y}^t (\mathbf{I} - \mathbf{Y}(\mathbf{Y}^t \mathbf{Y})^{-1} \mathbf{Y}^t) \mathbf{g} = 0 \quad (7.2.45)$$

Equation (7.2.45) states that, for a calibrated model, residuals are normal to model outputs. The calibration process can thus be viewed as having extracted as much information from the measurement dataset as it can by fitting this dataset as well as it can. That part of the calibration dataset that has not been fit is simply out of reach of the columns of \mathbf{Y} because there is nothing left in the measurement dataset \mathbf{g} that can project onto the columns of \mathbf{Y} . Hopefully, this is because all that remains in the dataset is measurement noise – noise that cannot be expressed as a linear combination of the columns of \mathbf{Y} because of the random nature of noise. In practice, however, it will normally be the incapacity of the model to simulate all details of an environmental system's behaviour that is responsible for the greater part of a model's incapacity to fit a calibration dataset exactly. Part of that behaviour is inexpressible as a linear combination of the columns of \mathbf{Y} because of the inadequacy of \mathbf{Y} as a descriptor of real-world behaviour.

The fact that, in most real-world modelling circumstances, model-to-measurement misfit is dominated by so-called structural (rather than measurement) noise has many repercussions. Some of these will be examined in chapter 15. In the meantime it is worth pointing out that if such noise is to be treated as if it were measurement noise so that the theory presented in this chapter can be applied with rigour, then its covariance matrix must be known. This is required so that a measurement weighting strategy, embodied in a weight matrix \mathbf{Q} , can be determined from its inverse. Regrettably the covariance matrix of structural noise can rarely, if ever, be known. Even more sadly, it is likely to be singular, and hence noninvertible, for reasons that are discussed in chapter 15.

Let us continue with our analysis nevertheless. The matrix $\mathbf{Y}(\mathbf{Y}^t\mathbf{Y})^{-1}\mathbf{Y}^t$ is idempotent, as can be readily verified by multiplying it by itself. It is also symmetric, as can be readily verified by taking its transpose and recalling the fact that the inverse of a symmetric matrix is itself symmetric. Hence it is an orthogonal projection operator. Equation (7.2.44) states that outputs calculated by a calibrated model are the orthogonal projection of measurements comprising the calibration dataset onto the range space of the model (i.e. the space spanned by the columns of the model matrix \mathbf{Y}). As can also be readily verified, $\mathbf{I} - \mathbf{Y}(\mathbf{Y}^t\mathbf{Y})^{-1}\mathbf{Y}^t$ is an idempotent symmetric matrix too. It projects onto the orthogonal complement of the range space of \mathbf{Y} . Equation (7.2.42) states that residuals pertaining to the calibrated model are the orthogonal projection of measurements onto this orthogonal subspace.

7.3 Post-Calibration Analysis

7.3.1 Problem Well-Posedness

7.3.1.1 Nonuniqueness

All of the equations presented in the previous section which involve estimation of parameters require that a matrix be inverted. Where no prior information is included in the parameter estimation process, this matrix is $\mathbf{X}^t\mathbf{Q}\mathbf{X}$.

If \mathbf{X} has a null space then $\mathbf{X}^t\mathbf{Q}\mathbf{X}$ is not invertible. If PEST is asked to solve an ill-posed inverse problem using no mathematical regularization device (which occurs if it runs in “estimation” mode), it may not work. The objective function may not fall. Problems may also be encountered if the $\mathbf{X}^t\mathbf{Q}\mathbf{X}$ matrix is near-singular. It can be shown that attempted inversion of a near-singular matrix (or attempted solution of a set of linear equations in which near-singular matrix inversion is implied) leads to gross amplification of tiny numerical errors that are incurred when matrix and vector elements are calculated and then stored in a computer. This amplification of numerical noise can easily dominate the “solution” to the inverse problem; the solution may therefore be meaningless. (In the next chapter we show why manual or numerical noise is amplified when attempting to solve an inverse problem that is ill-posed or nearly ill-posed.)

In spite of noninvertibility $\mathbf{X}^t\mathbf{Q}\mathbf{X}$, PEST may nevertheless derive a parameter set that fits a calibration dataset well. This occurs because addition of the Marquardt lambda (see later) to the diagonal of the $\mathbf{X}^t\mathbf{Q}\mathbf{X}$ matrix may facilitate a “false inversion” of this matrix, and the achievement of a set of parameters which minimize the objective function. The set of parameters that is obtained through this means is not unique however. Under a slightly different weighting scheme, or through use of a different Marquardt lambda selection strategy, an entirely different set of parameters may be obtained which also minimizes the objective function. It is unlikely that either of these parameter sets will approach that of minimized post-calibration error variance; recall that this was the goal that we set ourselves in section 6.4.

Use of equation (7.2.14) is not the only way to seek a parameter set \mathbf{p} that yields a good fit between field measurements and corresponding model outputs. (However it is generally the most model-run-efficient way.) Where inversion is implemented using so-called “global” methods, the attainment of different parameter sets that fit a calibration dataset equally well is sometimes offered as proof of the existence of local minima in the objective function – a phenomenon that can indeed accompany attempts to calibrate some types of highly nonlinear models. However, the discovery of multiple sets of parameters that fit a calibration dataset equally well may illustrate formulation of an ill-posed inverse problem rather than the existence

of local optima. Chance plays a large role in “solution” of an ill-posed inverse problem where regularization is eschewed. Of the many solutions that can be obtained to that problem, only one can be of minimum error variance.

7.3.1.2 Parameter Variance-Covariance Analysis

Invertibility of the $\mathbf{X}^t \mathbf{Q} \mathbf{X}$ matrix is a necessary, but not a sufficient, condition for success of the manually-regularized parameter estimation process. Another requirement is that post-calibration error variances that are computed for estimated parameters be reduced below their pre-calibration levels. The latter are their prior uncertainties. (“Pre-calibration error variance” characterizes the extent to which parameters may depart from their prior expected values. This, of course, is nothing other than a characterization of parameter uncertainty. Recall that variances constitute the diagonal elements of a covariance matrix, and that the variance of a random variable is the square of its standard deviation.)

If an inversion dataset includes prior information on all parameters which specifies that each parameter is “observed” to have a value equal to its prior mean value then, provided that this prior information is weighted appropriately, the error variance of no parameter should be assessed to rise through the inversion process. In this case parameter estimation is based on equation (7.2.38). In practice, as has already been mentioned, prior information is often omitted from formulation of a manually-regularized inverse problem. There are a number of reasons for this. One reason is that modellers often do not want to commit themselves to prior expected parameter values, especially if an inverse problem is “well-posed enough” for the information content of a calibration dataset to dominate that of prior expert knowledge. Another reason is that it is often difficult to derive a weighting scheme which adequately balances the uncertainties of prior information against those of measurement error. This is especially the case where “measurement error” is actually dominated by model structural error (as it often is); this is a topic to which we will return later in this book when we consider Tikhonov regularization.

Manually-regularized parameter estimation is thus often accomplished using equation (7.2.14) without prior information. Post-calibration parameter error variance is then assessed using equation (7.2.18). If the post-calibration error variances of some parameters are unacceptably high, it is then incumbent on a modeller to repeat the parameter estimation process with manual regularization strengthened. This may require that some parameters be fixed at their prior expected values and/or that other parameters be combined into composite parameters. The number of estimated parameters is thereby reduced. Presumably, parameters that are eliminated from the revised inversion process are those whose error variances rose, rather than fell, during the previous inversion process. These will have been revealed by studying the post-calibration parameter error covariance matrix that emerged from that process.

While the covariance matrix of post-calibration parameter error may indeed reveal that post-calibration error variances for some parameters are unacceptably high, the reasons for inflated post-calibration error variance may not be readily apparent from an inspection of this matrix. There are, in fact, two reasons why the post-calibration error variance of a parameter may be high. One is that the calibration dataset contains little or no information about this parameter. The second is that information in the calibration dataset that pertains to this parameter is shared with one or a number of other parameters; this information may be insufficient to resolve each parameter individually. These two situations are illustrated for a simple two-parameter inverse problem in figure 7.4.

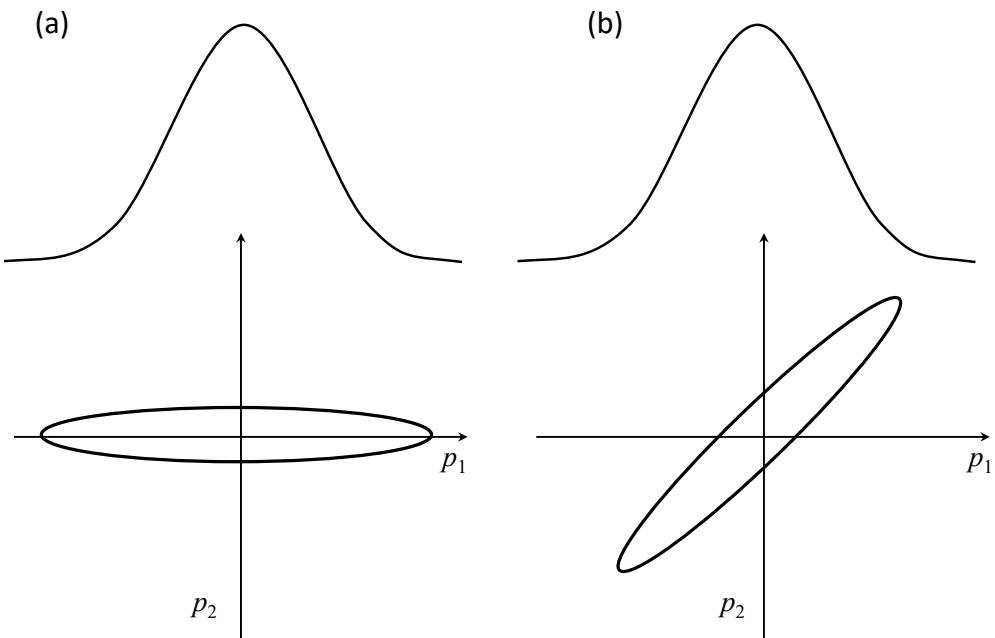


Figure 7.4. A contour of post-calibration parameter error variance, and the post-calibration marginal probability distribution of error for parameter p_1 . In part (a) of this figure, parameter p_1 has a high post-calibration error variance because the calibration dataset has little information that pertains to this parameter. In part (b) of this figure, the calibration dataset does contain information pertaining to parameter p_1 . However this information is shared with parameter p_2 .

A modeller can readily distinguish between the two cases that are depicted in figure 7.4 by inspecting the post-calibration parameter error correlation coefficient matrix. Correlation coefficients are calculated from covariances using equation (3.5.3). Recall that the diagonal elements of a correlation coefficient matrix are all 1. The matrix is symmetrical. Off-diagonal elements of this matrix lie between 1 and -1. The closer that an off-diagonal element approaches 1 or -1, the greater is the degree to which pertinent parameters are correlated; hence the greater the degree to which a contour of the post-calibration parameter error probability density function looks like part (b) of figure 7.4 rather than part (a) of this figure. As correlation coefficients approach 1 or -1, error variance contours become narrower than illustrated in this figure. In the limit that they are actually 1 or -1, the contours are so narrow that they become straight lines. In the limit they do not close. At this stage the matrix \mathbf{X} has a null space. Its null space lies along the direction of elongation of the ellipse. (This is the direction of infinitely many solutions to the inverse problem.)

The contour that is depicted in part (b) of figure 7.4 illustrates positive parameter correlation, and hence a correlation coefficient that is greater than zero. An ellipse that slopes upwards to the left indicates negative correlation between parameters. A circle, or an ellipse which is oriented along a parameter axis, indicates zero correlation between respective parameters.

The correlation coefficient matrix thus informs a modeller whether excessive correlation with other parameters may be the reason for high post-calibration error variance of a particular parameter. Unfortunately, however, there is a practical problem associated with use of this matrix. The closer a covariance matrix approaches singularity (a matrix is singular when it cannot be inverted), the greater is the amount of numerical error that is incurred when calculating correlation coefficients. This applies especially to correlation coefficients that approach 1 and -1. High levels of post-calibration parameter correlation may therefore go

unnoticed.

Fortunately, the same does not apply to calculation of eigenvectors and eigenvalues of a post-calibration parameter error covariance matrix. Hence these are probably the best source of information on well- or ill-posedness of a parameter estimation problem and, if ill-posed, the reason for its ill-posedness. If the ratio of highest to lowest eigenvalue of a post-calibration parameter error covariance matrix is large, this indicates that post-calibration parameter error probability ellipses are unduly long in at least one direction of parameter space; this direction is that of the eigenvector corresponding to the highest eigenvalue. If a number of eigenvalues are unduly large in comparison to others, then there are a number of directions of probability ellipse elongation; once again, elongation is in the direction of the corresponding eigenvectors. (Recall that eigenvalues and eigenvectors are matched to each other.) If the eigenvector corresponding to an unduly large eigenvalue is dominated by a single parameter, then the situation is as depicted in figure 7.4a. However if a number of parameter components are equally large, then the situation is as depicted in figure 7.4b. Eigenvectors listed on PEST printouts are normalized so that their magnitudes are 1.0. Hence a modeller can see at a glance whether an eigenvector is dominated by a single parameter, or whether a number of parameters feature strongly in this vector.

It is apparent that the information content of the eigenvectors and eigenvalues of the post-calibration parameter error covariance matrix is large. Furthermore, this information is formulated in a way that is readily digestible by a modeller. If there is a single number that can characterize the health (or otherwise) of a manually-regularized inversion process, it is the ratio of highest to lowest eigenvalue of the post-calibration parameter error covariance matrix. This is equal to the so-called “condition number” of this covariance matrix. It is apparent from (7.2.14) that it is also the condition number of the matrix that must be inverted to estimate parameter values (it is easy to show that the condition number of the inverse of a matrix is the same as that of the original matrix). Matrices that have a high condition number are difficult to invert because of their near-singularity; this is because numerical errors are amplified in their inversion. In PEST usage, the predominant source of numerical error in the $\mathbf{X}^t \mathbf{Q} \mathbf{X}$ matrix is the calculation of finite-difference derivatives required to fill the \mathbf{X} matrix.

Experience in using PEST suggests that if the ratio of highest to lowest eigenvalue of the post-calibration parameter error covariance matrix is greater than about 5×10^7 , the numerical integrity of the matrix inversion process is highly questionable. (Parameter outcomes of the inversion process generally lose their integrity at much lower eigenvalue ratios than this.) Refinement of the manual regularization strategy then becomes essential. If the eigenvector associated with the highest eigenvalue is dominated by a single parameter, then that parameter should not be estimated in the reformulated inverse problem. Instead it should be held at its prior expected value. If, on the other hand, a number of parameters feature strongly in this vector, then one of them should probably be fixed at its prior expected value in the reformulated inverse problem. The parameter, or parameters, with which it is correlated will then have lower post-calibration error variances. This is illustrated in figure 3.4. However this same figure illustrates that the values estimated for parameters that are correlated with the newly fixed parameter may be highly dependent on the value at which this parameter is fixed. Hence their apparent post-calibration certainty may be illusory. Obviously, when choosing the parameter to fix from a group of correlated parameters, the parameter with least prior uncertainty should be chosen, as this is the one that a modeller is least likely to fix at the wrong value.

Unfortunately, a fundamental problem with analysis of the post-calibration parameter error covariance matrix is that this matrix cannot be computed if the $\mathbf{X}^t \mathbf{Q} \mathbf{X}$ matrix of equation

(7.2.14) cannot be inverted. This is most unfortunate. Just when the information that this matrix can provide is needed most in order to shed light on the origins of inverse problem ill-posedness, it cannot be calculated. Later in this book, singular value decomposition is used to solve ill-posed inverse problems. It provides the same information that the post-calibration parameter error covariance matrix would have provided if it could be calculated. However this information is available prior to solving the inverse problem instead of afterwards. Furthermore, at the same time as it identifies the causes of inverse problem ill-posedness, singular value decomposition eliminates these causes prior to estimating parameters.

7.3.1.3 Composite Parameter Sensitivities

One means through which parameters which give rise to the situation depicted in figure 7.4a can be identified is through calculation of so-called “composite” sensitivities of all parameters that feature in the manually-regularized inverse problem, and then comparing these sensitivities with each other. Those parameters whose composite sensitivities are low relative to those of other parameters are candidates for removal from the inversion process if that process seems to be suffering from the deleterious effects of incipient inverse problem ill-posedness. A disadvantage of this statistic however is that it does not identify the problem that is depicted in figure 7.4b, for a parameter may have a high composite sensitivity but still be highly correlated with other parameters. However an advantage of the composite sensitivity statistic over eigenanalysis of the post-calibration parameter error covariance matrix is that it does not require calculation of that covariance matrix. Hence it can be used even if the $\mathbf{X}^t \mathbf{Q} \mathbf{X}$ matrix cannot be inverted.

The composite sensitivity of a parameter is the magnitude of the column of the \mathbf{X} matrix pertaining to that parameter, with weighting taken into account. This is generally then divided by the number of observations featured in the parameter estimation process; that is, it is divided by the number of rows of \mathbf{X} . If the weighting matrix (i.e. the \mathbf{Q} matrix) is diagonal, the composite sensitivity of parameter i (i.e. csp_i) can therefore be calculated as

$$csp_i = \frac{\left[\sum_j^n q_{jj} x_{ji}^2 \right]^{\frac{1}{2}}}{n} \quad (7.3.1a)$$

Where \mathbf{Q} is non-diagonal the composite sensitivity of parameter i can be calculated as

$$csp_i = \frac{[\mathbf{X}^t \mathbf{Q} \mathbf{X}]_{ii}^{\frac{1}{2}}}{n} \quad (7.3.1b)$$

In both of these equations n is, as usual, the number of observations comprising the calibration dataset.

Hill and Tiedeman (2007) suggest normalization of each composite parameter sensitivity by dividing it by the value of the pertinent parameter. If a parameter is log transformed for the purpose of parameter estimation (see below), this happens automatically. If parameters are subjected to Kahunen-Loëve transformation prior to estimation then, once again, scaling occurs automatically.

Experience in using PEST suggests that if the ratio of highest to lowest composite parameter sensitivity is of the order of 200 or greater, then those parameters with lowest composite sensitivities should be removed from the manually-regularized inversion process because the $\mathbf{X}^t \mathbf{Q} \mathbf{X}$ matrix is probably approaching singularity and is therefore invertible only with difficulty.

7.3.2 Analysis of Residuals

7.3.2.1 Goodness of Fit

Ideally, the goodness of fit attained through solution of an inverse problem should be commensurate with the level of measurement noise associated with the calibration dataset. If this level of fit cannot be attained, this suggests that the model may be somewhat defective in simulating environmental processes. To most environmental modellers this should come as no surprise. On the other hand, attainment of a higher level of fit than that expected on the basis of measurement noise is to be treated with extreme caution. Such “over-fitting” is often an indication of under-regularization and consequential inflation of post-calibration parameter error variance. The reason for this will become clear when regularization implemented through singular value decomposition is discussed.

Where regularization is manual, it may not be model defects that are responsible for a less-than-pleasing fit between model outcomes and field measurements. Unsatisfactory model-to-measurement fit may result from excessively heavy-handed manual regularization whereby too few parameters are estimated in order to avoid problem ill-posedness. Spatial and temporal details of consequential parameter-simplification-induced misfit may suggest to a modeller how best to revise parameter definition in order to reduce this misfit on the next inversion run, without at the same time inducing inverse problem ill-posedness.

An often-used measure of fit between model outcomes and field measurements is the correlation coefficient (which should not be confused with the correlation coefficient defined in equation 3.5.3). It is calculated as

$$R = \frac{(\mathbf{Q}^{1/2}\mathbf{h} - \boldsymbol{\mu}_h)^t (\mathbf{Q}^{1/2}\mathbf{o} - \boldsymbol{\mu}_o)}{\left[(\mathbf{Q}^{1/2}\mathbf{h} - \boldsymbol{\mu}_h)^t (\mathbf{Q}^{1/2}\mathbf{h} - \boldsymbol{\mu}_h) (\mathbf{Q}^{1/2}\mathbf{o} - \boldsymbol{\mu}_o)^t (\mathbf{Q}^{1/2}\mathbf{o} - \boldsymbol{\mu}_o) \right]^{1/2}} \quad (7.3.2)$$

where

$$\boldsymbol{\mu}_{hi} = \frac{\sum_{j=1}^n (\mathbf{Q}^{1/2}\mathbf{h})_j}{n} \quad \text{for all } i \quad (7.3.3a)$$

and

$$\boldsymbol{\mu}_{oi} = \frac{\sum_{j=1}^n (\mathbf{Q}^{1/2}\mathbf{o})_j}{n} \quad \text{for all } i \quad (7.3.3b)$$

R defined through (7.3.2) provides a measure of the extent to which variability of field measurements is explained by the calibrated model compared to that which can be construed as purely random. Some modellers recommend that success of a calibration exercise should be judged by whether or not a certain correlation coefficient is attained (for example 0.9). It is the author’s opinion however, that this statistic (and similar fit/misfit statistics) provides little assistance to the necessarily subjective process of judging the success or otherwise of a parameter estimation process. Its usefulness is further degraded where observations are of different types, and where observation weighting includes strategies that are discussed in chapter 15 of this book that attempt to mitigate the potential for parameter and predictive bias induced by structural noise.

7.3.2.2 Statistical Analysis of Residuals

Many texts on parameter estimation discuss methods through which residuals can be analysed. These analyses are often undertaken in order to verify that assumptions regarding the statistical characterization of measurement noise are correct. When working with well-posed inverse

problems (which, in the environmental context, are often a consequence of heavy-handed manual regularization) the outcomes of the inversion process may be highly dependent on this characterization. Ideally a weight matrix \mathbf{Q} is formulated in a way that complements statistical characterization of measurement noise; see equation (7.2.6). Parameters are then estimated using formulas that include this weight matrix. The uncertainties associated with these estimates, as well as the uncertainties of model predictions, are also evaluated using formulas that include this weight matrix. Ultimately, parameter and predictive uncertainty are attributed in large part, if not entirely, to measurement noise. Proper statistical characterization of measurement noise is therefore essential.

As has been alluded to already, and as will be discussed later in this book, in most environmental modelling contexts the bulk of model-to-measurement misfit cannot be attributed to noise associated with the making of measurements. In fact, it is attributable to shortcomings of a numerical model as a simulator of complex environmental behaviour. If calibration is based on manual regularization, parameter simplification can also introduce substantial amounts of structural noise. The covariance matrix of structural noise is unknown, and is probably singular. Statistical characterization of what is mostly a deterministic quantity is generally impossible. Much of the theory on which manually-regularized parameter estimation and uncertainty quantification is based is therefore invalidated.

For complex models that simulate system state, and changes to system state, at many different locations in a three-dimensional model domain, it is unlikely that quantitative post-calibration analysis of residuals will do anything other than invalidate statistical characterizations of measurement noise upon which construction of a weight matrix is based. Hence methodologies for residuals analysis are not discussed herein. This is not meant to suggest that a modeller should not inspect residuals carefully in space and in time, and draw whatever conclusions that he/she can about model or parameterization inadequacy from differences between model outputs and true system behaviour. This, in fact, is an essential component of good calibration practice. It is suggested however that a serious statistical analysis of residuals is unlikely to yield results that are useful to the history-matching process.

While this may appear at first to be a rather depressing conclusion, the situation should be seen in context. An important feature of this context is that for many predictions of management interest, predictive uncertainty is not so much a function of measurement noise as it is of information insufficiency in the calibration dataset. Integrity of characterization of predictive uncertainty therefore depends more on strategic formulation of an ill-posed inverse problem and on proper characterization of prior parameter uncertainty than it does on stochastic characterization of measurement noise.

A second redeeming feature is that formulation of an objective function in most real-world modelling contexts will often be a matter that requires considerable subjectivity and more than a little creativity. The objective function will often be at least partly formulated in such a way as to compare observations which are processed in some way with respective model outputs which are processed in the same manner. As will be discussed in chapter 15, such a strategy can mitigate some of the more deleterious impacts of structural noise by partially filtering out those aspects of it which have too much structure and too little stochasticity. Weighting of different components of an objective function formulated in this way is a necessarily subjective matter. Some of this subjectivity is unavoidably inherited by values ascribed to estimated parameters through the inversion process, and by assessments of the uncertainties associated with these estimates. With acknowledgement of the necessarily high subjective input into history-matching in particular, and environmental modelling in general, comes the freedom to

dispense with certain aspects of quantitative analysis that are unlikely to add much value to the overall process.

7.3.3 Influence Statistics

7.3.3.1 Background

A number of statistics have been developed for comparing the relative information content of different observations comprising a calibration dataset. These have been designed for use in the over-determined parameter estimation context that is the subject of the present chapter. As such, they inherit limitations which pertain to manual regularization as a device for solution of inverse problems that arise in decision-support environmental modelling. These limitations will be discussed later in this chapter.

It is of interest to note that some of the statistics that are discussed in the present and previous sections of this chapter (and many statistics that will be discussed in future chapters of this book) require only sensitivities of model outputs with respect to adjustable model parameters as encapsulated in the \mathbf{X} matrix; they require neither the values of parameters nor of measurements. These statistics can therefore be employed even if a model has not yet been calibrated, and even if data that will comprise an hypothesized calibration dataset has not yet been gathered. They can therefore be used to examine the benefits of a future calibration process, and inquire into the sources of these benefits. Where required, estimates of the objective function that will be achieved in a future calibration process can be guessed; this guess can be supported by an assessment of the amount of measurement noise that the calibration dataset is likely to host.

In real-world modelling practice, where earth system behaviour and the models which simulate this behaviour are nonlinear, the \mathbf{X} matrix on which calibration-pertinent statistics are based must be replaced by a Jacobian matrix. Ideally this matrix should be calculated using a parameter set that is the outcome of model calibration. If a model is truly linear the Jacobian matrix does not, of course, vary with parameter values. If a model is nonlinear this is not the case. Nevertheless, if a model has not yet been calibrated, an \mathbf{X} matrix that is calculated using non-calibrated parameters can suffice for formulation of some of the statistics that are discussed in the present section, and many of those that are discussed in later chapters of this book. Prior parameter expected values constitute a sound basis for calculation of the \mathbf{X} matrix under these circumstances.

The following discussion is brief. Those interested in this topic should refer to Hadi (1982), Yager (1998), Cook and Weisberg (1982), and Belsley et al (1980) for further details.

7.3.3.2 The Statistics

Composite observation sensitivity constitutes a crude influence statistic. This complements the composite parameter sensitivity statistic discussed above. It is the magnitude of the row of the \mathbf{X} matrix that corresponds to the observation of interest, divided by the number of parameters. Where the weight matrix is diagonal, the composite sensitivity of observation i can be written as follows.

$$CSO_i = \frac{\left[q_i \sum_j^m x_{ij}^2 \right]^{\frac{1}{2}}}{m} \quad (7.3.4a)$$

Here m is the number of parameters featured in the manually-regularized calibration process. More generally

$$cso_i = \frac{[\mathbf{Q}^{\frac{1}{2}} \mathbf{X} \mathbf{X}^T \mathbf{Q}^{\frac{1}{2}}]_{ii}^{\frac{1}{2}}}{m} \quad (7.3.4b)$$

If the composite sensitivity of an observation is low, then the model output which corresponds to that observation is sensitive to no parameter. As a consequence the observation carries little information with respect to any parameter.

For the remainder of this section we will employ the alternative mathematical characterization of the inversion process that was introduced above. This endows (transformed) measurement noise with a covariance matrix that is proportional to \mathbf{I} ; see equation (7.2.25) and subsequent equations. This simplifies the equations to follow.

Let the matrix \mathbf{H} be defined as follows:

$$\mathbf{H} = \mathbf{Y}(\mathbf{Y}^T \mathbf{Y})^{-1} \mathbf{Y} \quad (7.3.5)$$

Recall from section 7.2.6 that \mathbf{H} is an orthogonal projection operator onto model output space. Equation (7.2.44) can now be re-written as

$$\underline{\mathbf{e}} = \mathbf{H}\mathbf{g} \quad (7.3.6)$$

Recall also that $\underline{\mathbf{e}}$ is the vector of calibration-pertinent outputs calculated by the calibrated model, while \mathbf{g} comprises the calibration dataset itself. The covariance matrix of $\underline{\mathbf{e}}$ can be calculated as

$$\mathbf{C}(\underline{\mathbf{e}}) = \mathbf{H}\mathbf{C}(\mathbf{g})\mathbf{H}^T \quad (7.3.7)$$

The covariance matrix of \mathbf{g} is simply the covariance matrix of ξ , that is the covariance matrix of weight-transformed measurement noise. This is $\sigma_r^2 \mathbf{I}$. Hence

$$\mathbf{C}(\underline{\mathbf{e}}) = \sigma_r^2 \mathbf{H}\mathbf{H}^T = \sigma_r^2 \mathbf{H} \quad (7.3.8)$$

Similarly, for residuals $\underline{\mathbf{f}}$ pertaining to the calibrated model (see equation 7.2.42)

$$\mathbf{C}(\underline{\mathbf{f}}) = \sigma_r^2 (\mathbf{I} - \mathbf{H}) \quad (7.3.9)$$

The variance of the i 'th element of $\underline{\mathbf{e}}$ divided by the variance of the i 'th element of $\underline{\mathbf{f}}$ is denoted as the “leverage” of observation i . From (7.3.8) and (7.3.9) it is readily calculated as

$$L_i = \frac{h_{ii}}{1-h_{ii}} \quad (7.3.10)$$

where, of course, h_{ii} is the i 'th diagonal element of \mathbf{H} .

The Cook's D statistic takes the value of a weighted residual specifically into account when computing the influence of the corresponding observation. The Cook's D statistic for observation i can be thought of as the relative distance between the centre of confidence ellipsoids for $\underline{\mathbf{p}}$ based on the full calibration dataset on the one hand, and the centre of the confidence ellipsoids for $\underline{\mathbf{p}}$ based on a calibration dataset with observation i omitted from it on the other hand. (This concept is easily extended to the omission of more than one observation from a calibration dataset.) It is calculated as

$$D_i = \frac{1}{m} v_i^2 \frac{h_{ii}}{1-h_{ii}} \quad (7.3.11)$$

Where m is the number of parameters, and v_i is the weighted “Studentized” residual corresponding to observation i . The latter is computed as

$$v_i = \frac{f_i}{\sigma_r \sqrt{1-h_{ii}}} \quad (7.3.12)$$

$\underline{\sigma}_r$ is the square root of the reference variance $\underline{\sigma}_r^2$ estimated from the minimized objective function using equation (7.2.23).

Focusing on outliers, Hadi (1982) suggests the following statistic for characterizing the influence of the i 'th observation on estimated parameters as a whole

$$H_i^2 = \frac{m}{1-h_{ii}} \frac{f_i^2}{(f^t f - f_i^2)} + \frac{h_{ii}}{1-h_{ii}} \quad (7.3.13)$$

DFBETAS $_{ij}$ is a measure of the influence of observation i on estimated parameter j . It is calculated as

$$\beta_{ij} = \frac{c_{ji}}{\sqrt{\sum_{k=1}^n c_{jk}^2}} \frac{f_i}{s(i)(1-h_{ii})} \quad (7.3.14)$$

where c_{jk} is an element of the matrix product $\mathbf{C} = (\mathbf{Y}^t \mathbf{Y})^{-1} \mathbf{Y}^t \cdot s(i)$ is calculated as follows (Belsley et al, 1980)

$$s^2(i) = \frac{1}{(n-m-1)} \left[(n-m) \underline{\sigma}_r^2 - \frac{f_i^2}{1-h_{ii}} \right] \quad (7.3.15)$$

Much more will be said on the subject of data worth later in this book. There the discussion will focus on assessment of data worth in highly parameterized contexts wherein the null space may contribute substantially more to parameter and predictive uncertainty than measurement noise.

7.4 Nonlinear Models

7.4.1 The Jacobian Matrix

Theory presented so far in this chapter assumes that relationships between parameters \mathbf{p} of a model and model outputs \mathbf{o} to which there are corresponding field measurements \mathbf{h} are linear, and hence can be represented by a matrix operation. In the over-determined context made possible by manual regularization, this matrix has been designated as \mathbf{X} . Theory presented above has demonstrated how a parameter set \mathbf{p} that supposedly satisfies the calibration metric stated in section 6.4 can be estimated from \mathbf{h} . (We will review whether this metric can really be met in the manually-regularized context shortly.) Formulas were derived for calculation of the post-calibration parameter error covariance matrix $\mathbf{C}(\mathbf{p} - \mathbf{p})$. Strategies were presented for discerning the well-posedness or otherwise of the manually-regularized inverse problem. Means through which the influence of different observations on the parameter estimation process could be assessed were provided. Equations which underpin all of these calculations use the matrix \mathbf{X} .

A nonlinear model must be linearized before any of the theory presented above can be implemented. This is achieved through replacing the \mathbf{X} matrix featured in the above equations with the Jacobian matrix. See section 6.5.1 for specifications of a Jacobian matrix. For a nonlinear model, this matrix is a function of parameter values. The parameter values on which its computation should be based depend on what it is being computed for. If it is being computed during one iteration of an objective function minimization procedure, then parameter values pertaining to the current iteration of that procedure should be used in its calculation. If it is being computed in order to calculate the post-calibration parameter error covariance matrix, then parameter values estimated by the calibration process should be used in its

calculation. The same applies to calculation of observation influence statistics. In some cases, formulas for the latter require values of minimized residuals as well; ideally, their calculation should be preceded by a calibration exercise in which both optimized parameters and optimized residuals are calculated.

There may be occasions, however, where some of the equations presented above are deployed in contexts that are not preceded by estimation of parameters. This may occur, for example, where a modeller would like to obtain estimates of calibration-constrained parameter error variance, and/or establish the well-posedness or otherwise of a manually-regularized inverse problem, before actually attempting the inversion. In this case Jacobian matrix calculation may be based on prior expected parameter values.

Statistics such as post-calibration parameter error variance calculated for a nonlinear model using a linearized substitute for that model are, of course, likely to be approximate. Nevertheless they may be useful. If, for example, the ratio of highest to lowest eigenvalue of a parameter error covariance matrix indicates ill-posedness of an inverse problem, then the problem is indeed ill-posed as charged. Parameters identified as being responsible for its ill-posedness through being either individually insensitive, or through being part of a group of parameters identified as making major contributions to an eigenvector whose corresponding eigenvalue is very large, are indeed problematic. Attention must therefore be given to these parameters when formulating a revised manual regularization strategy. Linear estimates of post-calibration parameter error variance may indeed be approximately correct; as will be discussed in later chapters however, the linearized post-calibration parameter error covariance matrix can provide a useful starting point for more robust nonlinear analysis. Meanwhile comparative parameter error variances provided by this matrix are likely to be valid.

It is in the inversion process itself where the need for modification of linear theory is most urgent. It is to this topic that we now turn.

7.4.2 Nonlinear Parameter Estimation

The means by which the linear theory presented so far can be adapted to estimation of parameters of a nonlinear model is now addressed.

7.4.2.1 Formulation of Equations

As has already been discussed, where a model is nonlinear the parameter estimation process becomes an iterative procedure. This is schematized for a two-parameter model in figure 7.5.

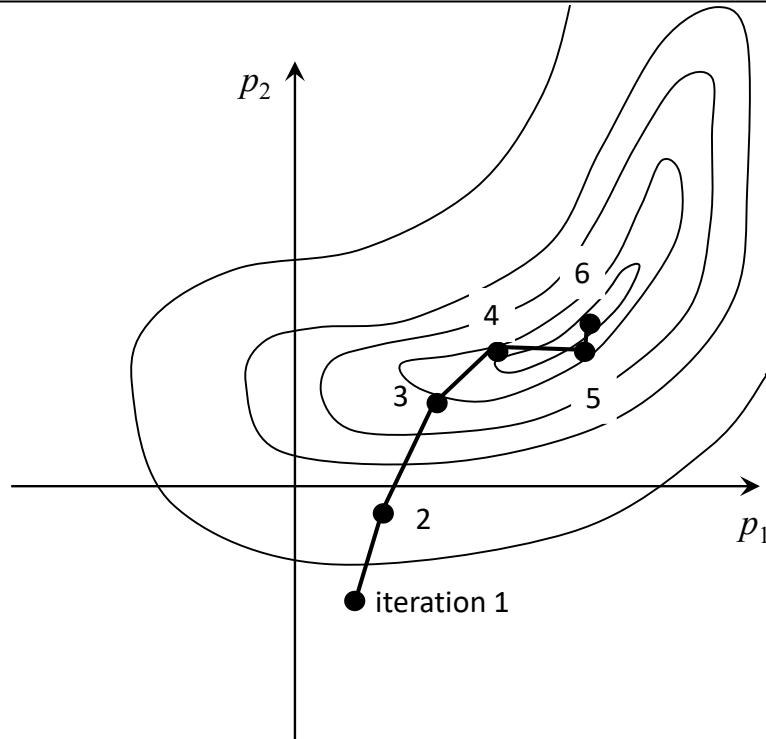


Figure 7.5. Iterative adjustment of parameters toward the objective function minimum during a nonlinear parameter estimation process.

During each iteration of the objective function minimization process, equation (7.2.14) is used to improve values assigned to parameters from those which they inherited from the previous iteration. Parameter improvements are calculated using the following formula

$$\mathbf{p} - \mathbf{p}_0 = (\mathbf{J}^t \mathbf{Q} \mathbf{J})^{-1} \mathbf{J}^t \mathbf{Q} \mathbf{r} \quad (7.4.1)$$

In equation (7.4.1) \mathbf{p}_0 designates parameter values at the start of the iteration while \mathbf{r} designates model-to-measurement residuals calculated using the \mathbf{p}_0 parameter set. \mathbf{J} is the Jacobian matrix calculated using the \mathbf{p}_0 parameter set.

7.4.2.2 The Marquardt Lambda

Use of equation (7.4.1) directly to estimate parameters for a nonlinear model can be numerically inefficient. In practice, the diagonal terms of the $\mathbf{J}^t \mathbf{Q} \mathbf{J}$ matrix should be enhanced through use of the following equation to calculate parameter improvements

$$\mathbf{p} - \mathbf{p}_0 = (\mathbf{J}^t \mathbf{Q} \mathbf{J} + \lambda \mathbf{I})^{-1} \mathbf{J}^t \mathbf{Q} \mathbf{r} \quad (7.4.2)$$

λ is the so-called “Marquardt parameter”, or simply “Marquardt lambda”, named after Marquardt (1963) who employed this strategy. However use of this parameter was, in fact, pioneered by Levenberg (1944).

Suppose that λ is very high. Then (7.4.2) becomes

$$\mathbf{p} - \mathbf{p}_0 = \lambda^{-1} \mathbf{J}^t \mathbf{Q} \mathbf{r} \quad (7.4.3)$$

It is not difficult to show that the gradient of the objective function in parameter space is given by

$$\nabla \Phi = \begin{bmatrix} \frac{\partial \Phi}{\partial p_1} \\ \frac{\partial \Phi}{\partial p_2} \\ \vdots \\ \frac{\partial \Phi}{\partial p_m} \end{bmatrix} = -2\mathbf{J}^T \mathbf{Q} \mathbf{r} \quad (7.4.4)$$

Equation (7.4.3) states that if the Marquardt lambda is very high, then the parameter estimation process simply moves parameters down the objective function gradient. Also, the higher is the Marquardt lambda, the less movement takes place.

Moving down-gradient makes sense early in a parameter estimation process. This accommodates non-ellipticity of objective function contours arising from model nonlinearity. However as the objective function minimum is approached, efficiency of the objective function minimization process requires that the Marquardt lambda be decreased. This avoids the phenomenon of “hemstitching”, which is shown schematically in figure 7.6. The threat posed by hemstitching is especially high where objective function contours are long and narrow. A strength of equation (7.4.1) is that parameter estimation is able to proceed with great efficiency even where the objective function surface is characterized by this type of topography.

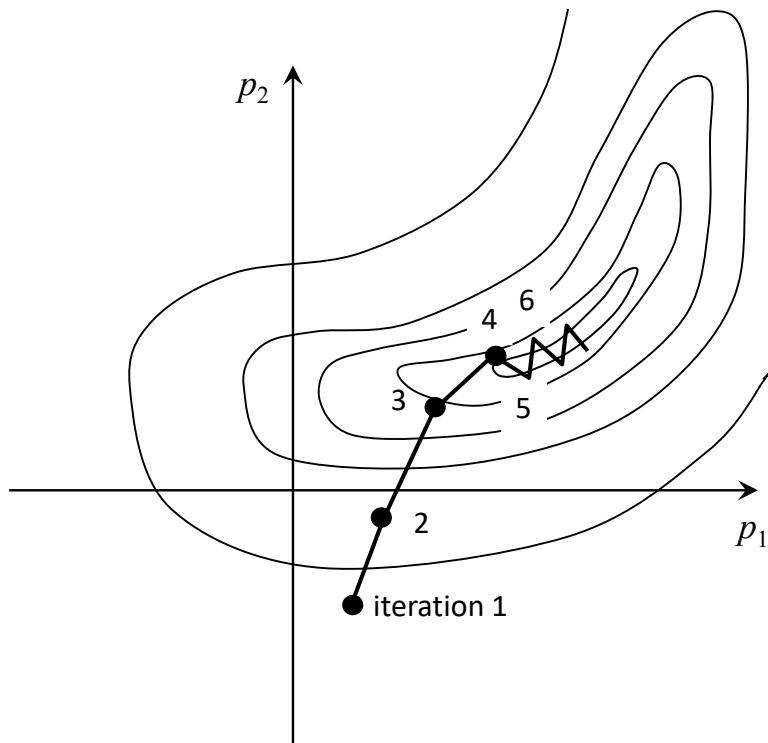


Figure 7.6. “Hemstitching” in search of the objective function minimum.

It follows from the above discussion that a nonlinear parameter estimation process should commence with the Marquardt lambda set to a value which is large enough to promote a reasonable component of direct down-gradient parameter movement. However as the parameter estimation process progresses, the Marquardt lambda should ideally decrease in value; this allows the parameter estimation process to squeeze into narrow objective function valleys if this is required.

Parameter estimation based on equation (7.4.2) is variously referred to as the “Gauss-Marquardt-Levenberg”, “Levenberg-Marquardt”, or “damped least squares” method. Alternative implementations of this method differ in the way that the Marquardt lambda is

calculated at different stages of the parameter estimation process. PEST implements a trial-and-error procedure that makes maximum use of spare capacity in a parallel computing environment.

Some numerical advantages can be gained by modifying equation (7.4.2) to prevent corruption of the parameter estimation process by numerical round-off errors. Let \mathbf{S} be a square $m \times m$ matrix (where m is the number of parameters featured in the parameter estimation process) with diagonal elements only. Let the i 'th diagonal element of \mathbf{S} be given by

$$s_{ii} = (\mathbf{J}^t \mathbf{Q} \mathbf{J})_{ii}^{-1/2} \quad (7.4.5)$$

If \mathbf{S} is introduced into equation (7.4.2) the following equation is obtained for $\mathbf{S}^{-1}(\mathbf{p} - \mathbf{p}_0)$

$$\mathbf{S}^{-1}(\mathbf{p} - \mathbf{p}_0) = ((\mathbf{J} \mathbf{S})^t \mathbf{Q} \mathbf{J} \mathbf{S} + \lambda \mathbf{S}^t \mathbf{S})^{-1} (\mathbf{J} \mathbf{S})^t \mathbf{Q} \mathbf{r} \quad (7.4.6)$$

It can be shown that although equation (7.4.6) is mathematically equivalent to equation (7.4.2) it is numerically superior. If λ is zero, all of the diagonal elements of $(\mathbf{J} \mathbf{S})^t \mathbf{Q} \mathbf{J} \mathbf{S} + \lambda \mathbf{S}^t \mathbf{S}$ are equal to unity. For a non-zero λ the diagonal elements of $(\mathbf{J} \mathbf{S})^t \mathbf{Q} \mathbf{J} \mathbf{S} + \lambda \mathbf{S}^t \mathbf{S}$ are greater than unity, though in general they will not be equal. Let the largest element of $\lambda \mathbf{S}^t \mathbf{S}$ be denoted as λ_p . This is the value that PEST reports as its Marquardt lambda. Then the largest diagonal element of $(\mathbf{J} \mathbf{S})^t \mathbf{Q} \mathbf{J} \mathbf{S} + \lambda \mathbf{S}^t \mathbf{S}$ is equal to $1 + \lambda_p$.

As well as being useful in adjusting the direction of parameter improvement during the early stages of a parameter estimation process where that process is not able to “see” the closed contours of the objective minimum at the bottom of the objective function valley, the Marquardt lambda has a secondary role. This is a role that it probably should not have, but can be extremely valuable nevertheless.

If the $\mathbf{J}^t \mathbf{Q} \mathbf{J}$ matrix of equation (7.4.2) cannot be inverted, then the addition of positive terms to its diagonal renders it invertible. The Marquardt lambda thus becomes a primitive regularization device. If the Marquardt lambda is assigned values on a trial-and-error basis (as is done by PEST), the amount of diagonal strengthening can be adjusted to suit the need for regularization at any particular point during the iterative parameter estimation process. Ideally, regularization should not be applied in such an ad hoc fashion. Nevertheless, where model nonlinear behaviour is such that sensitivities encapsulated in the \mathbf{J} matrix can vary greatly from iteration to iteration, then singularity of $\mathbf{J}^t \mathbf{Q} \mathbf{J}$ can occur unexpectedly. Temporary non-invertibility of this matrix may then bring the parameter estimation process to a halt. However with the Marquardt lambda restoring the invertibility status of $\mathbf{J}^t \mathbf{Q} \mathbf{J}$, another set of parameters can be calculated for use in the next iteration of the parameter estimation process; in many contexts this allows the overall parameter estimation process to then proceed without further problems.

In some modelling contexts the Marquardt lambda can play yet another useful role. If an environmental model has problematic numerical behaviour (for example if its solver experiences convergence difficulties), the integrity of finite-difference derivatives that are used to fill the Jacobian matrix may be compromised. Through trial-and-error variation of the Marquardt lambda, a direction of parameter improvement may nevertheless be found in spite of the fact that no parameter improvement at all may be forthcoming from direct application of equation (7.4.1).

7.4.2.3 Optimum Length of the Parameter Upgrade Vector

A parameter upgrade vector is calculated during every iteration of the inversion process. Actually, PEST calculates a number of different parameter upgrade vectors, each based on a

different value of the Marquardt lambda; the parameter upgrade vector which yields the largest reduction in the objective function provides parameters with which to start the next iteration.

Let the (Marquardt lambda dependent) parameter upgrade vector be designated as \mathbf{u} . Thus

$$\mathbf{u} = \mathbf{p} - \mathbf{p}_0 \quad (7.4.7)$$

The optimal length of this vector can be calculated using information that is available through local model linearization. Let the new set of parameters be actually calculated as

$$\mathbf{p} - \mathbf{p}_0 = \alpha \mathbf{u} \quad (7.4.8)$$

By varying α , the magnitude of parameter movement in the direction of the parameter upgrade vector can also be varied. As has already been discussed, the objective function Φ , whose task it is for the parameter estimation process to minimize, is calculated as

$$\Phi = \mathbf{r}^t \mathbf{Q} \mathbf{r} \quad (7.4.9)$$

Suppose that model outputs are \mathbf{o}_0 when model parameters are \mathbf{p}_0 . When parameters are updated to \mathbf{p} using (7.4.8) a linear approximation to updated residuals can be computed as

$$\mathbf{r} = \mathbf{h} - [\mathbf{o}_0 + \alpha \mathbf{J}(\mathbf{p} - \mathbf{p}_0)] = \mathbf{r}_0 - \alpha \mathbf{J} \mathbf{u} \quad (7.4.10)$$

where the \mathbf{r}_0 vector denotes current residuals (i.e. residuals corresponding to parameters \mathbf{p}_0). The objective function then becomes

$$\Phi = (\mathbf{r}_0 - \alpha \mathbf{J} \mathbf{u})^t \mathbf{Q} (\mathbf{r}_0 - \alpha \mathbf{J} \mathbf{u}) \quad (7.4.11)$$

After multiplication and collecting terms, this becomes

$$\Phi = \mathbf{r}_0^t \mathbf{Q} \mathbf{r}_0 - 2\alpha \mathbf{r}_0^t \mathbf{Q} \mathbf{J} \mathbf{u} + \alpha^2 \mathbf{u}^t \mathbf{J}^t \mathbf{Q} \mathbf{J} \mathbf{u} \quad (7.4.12)$$

If this equation is now differentiated with respect to α , and the differential set to zero (as this minimizes Φ), we obtain the optimal factor by which to multiply the length of the parameter upgrade vector. This is

$$\alpha = \frac{\mathbf{r}_0^t \mathbf{Q} \mathbf{J} \mathbf{u}}{\mathbf{u}^t \mathbf{J}^t \mathbf{Q} \mathbf{J} \mathbf{u}} \quad (7.4.13)$$

7.4.2.4 Broyden Jacobian Update

Each iteration of the nonlinear parameter estimation process is effectively divided into two parts. During the first part of each iteration, the Jacobian matrix is calculated. During the second part of each iteration, the optimal direction and length of the parameter upgrade vector are determined. As implemented in PEST, this second segment of each iteration requires that a number of model runs be undertaken as different values assigned to the Marquardt lambda are tested in order to establish that which is most effective in reducing the objective function. The best set of upgraded parameters, as well as the lambda that achieved it, are then inherited by the next iteration of the nonlinear parameter estimation process.

When a model is run in order to test a parameter upgrade, this provides an opportunity to compare changes to model outputs computed using the Jacobian matrix to those which actually occur. The difference between predicted and actual changes can then be used to partially correct the Jacobian matrix so that it is more effective in upgrading parameters in directions that are similar to that of the parameter upgrade that was just attempted. If the directions of parameter upgrades computed on the basis of different Marquardt lambdas are not too different from each other, improvements made to the Jacobian matrix based on lessons learned in testing one

Marquardt lambda may make calculation of the next upgrade vector based on the next Marquardt lambda more effective.

Let $(\mathbf{p}_n - \mathbf{p}_0)$ denote the upgrade vector calculated using one particular value of the Marquardt lambda. The subscript “n” stands for “new” while the “0” subscript indicates current parameter values. If the model were linear, altered model outputs could be calculated from the existing Jacobian matrix as

$$\mathbf{o}_n = \mathbf{o}_0 + \mathbf{J}(\mathbf{p}_n - \mathbf{p}_0) \quad (7.4.14)$$

Let model outputs as calculated by the actual model be designated as \mathbf{o}_m . The difference between actual and linear-predicted model outputs, divided by the norm of the difference in parameter values that was responsible for the change in outputs, is calculated as

$$\mathbf{w} = \frac{\mathbf{o}_m - \mathbf{o}_0 - \mathbf{J}(\mathbf{p}_n - \mathbf{p}_0)}{(\mathbf{p}_n - \mathbf{p}_0)^t(\mathbf{p}_n - \mathbf{p}_0)} \quad (7.4.15)$$

An updated Jacobian matrix \mathbf{J}_n can then be calculated as

$$\mathbf{J}_n = \mathbf{J} + \mathbf{w}(\mathbf{p}_n - \mathbf{p}_0)^t \quad (7.4.16)$$

For the i ’th model output o_i of \mathbf{o} , equation (7.4.16) distributes the correction to the Jacobian matrix among parameters in proportion to the relative amount that each parameter is featured in the upgrade direction.

7.4.2.5 Restricting Parameter Movement

Notwithstanding use of the Marquardt lambda to accommodate nonlinear model behaviour, calculation of parameter upgrades still relies heavily on the assumption of model linearity. This assumption is workable if parameters do not need to change too much during any iteration of the nonlinear parameter estimation process. However if a parameter change vector is calculated whose length grossly exceeds the region of applicability of the linearity assumption on which its calculation rests, then the objective function may rise instead of fall as the upgrade vector crosses the objective function valley and proceeds up the other side of the objective function hill. On the next iteration of the nonlinear parameter estimation process, a new parameter upgrade vector will be calculated which tries to correct the previous overshoot. The result may be over-correction and another overshoot of the objective function valley.

This problem can be prevented by placing limits on the length of the parameter upgrade vector. Different parameter estimation packages differ in how they impose this limit. PEST offers a number of options. The length of the upgrade vector may be limited such that

- the value of no parameter alters by more than a user-specified factor;
- the value of no parameter alters by more than a user-specified fraction of its current value;
- no parameter changes by more than a user-specified, parameter-specific amount.

In all cases, unless a user specifies otherwise, the direction of the parameter upgrade vector as calculated by the procedures described above is unaltered. Only the length of the upgrade vector is changed.

Another methodology for preventing parameter overshoot on successive iterations of a nonlinear inversion process requires detection of parameters whose changes on successive iterations are of opposite sign, and then subjecting these parameters to especially strong damping. PEST implements damping of this kind using a method described by Cooley (1983)

and implemented in Hill (1992).

7.4.2.6 Parameter and Observation Transformations

The more linear is a model's behaviour, the more efficient is the Gauss-Marquardt-Levenberg method in minimizing the objective function on which the calibration process is based. There are many modelling contexts in which inverse problem linearity can be more closely approached if parameters and/or observations are appropriately transformed. Logarithmic transformation of parameters is often highly effective in this regard, particular for parameters such as permeability and electrical resistivity whose values can span many orders of magnitude.

Sometimes particular care needs to be taken in definition of observations. For example, rate of springflow may comprise part of a measurement dataset used to calibrate a groundwater model. However, if model-calculated heads fall below the level of the spring, then model-calculated springflow is zero; at the same time, the sensitivity of model-calculated springflow to changes in parameter values is zero because no springflow actually occurs as far as the model is concerned. Problems associated with "asymmetric observations" of this type, and means by which they can be ameliorated, are discussed in chapter 16 of this book. For the spring example, the actual springflow observation may need to be supplemented with an observation that groundwater head is at the elevation of the spring. The residual for this supplementary observation will be positive if the model-calculated head is below the level of the spring, and the spring is therefore not flowing; this ancillary observation (which is spatially coincident with the springflow observation) will indeed be sensitive to at least some model parameters when the springflow observation is not.

7.4.2.7 Parameter Bounds

Bounds may be placed on values taken by parameters during the inversion process for one of two reasons. One reason is to prevent them from adopting values that will cause a model to crash. For example, a groundwater or subsurface reservoir model is unlikely to tolerate a negative value supplied for porosity; similarly, a surface water model is unlikely to tolerate a negative value supplied for soil moisture store volume. The other reason for which a modeller may wish to place bounds on values taken by parameters is to prevent the parameter estimation process from endowing parameters with values which a modeller believes to be unrealistic.

Numerically, the imposition of bounds on parameters requires more than simply replacing the value of a parameter that has transgressed its bound with the value of the bound that it has transgressed. The bounds imposition algorithm implemented by PEST is more sophisticated than this. If a parameter upgrade vector \mathbf{u} is calculated which would cause a parameter to move beyond its upper or lower bound, PEST reduces the length of the upgrade vector \mathbf{u} such that this does not occur. The offending parameter is thereby placed at its bound. On later iterations, special treatment is then provided for this, and any other parameters, which are at their allowed limits. If the components of both the upgrade vector and the negative of the gradient vector pertaining to a parameter which is at its upper or lower limit are such as to take the parameter further out of bounds, then the parameter is temporarily frozen at its bound and the parameter estimation problem is reformulated accordingly with that parameter omitted from the reformulated inverse problem; hence the new upgrade vector will not result in any adjustment of that parameter. If, after reformulation of the problem in this manner, there are other parameters at their limits for which the parameter upgrade vector still points outward, but for which the negative of the gradient vector points inward, then these parameters, too, are temporarily frozen at their respective bounds. This process continues until a parameter upgrade

vector is calculated which either moves parameters from their bounds back into the allowed parameter domain, or leaves them fixed.

The strength of this strategy is that it allows PEST to search along the boundaries of the bounded parameter domain for the smallest objective function to which it has access, in spite of the fact that the global minimum of the objective function may lie outside that part of parameter space to which PEST is granted access.

At the beginning of each new iteration of the nonlinear parameter estimation process all temporarily-frozen parameters are freed. This allows them to move back inside the permitted part of parameter space if solution of equation (7.4.2) deems this necessary. The stepwise, temporary freezing of parameters is then repeated as described above.

Figure 7.7 schematizes parameter trajectories where the objective function minimum lies outside that part of parameter space wherein parameters are allowed to roam.

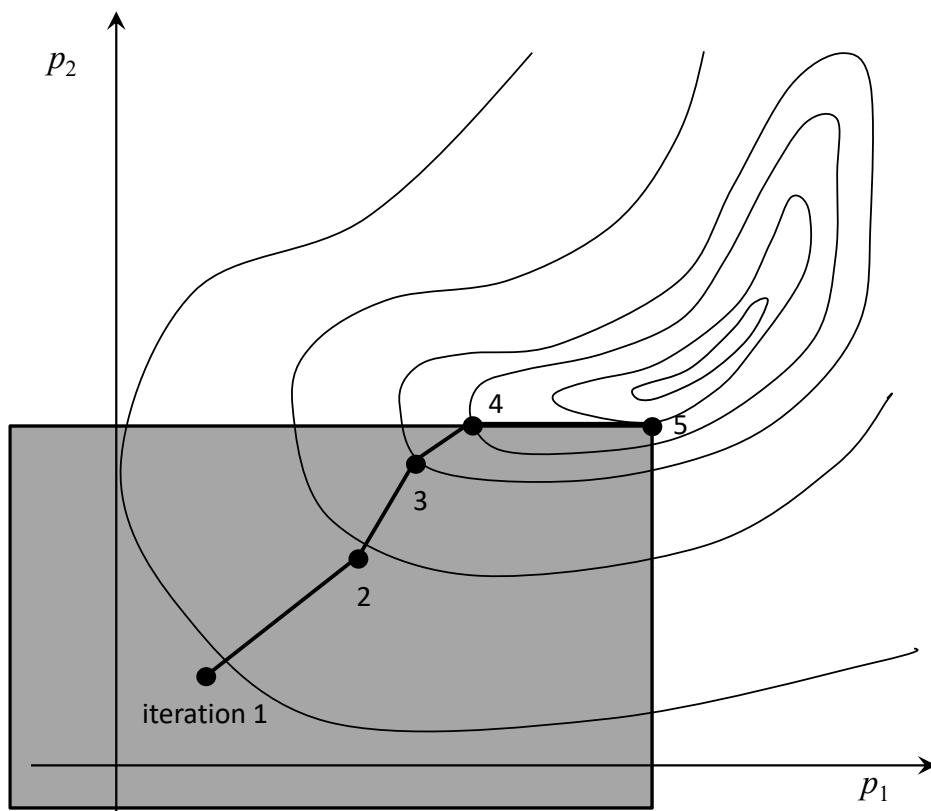


Figure 7.7. The shaded region shows that part of parameter space that is bounded by user-supplied upper and lower parameter limits. Also shown is the trajectory taken by parameters during the parameter estimation process.

Care should be taken when imposing bounds on parameters. In many instances of manually-regularized parameter estimation, bounds act as a de-facto regularization device. It is often problematic parameters (i.e. parameters whose composite sensitivities are small, or that share occupancy of a parameter covariance matrix eigenvector whose eigenvalue is unduly large) that equation (7.4.2) deems to require the largest alterations. This, of course, follows from the fact that large alterations to individually or collectively insensitive parameters are required if the objective function is to improve. More often than not, these parameters are virtually inestimable, and should have been fixed at their prior expected values prior to embarking on the parameter estimation process. However if a wayward parameter encounters its bound, and

is effectively removed from the parameter estimation process, this has the same regularizing effect as that which a modeller would introduce to the parameter estimation process if he/she had fixed the parameter him/herself. However fixing the parameter at its limiting value constitutes a far inferior form of regularization than fixing the parameter at its prior expected value. “Regularization by bounds” is thus not a recommended procedure, as estimated parameter values can hardly then claim a status of minimum error variance.

7.5 Critique of Manual Regularization

In the next chapter of this book the focus shifts to mathematical regularization. Meanwhile, the present chapter has assumed that parameter uniqueness is achieved through manual regularization. In the author’s opinion, mathematical regularization provides a far better path to parameter uniqueness. This is because a solution to an inverse problem achieved through mathematical regularization can normally lay far greater claim to a status of minimum error variance than that achieved through manual regularization. Furthermore, when regularized mathematically, model calibration can take place in a highly parameterized context (instead of a parsimoniously-parameterized parameter context) wherein the process complexity of environmental simulators is partnered with parameterization complexity that respects the natural heterogeneity of the real world.

In addition to this, representation of parameterization complexity on which a prediction may depend is fundamental to quantifying the uncertainty of that prediction. As has already been stated, parameters should be represented in a model whether they are estimable or not. It is their relevance to decision-critical model predictions that matters most when deciding whether to include or omit parameters from an inversion process.

Prior to introducing mathematical regularization, the shortcomings of manual regularization are first discussed in greater detail.

7.5.1 Man-Made Structural Noise

This chapter has discussed how inverse problem well-posedness can be achieved manually. It has also discussed how some of the vectors and matrices that emerge from a manually-regularized inversion process can be analysed to check whether the problem is in fact well-posed. Analyses presented in this chapter are based on the concept that the following equation describes the action of a model under calibration conditions:

$$\mathbf{h} = \mathbf{Xp} + \boldsymbol{\varepsilon} \quad (7.5.1)$$

They also assume that a statistical description of measurement noise is available, and that this is encapsulated in a covariance matrix $C(\boldsymbol{\varepsilon})$. As has been discussed, knowledge of $C(\boldsymbol{\varepsilon})$ is important. It provides the basis for choice of a weighting matrix \mathbf{Q} that supports the search for a minimum error variance solution \mathbf{p} to the manually regularized inverse problem. It also underpins quantification of that error variance. The latter is expressed through the post-calibration parameter error covariance matrix $C(\mathbf{p} - \mathbf{p})$.

The present chapter has also shown how the Gauss-Marquardt-Levenberg method can be used to solve the inverse problem that is posed by equation (7.5.1). Other methods can also be used to seek solutions to this problem, including so-called global methods that do not rely on calculation of a Jacobian matrix. Generally these are not nearly as numerically efficient as the Gauss-Marquardt-Levenberg method; however some of them can accommodate high degrees of model nonlinearity and poor model numerical behaviour.

Alternatively, a modeller may dispense with the quest for uniqueness and employ a Bayesian

methodology such as Markov chain Monte Carlo to sample the posterior probability distribution of \mathbf{p} ; the mean of this posterior distribution is the \mathbf{p} that is sought in the present chapter.

Regardless of the method that is employed to solve the problem posed by (7.5.1), and regardless of the methodology that is adopted to quantify the uncertainty, or potential for error, in that solution, the imperative to weight individual members of a calibration dataset in a way that reflects their relative credibility remains. Thus a modeller cannot avoid the task of adopting either a $C(\boldsymbol{\varepsilon})$ matrix, or a matrix that is proportional to it; in the latter case the constant of proportionality σ^2_r can be estimated through the inverse problem solution procedure. Normally $C(\boldsymbol{\varepsilon})$ is specified as a diagonal matrix; this choice rests on the assumption that errors afflicting measurements comprising a calibration dataset are independent of each other.

Calibration of just about any environmental model inevitably leads to the discovery that the level of model-to-measurement misfit achieved through the calibration process is far greater than can be explained by measurement noise alone. In fact, misfit is often dominated by a so-called structural component. This reflects the fact that a numerical model is a less-than-perfect simulator of complex environmental processes that are operative at a particular study site. It also reflects the fact that the parameterization scheme employed by an environmental model does not encapsulate the complex spatial variability of physical/chemical properties that characterize the simulated system.

Equation (7.1.10) depicts structural noise arising from parameter simplification. A challenge that faces manual regularization is how to make this term as small as possible while still achieving a matrix \mathbf{X} to use in place of the “real” model matrix \mathbf{Z} , keeping in mind that the matrix $\mathbf{X}^t \mathbf{Q} \mathbf{X}$ must be invertible (and that \mathbf{X} does therefore not possess a null space). However this is not the only condition that \mathbf{X} must satisfy. As discussed above, not only must $\mathbf{X}^t \mathbf{Q} \mathbf{X}$ be invertible; post-calibration parameter error variances must be smaller than pre-calibration parameter variances.

Manual attainment of an appropriate \mathbf{X} matrix is often hard work, requiring much laborious trial and error. On the one hand, a modeller wants to reduce structural noise to a minimum. On the other hand, all attempts to do so by increasing the number of parameters that are estimated through the inversion process increase the chances that the resulting \mathbf{X} matrix will either be rank-deficient (i.e. it will possess a null space), or that the error variance of some parameters will be amplified rather than diminished through the inversion process.

So the question must be asked “why manufacture structural noise through manual parameter simplification when mathematical regularization can promise far lower levels of regularization-induced structural noise?” Of course parameter simplification is not the only model structural defect whose presence can engender structural noise; after all, the model itself is a simplification of reality. To mitigate structural noise arising from model defects, modellers often deploy complex simulators that are sometimes characterized by large run times and problematic numerical behaviour. The above question then becomes even more salient. Why drape a simplistic parameterization scheme over the domain of a complex model, thereby eroding the capacity of that model to replicate the historical behaviour of the environmental system that it purports to represent?

In all of the theory that is presented in this chapter, regularization-induced structural noise has been ignored. As previously stated, this is represented by $\boldsymbol{\eta}$ of equation (7.1.10). If an appropriate weighting matrix \mathbf{Q} must be developed in order to achieve parameter estimates of minimum error variance, then it is $C(\boldsymbol{\eta} + \boldsymbol{\varepsilon})$ (i.e. $C(\boldsymbol{\tau})$ from equation (7.1.11)) that must be

characterized and then inverted, and not $C(\boldsymbol{\varepsilon})$. This requires that $C(\boldsymbol{\eta})$ be specified.

Sadly, calculation of $C(\boldsymbol{\eta})$ is rarely possible. From (7.1.10) $C(\boldsymbol{\eta})$ can be conceptually calculated as

$$C(\boldsymbol{\eta}) = (\mathbf{Z}\mathbf{L} - \mathbf{X})C(\mathbf{p})(\mathbf{Z}\mathbf{L} - \mathbf{X})^t \quad (7.5.2)$$

If $C(\boldsymbol{\eta})$ is calculated using equation (7.5.2) pertinent Jacobian matrices must replace \mathbf{Z} and \mathbf{X} . \mathbf{X} is easily calculated; in fact it must be calculated routinely to estimate parameters \mathbf{p} through the Gauss-Marquardt-Levenberg method. The same is not the case for \mathbf{Z} as \mathbf{Z} includes all model parameters – not just those which are adjusted through the calibration process; it includes those that are fixed and tied through manual regularization. Its calculation is therefore much more laborious than that of \mathbf{X} . And even if it were to be calculated, another question immediately arises. Why not use mathematical regularization methods discussed in the following chapter to estimate \mathbf{k} of equation (7.1.1) instead of a parsimonious set of parameters \mathbf{p} based on a regularization strategy that may be far from optimal?

The $C(\mathbf{p})$ term that is featured in equation (7.5.2) is even more problematic. As was discussed in section 7.1, the relationship between \mathbf{p} and the real-world parameter set \mathbf{k} is unknown. Unfortunately, it is \mathbf{k} to which expert knowledge can be applied, and not \mathbf{p} . Hence while it is conceptually possible for an expert to arrive at a suitable $C(\mathbf{k})$ matrix, the same cannot be said for $C(\mathbf{p})$. We will return to this subject shortly.

There is yet another problem. Doherty and Welter (2010) show that $C(\boldsymbol{\eta})$ is generally singular. This is not really surprising as it is not noise at all, but (as Doherty and Welter point out), “information with nowhere to go”. Information has structure; if it were random, it would not be information.

The singular status of $C(\boldsymbol{\eta})$ can also be gleaned intuitively from equation (7.5.2). Presumably the small-dimensional covariance matrix $C(\mathbf{p})$ is non-singular, for otherwise it cannot be a covariance matrix. The vector $\boldsymbol{\eta}$ must possess more elements than that of \mathbf{p} , as observations must outnumber parameters for an inverse problem to be well-posed. If \mathbf{p} is random then $\boldsymbol{\eta}$ is random. However it is conceptually impossible for its variability to possess more degrees of freedom than that of \mathbf{p} because it is calculated from \mathbf{p} . $C(\boldsymbol{\eta})$ must therefore be singular.

Fortunately it is $C(\boldsymbol{\eta} + \boldsymbol{\varepsilon})$, i.e. $C(\boldsymbol{\tau})$, that requires inversion for formulation of a structural-noise-accommodating \mathbf{Q} and not $C(\boldsymbol{\eta})$. Nevertheless, $C(\boldsymbol{\tau})$ will approach singularity where measurement noise is small. Numerical calculation of its inverse $C^{-1}(\boldsymbol{\tau})$ may then be accompanied by numerical error.

In light of all of the above considerations, achievement of a minimum error variance solution to a manually regularized inverse problem, and quantification of post-calibration parameter error variance, is a theoretically challenging undertaking. It may certainly be possible to achieve a good fit between model outputs and field measurements using the Gauss-Marquardt-Levenberg method (or some other method). But assessment of how wrong estimated parameters may be becomes subjective at best, and impossible at worst.

It is worth noting that these problems are not confined to calibration; they apply equally to direct posterior uncertainty analysis. Suppose that a Markov chain Monte Carlo (MCMC) methodology is employed to directly sample the posterior probability distribution of \mathbf{p} . This may be preceded by parameter simplification in order to enable use of MCMC without a prohibitively large number of model runs. For the same reasons as those outlined above, the integrity of MCMC sampling of the posterior probability distribution of \mathbf{p} is questionable. Failure to characterize the prior probability distribution of \mathbf{p} , and failure to properly

characterize the likelihood function, given the impact of structural noise on this function, are the reasons for this loss of integrity.

7.5.2 Parameter and Predictive Error Variance at the Scale that Matters

Instead of achieving parameter simplification by defining \mathbf{p} in terms of \mathbf{k} , equation (7.1.2) achieves parameter simplification by defining \mathbf{k} in terms of \mathbf{p} . This is a matter of practicality, and cannot be avoided. A complex model requires \mathbf{k} parameters; these can number in the hundreds or even thousands (or even greater if they pertain to the scale of a model grid). As was discussed in section 7.1, a simplistic \mathbf{p} -based parameterization scheme may comprise, for example, zones of piecewise constancy (in which case groups of \mathbf{k} -parameters are assumed to be equal to each other), or parameters combined according to fixed ratios (in which case groups of \mathbf{k} -parameters are assumed to vary in harmony), or parameters which are fixed (in which case some \mathbf{k} -parameters are not estimated at all). In all cases a modeller defines the groups, supplies the ratios, and fixes non-adjustable \mathbf{k} parameters at appropriate values. Each adjustable parameter or parameter group that emerges from this process defines an element p_i of \mathbf{p} . The inversion process then assigns values to elements of \mathbf{p} . Corresponding values for \mathbf{k} are calculated for use by the model on the basis of the user-specified \mathbf{L} matrix of equation (7.1.2).

Suppose that \mathbf{p} has been estimated using equation (7.2.14) or one of its nonlinear counterparts. A modeller will have chosen a suitable weight matrix \mathbf{Q} , notwithstanding the conceptual difficulties implied in such a choice discussed above, and will have proceeded with the parameter estimation process using this \mathbf{Q} . Once \mathbf{p} has been estimated, the \mathbf{k} -scale estimated parameter set \mathbf{k} that is actually provided to the model is calculated as

$$\underline{\mathbf{k}} = \mathbf{L}\underline{\mathbf{p}} = \mathbf{L}(\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \mathbf{h} = \mathbf{G} \mathbf{h} \quad (7.5.3)$$

The matrix \mathbf{G} through which $\underline{\mathbf{k}}$ is estimated from \mathbf{h} is defined through the above equation. If (7.1.1) is substituted for \mathbf{h} in (7.5.3) we obtain

$$\underline{\mathbf{k}} = \mathbf{L}\underline{\mathbf{p}} = \mathbf{L}(\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \mathbf{Z} \mathbf{k} + \mathbf{L}(\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \boldsymbol{\epsilon} \quad (7.5.4)$$

Suppose that measurement noise (“real” measurement noise this time, and not the mixture of measurement and parameter-simplification-induced structural noise that dominates model-to-measurement misfit under a regime of \mathbf{p} -scale manual regularization) is zero. Equation (7.5.4) then becomes

$$\underline{\mathbf{k}} = \mathbf{L}(\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \mathbf{Z} \mathbf{k} \quad (7.5.5)$$

This equation shows the relationship between estimated parameters $\underline{\mathbf{k}}$ and true system parameters \mathbf{k} . These, and not the parsimonized parameters \mathbf{p} , are the parameters that matter, for these are the parameters that the model uses when it makes predictions of future system behaviour.

Equation (7.5.5) defines the resolution matrix that arises from manual regularization. Specifically

$$\mathbf{R} = \mathbf{L}(\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \mathbf{Z} \quad (7.5.6)$$

Recall from section 2.9 that the resolution matrix defines the relationship between estimated and real parameters. Recall also that each row of the resolution matrix defines the conceptual parameter averaging process that is implied in estimating each element k_i of $\underline{\mathbf{k}}$ through an ill-posed inverse problem. It demonstrates that the value of each such estimated element is, in fact, an average over many real parameter elements. A modeller cannot know \mathbf{k} ; all that he/she can obtain through the calibration processes is a “blurred”, “averaged” version of parameter reality,

and not reality itself.

From (7.5.4) the error in estimated \mathbf{k} -scale parameters can be calculated as

$$\underline{\mathbf{k}} - \mathbf{k} = -(\mathbf{I} - \mathbf{L}(\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \mathbf{Z}) \mathbf{k} + \mathbf{L}(\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \boldsymbol{\varepsilon} \quad (7.5.7)$$

That is

$$\underline{\mathbf{k}} - \mathbf{k} = -(\mathbf{I} - \mathbf{R}) \mathbf{k} + \mathbf{G} \boldsymbol{\varepsilon} \quad (7.5.8)$$

where \mathbf{R} is defined above and

$$\mathbf{G} = \mathbf{L}(\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \quad (7.5.9)$$

If \mathbf{k} is statistically independent of $\boldsymbol{\varepsilon}$ (which is a reasonable assumption), then the covariance matrix of post-calibration parameter error can be calculated as

$$\mathbf{C}(\underline{\mathbf{k}} - \mathbf{k}) = (\mathbf{I} - \mathbf{R}) \mathbf{C}(\mathbf{k}) (\mathbf{I} - \mathbf{R})^t + \mathbf{G} \mathbf{C}(\boldsymbol{\varepsilon}) \mathbf{G} \quad (7.5.10)$$

Equation (7.5.10) is completely general. It does not depend on the type of regularization employed to achieve parameter uniqueness. The \mathbf{R} and \mathbf{G} matrices are regularization-specific; however the equation is not. A choice for \mathbf{R} and \mathbf{G} can be said to be optimal when the trace of $\mathbf{C}(\underline{\mathbf{k}} - \mathbf{k})$ is minimized. (Actually when model defects are considered, perhaps a more practical definition of optimality of regularization is that the error variance of predictions of interest are minimized. As is shown in chapter 15 of this text, where model defects are considered, minimization of parameter error variance and minimization of predictive error variance become two different things. For the moment however, we consider the \mathbf{Z} model to have no defects.)

What is immediately apparent from equation (7.5.10) is that computation of parameter error variance at the scale that matters - the scale of \mathbf{k} rather than the scale of \mathbf{p} - requires two terms instead of the single term embodied in equation (7.2.18). If \mathbf{Q} is chosen as $\mathbf{C}^{-1}(\boldsymbol{\varepsilon})$ (so that the reference variance is unity), the single term of equation (7.2.18) is indeed present as the second term on the right side of equation (7.5.10). However equation (7.5.10) has another term on its right side. This is the “cost of uniqueness” term. This term emerges from the necessity to simplify parameters prior to calibration in order to achieve calibration uniqueness. This term is zero only if \mathbf{R} is equal to \mathbf{I} (i.e. the identity matrix). This can happen only if \mathbf{k} is the same as \mathbf{p} (and \mathbf{Z} is the same as \mathbf{X} because \mathbf{L} is equal to \mathbf{I}), and hence the world is as simple as the model. Under any other circumstances, the cost of uniqueness must be taken into account in assessing the potential for parameter error at the scale that matters, i.e. the \mathbf{k} -scale. If it is not taken into account, then the error variances of at least some \mathbf{k} -scale parameters are likely to be grossly under-estimated because the first term of (7.5.10) is thereby ignored. For example, where a spatial model domain is endowed with zones of piecewise constancy as a manual regularization device, \mathbf{k} -scale error arises from the erroneous assumption that real world system properties are actually piecewise constant. To the extent that they are not, enforced piecewise homogeneity of \mathbf{k} -scale parameters creates a potential for \mathbf{k} -scale error. This is in addition to errors that arise in estimation of \mathbf{p} -scale parameters from the presence of noise in the calibration dataset.

Let s be a prediction of interest. Let the vector \mathbf{y} express the sensitivity of this prediction to parameters \mathbf{k} . Then the prediction is made as

$$s = \mathbf{y}^t \mathbf{k} \quad (7.5.11a)$$

In contrast, predictions made by the calibrated model are actually made using the estimated parameter set $\underline{\mathbf{k}}$. Let this prediction be characterized as \underline{s} . Hence

$$\underline{s} = \mathbf{y}^t \underline{\mathbf{k}} \quad (7.5.11b)$$

Predictive error is thus

$$\underline{s} - s = \mathbf{y}^t(\underline{\mathbf{k}} - \mathbf{k}) \quad (7.5.12)$$

From (7.5.10), the variance of predictive error can be written as

$$\sigma_{s-s}^2 = \mathbf{y}^t \mathbf{C}(\underline{\mathbf{k}} - \mathbf{k}) \mathbf{y} = \mathbf{y}^t (\mathbf{I} - \mathbf{R}) \mathbf{C}(\mathbf{k}) (\mathbf{I} - \mathbf{R})^t \mathbf{y} + \mathbf{y}^t \mathbf{G} \mathbf{C}(\boldsymbol{\varepsilon}) \mathbf{G} \mathbf{y} \quad (7.5.13)$$

Hence the cost of parameter simplification flows through to model predictions, and computation of potential errors associated therewith. The extent to which this affects any prediction is dependent on the prediction itself; that is, it is dependent on the sensitivity \mathbf{y} of that prediction to the different parameters \mathbf{k} used by the model.

If we continue to use zonal regularization as an example, then to the extent that any prediction s is sensitive to within-zone heterogeneity, that prediction may incur an error through suppression of recognition of that heterogeneity in the calibrated model. To be sure, within-zone heterogeneity may not be estimable regardless of the regularization device deployed for attainment of parameter uniqueness; the information content of the calibration dataset \mathbf{h} may simply not support such estimation. What is important, however, is that the parameters that express the potential for within-zone heterogeneity (i.e. the \mathbf{k} parameters) be retained in the model so that their inestimability can be taken into account when computing predictive error variance. This is done in equation (7.5.13). It is not done in equation (7.2.22) which is typically used to compute predictive error variance in the manual regularization context; equation (7.2.22) includes only the second term on the right of (7.5.13) and not the first, important, cost-of-uniqueness term.

Omission of \mathbf{k} parameters, and their replacement with \mathbf{p} parameters, carries a further disadvantage. In general, it cannot be known in advance whether the calibration dataset does indeed carry at least some information that can inform parameterization detail at the \mathbf{k} scale. Optimality of the calibration process cannot be assured if such information is refused entry into that process.

While the above discussion focuses on contexts wherein a model is first calibrated and parameter and predictive error variance are then quantified as surrogates for parameter and predictive uncertainty, arguments pertaining to parameter simplification are no less salient where simplification is undertaken as a precursor to using methods such as Markov chain Monte Carlo (MCMC), or even ensemble-based methods that are discussed later in this book, for direct exploration of posterior parameter and predictive uncertainty. As has already been discussed, the temptation to simplify parameterization in this context is sometimes overwhelming, as methods such as MCMC have difficulty in accommodating large numbers of parameters, especially where model run times are long. The cost of using a parameter set \mathbf{p} instead of a more complex parameter set \mathbf{k} which embodies all aspects of parameterization detail to which a prediction of interest may be sensitive, is expressed by the first term of equation (7.5.10) regardless of how uncertainty analysis is undertaken. As equation (7.5.13) demonstrates, this term contributes to predictive error. Its exclusion from a parsimoniously parameterized model results in underestimation of predictive uncertainty. If model predictive uncertainty analysis must be preceded by a set of parameter simplification assumptions in order for it to become numerically feasible, then sadly, its feasibility may have been purchased at the cost of its validity.

7.5.3 What is being Estimated?

In the previous section, the relationship between estimated $\underline{\mathbf{k}}$ and real \mathbf{k} is quantified. The present section addresses the issue of just what is the simplified parameter set \mathbf{p} and what is its

estimated counterpart $\underline{\mathbf{p}}$.

In a similar manner to what was done in the last section we can substitute (7.1.1) into (7.2.14) to obtain

$$\underline{\mathbf{p}} = (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \mathbf{Z} \mathbf{k} + (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \boldsymbol{\varepsilon} \quad (7.5.14)$$

Where measurement noise is zero, this becomes

$$\underline{\mathbf{p}} = (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \mathbf{Z} \mathbf{k} = \mathbf{R}' \mathbf{k} \quad (7.5.15)$$

\mathbf{R}' is a kind of resolution matrix for $\underline{\mathbf{p}}$. However, unlike a real resolution matrix \mathbf{R} which links $\underline{\mathbf{k}}$ to \mathbf{k} , \mathbf{R}' is not square. Nevertheless each row of the \mathbf{R}' matrix provides the averaging coefficients through which an element of $\underline{\mathbf{p}}$ is formed from the elements of \mathbf{k} .

The matrix \mathbf{L} appears implicitly in equation (7.5.15) through equation (7.1.5) which is repeated here for convenience

$$\mathbf{X} = \mathbf{Z} \mathbf{L} \quad (7.5.16)$$

It is important to note that the relationship between $\underline{\mathbf{p}}$ and \mathbf{k} is “discovered” through the calibration process; it is not pre-assigned (unless the methodology discussed in the following section is followed). Hence if, for example, \mathbf{L} is a matrix such as that depicted in equation (7.1.3) which forms the basis for subdivision of a model domain into zones of piecewise constancy, it does not follow that the averaging of elements of \mathbf{k} specified by the rows of \mathbf{R}' endows elements of $\underline{\mathbf{p}}$ with a status of average system property within their respective zones. In fact there is no reason whatsoever why the averaging process specified by any row of \mathbf{R}' should preclude contributions to the value of an element p_i of $\underline{\mathbf{p}}$ from elements k_i of \mathbf{k} that lie outside of its zone. See Moore and Doherty (2006), Gallagher and Doherty (2007), Watson et al (2013) and Manewell et al (2023a) for a further discussion of this issue, as well as examples. It is wrong, therefore, to inspect the outcomes of a calibration process which features zones of piecewise constancy and infer that the calibration process has provided estimates of the average system property within each zone. Similar considerations apply to regularization devices other than zones of piecewise constancy.

These considerations have consequences for parsimonization in general, whether this is undertaken as a form of manual regularization, or whether it is undertaken as a means of bestowing numerical feasibility on other types of analysis. The problem of defining a likelihood function following parameter simplification was discussed above. Another problem is that of defining the prior probability distribution of $\underline{\mathbf{p}}$. From the above discussion, it is obvious that a pre-calibration relationship between $\underline{\mathbf{p}}$ and \mathbf{k} cannot be assumed. The prior probability distribution of $\underline{\mathbf{p}}$ cannot therefore be calculated from that of \mathbf{k} . The two terms on the right of Bayes equations (see equation 6.2.4) are therefore compromised.

Where calibration-constrained parameter error analysis is undertaken by first manually regularizing a model’s parameter set, by then calibrating the model, and by finally analysing the post-calibration propensity for parameter error through calculation of $C(\underline{\mathbf{p}} - \mathbf{p})$ (see equation 7.2.18), similar considerations apply. The question of “just what is this $\underline{\mathbf{p}}$ whose propensity for error we are quantifying” does not go away. In many contexts it matters little, as most environmental models are built to make predictions; they are not necessarily built for estimation of system properties. However if this is the case, then to the extent that a prediction of interest is sensitive to parameterization detail that is represented in \mathbf{k} but not in $\underline{\mathbf{p}}$, \mathbf{k} must feature in any strategy that is employed to quantify the uncertainty of that prediction.

7.5.4 Cooley's Strategy

While focusing on calibration of a groundwater model with a heterogeneous hydraulic property field, Cooley (2004) and Cooley and Christensen (2006) proposed a calibration strategy that can avoid some of the problems outlined above, while still employing manual regularization to achieve inverse problem well-posedness. Their theory also accommodates nonlinear model behaviour when applied to calculation of post-calibration parameter and predictive error variance.

As usual, let \mathbf{k} designate the “true parameterization” of a model domain, and let \mathbf{p} designate a parsimonious set of parameters that is used in its place. Suppose that we desire that \mathbf{p} have a pre-defined relationship to \mathbf{k} . Thus this relationship is designated before calibration; it is not something that can only be “discovered” after calibration using a resolution matrix. Suppose that we specify this relationship as follows

$$\mathbf{p} = \mathbf{N}\mathbf{k} \quad (7.5.17)$$

For the case where \mathbf{p} pertains to zones of piecewise constancy, each element p_i of \mathbf{p} may be pre-defined as the average of all elements k_i of \mathbf{k} within its zone; the \mathbf{N} matrix of (7.5.17) must therefore perform the desired spatial averaging.

More generally, \mathbf{N} may be related to \mathbf{L} of equation (7.1.2) through

$$\mathbf{N} = (\mathbf{L}^t \mathbf{L})^{-1} \mathbf{L} \quad (7.5.18a)$$

or

$$\mathbf{N} = (\mathbf{L}^t \mathbf{Q} \mathbf{L})^{-1} \mathbf{L}^t \mathbf{Q} \quad (7.5.18b)$$

Direct substitution into equation (2.9.5) confirms that \mathbf{N} comprises a generalized inverse of \mathbf{L} . At the same time, from equation (7.2.14), \mathbf{p} can be considered to be a simplified parameter set for which the complementary model parameter field calculated from it using (7.1.2) is the best approximation (in the weighted least squares sense) to the original complex parameter field \mathbf{k} from which \mathbf{p} is derived using equation (7.5.17). That is to say, after \mathbf{p} is calculated from \mathbf{k} using (7.5.17), a simplified \mathbf{k} -scale parameter field \mathbf{j} can be calculated from \mathbf{p} using (7.1.2) as

$$\mathbf{j} = \mathbf{L}\mathbf{p} \quad (7.5.19)$$

Through (7.5.18), \mathbf{j} becomes the closest approximation to \mathbf{k} in the weighted least squares sense. Where \mathbf{p} represents zones of piecewise constancy, the elements of \mathbf{p} can thus represent the spatial average of those elements of \mathbf{k} that lie within the zone defined by each element of \mathbf{p} . At the same time, \mathbf{j} parameters are uniform within each \mathbf{p} -zone.

A workflow similar to the following is then implemented.

1. Generate a stochastic realization of \mathbf{k} using the $C(\mathbf{k})$ prior covariance matrix of \mathbf{k} .
2. Run the model using \mathbf{k} . Record model outputs \mathbf{o} corresponding to the measurement dataset \mathbf{h} that will be used to calibrate the model. Add a realization of measurement noise to \mathbf{o} to generate a “pseudo \mathbf{h} ”.
3. Calculate \mathbf{j} from \mathbf{k} using (7.5.17) followed by (7.5.19). Where zones of piecewise constancy are employed as a manual regularization device, this amounts to assigning all \mathbf{j} elements within a particular zone a value equal to the average of all \mathbf{k} elements within that same zone. Run the model using the \mathbf{j} parameter set and obtain values for model outputs that correspond to \mathbf{h} .
4. Repeat steps 1 to 3 many times. Then calculate an empirical $C(\mathbf{r})$ matrix based on differences between outputs of the \mathbf{k} -model and corresponding outputs of the \mathbf{j} -model.

(See equation 7.1.11 for definition of τ as measurement noise plus structural noise.)

5. Estimate \mathbf{p} from the real-world \mathbf{h} using (7.2.14), with \mathbf{Q} calculated as $\mathbf{C}^{-1}(\tau)$.
6. Calculate parameter and predictive error variance using equations (7.2.18) and (7.2.22), optionally with appropriate enhancements to accommodate model nonlinearity that are described by Cooley (2004) and Cooley and Christensen (2006).

Cooley (2004) and Cooley and Christensen (2006) show that an estimate of \mathbf{p} obtained in this way constitutes the minimum error variance estimate of \mathbf{Nk} . Where regularization is based on replacing geological heterogeneity by zones of piecewise constancy, a modeller thus specifies before undertaking the calibration process that he/she would like to estimate the average value of system properties within each zone. The minimum error variance estimate of these average properties is then obtained.

An advantage of this approach to manual regularization is that it recognizes the existence of system property complexity behind the parsimonious parameterization scheme that is used to calibrate a model. It attempts to make a formal, pre-defined connection between these two levels of parameterization.

Though conceptually interesting and mathematically insightful, the methodology has not found much use in everyday modelling practice. Reasons include the following.

- Modellers must decide on the number of \mathbf{p} parameters to estimate before the parameter estimation process begins rather than have mathematical regularization determine the optimal dimensionality of the calibration solution space itself. A modeller therefore cannot know until post-calibration statistics are available whether he/she is attempting to estimate too many or too few parameters.
- If measurement noise is small in comparison to structural noise induced through parameter simplification, many model runs are required to calculate a stable empirical estimate of the $\mathbf{C}(\tau)$ matrix because of near singularity of this matrix.
- While minimum error variance estimates of the elements of \mathbf{p} , defined according to a modeller's specifications, may be obtained through this process, this is often of secondary or tertiary importance to obtaining minimum error variance estimates of predictions of management interest. The above methodology does not try to minimize structural noise. Instead it embraces it. Error variances of important predictions may be larger than they need to be, given the information content of the calibration dataset and the information content of expert knowledge.
- While the methodology presented by Cooley (2004) and Cooley and Christensen (2006) quantifies structural noise induced by parameter simplification, other sources of structural noise are ignored. In real world modelling contexts, these may be considerably greater in magnitude than those induced by parameter simplification. A weight matrix that is specified as $\mathbf{C}^{-1}(\tau)$, where τ includes only parameter-simplification-induced structural noise, may therefore fail to guarantee minimum error variance of \mathbf{p} as an estimate of \mathbf{Nk} .

7.5.5 The Role of Prior Information

As has been discussed, a calibration dataset can be expanded to include expert knowledge of parameter values that exists prior to the calibration process. Knowledge of this type should always be available; however, uncertainties that are associated with it will, of course, vary from case to case. Ideally, there is no reason why expert knowledge should not be included in the parameter estimation process as a matter of course, especially where the information content of a calibration dataset is limited, and parameter nonuniqueness therefore prevails. Practically,

however, there are difficulties in doing this where regularization is undertaken manually.

The first (and probably most significant) problem with use of prior information in the manual regularization process is that of weighting it. As has been discussed extensively in this chapter, weights should ideally be inversely proportional to the “noise” associated with measurements comprising a calibration dataset. In the case of prior information this “noise” is the uncertainty of that information. This is where the first problem arises; how does an expert quantify the uncertainty associated with his/her expert knowledge? Where prior information assigns a preferred value to each parameter (presumably its prior mean), this uncertainty is (ideally) informed by the prior parameter probability distribution. However hyperparameters which govern this distribution are generally uncertain themselves.

The second problem is that of maintaining correct relativity between prior information weights and weights assigned to measurements of system state. The problem here is that model-to-measurement misfit is not an outcome of measurement noise; so specifications of measurement instruments cannot be used to estimate it. In fact, the amount of “measurement noise” that is associated with a calibration dataset is usually estimated through the calibration process itself; see equation (7.2.23). Furthermore, this estimate pertains to the calibration dataset as a whole, and not to different components of it. Hence even if a modeller can quantify the uncertainty of his/her expert knowledge, it is very difficult for him/her to ensure correct relativity between weights assigned to prior information on the one hand, and to members of a measurement dataset on the other hand when the noise associated with the latter is not known until the calibration process is complete.

This is an unfortunate situation because correct relativity of weighting between measurements of system behaviour on the one hand, and prior information on the other hand, that collectively comprise a calibration dataset is important. If prior information is weighted too heavily, the inversion process will ignore measurements of system behaviour in favour of prior information. If prior information is not weighted heavily enough, it may not exert a sufficiently strong regularizing influence on the inversion process. The latter may remain ill-posed; nonuniqueness will then prevail, as will numerical difficulties that attend inversion of near-singular matrices.

Yet another problematic aspect of using prior information in the manual regularization context is that of defining the meaning of individual members of the parsimonized parameter set that is being estimated. Equation (7.5.15) shows that individual members of the parsimonized parameter set bear complex relationships to the real-world parameters to which expert knowledge actually pertains. Furthermore, this relationship is rarely quantifiable as equation (7.5.15) includes a high-dimensional sensitivity matrix (the \mathbf{Z} matrix) that is generally not available because it pertains to fine scale parameters (i.e. \mathbf{k} -scale parameters) that are not actually featured in the manually regularized inversion process. In addition to this, the averaging process that is implied by rows of the pseudo-resolution matrix \mathbf{R}' of equation (7.5.15) may reference fine scale parameters of multiple types in the characterization of one particular, parsimonized \mathbf{p} -scale parameter. It follows that just because a simplified parameter used in the manually-regularized inversion process is named after a particular system property, this does not mean that it is being estimated as such. The difficulties of associating expert knowledge with such a parameter, and of weighting such an association, are therefore compounded.

The following chapters of this text abandon \mathbf{p} -scale parameters in favour of \mathbf{k} -scale parameters to which, hopefully, expert knowledge is more directly applicable. This creates the context wherein expert knowledge can indeed play an important role in regularized inversion.

7.5.6 Information Criteria Statistics

The following post-calibration statistics are sometimes employed to assist a modeller in deciding how many parameters to include in a parsimonious parameterization scheme on which manual regularization is based. They are also sometimes used to assist a modeller in ranking parsimoniously calibrated models. This ranking allows weights to be assigned to predictions made by these models so that they can be used collectively in a predictive uncertainty analysis exercise in which different system conceptualizations are taken into account. See Poeter and Hill (2007) for more details.

AIC: Akaike Information Criterion

$$\text{AIC} = n \ln(\sigma^2) + 2k \quad (7.5.20)$$

AICc: second order bias-corrected AIC

$$\text{AICc} = n \ln(\sigma^2) + 2k + \left(\frac{2k(k+1)}{n-k-1} \right) \quad (7.5.21)$$

BIC: Bayesian Information Criterion

$$\text{BIC} = n \ln(\sigma^2) + k \ln(n) \quad (7.5.22)$$

KIC: Kashyap Information Criterion

$$\text{KIC} = (n - (k - 1)) \ln(\sigma^2) - (k - 1) \ln(2\pi) + \ln|\mathbf{J}^t \mathbf{Q} \mathbf{X}| \quad (7.5.23)$$

In all of these formulas

- n is the number of observations and prior information equations employed in the parameter estimation process;
- k is equal to $m+1$ where m is the number of adjustable parameters involved in the parameter estimation process;
- σ^2 is Φ_{\min}/n , where Φ_{\min} is the minimized objective function;
- \mathbf{X} is the Jacobian matrix computed on the basis of optimised (or almost optimised) parameters;
- \mathbf{Q} is the weight matrix.

The formula for each of these statistics has two terms. The first decreases as the objective function falls, and the level of model-to-measurement fit thereby improves. The second rises as the number of parameters rises. Supposedly, the optimal number of parameters to employ in the calibration process is that which fits the measurement dataset well enough to have extracted its information, but not so well that over-fitting has occurred and parameters are thereby “informed” by measurement noise rather than by measurement information. This “sweet spot” is supposed to occur where the various information criteria listed above are minimized with respect to the number of parsimonized parameters.

In the author’s opinion, these formulas have little meaning in the environmental modelling context. The balance between under-fitting and over-fitting can be formalized using theory based on singular value decomposition that is presented in the next chapter. That theory employs **k**-level parameters rather than **p**-level parameters; hence it avoids problems associated with discerning the exact meaning of the latter. It is also directly applicable to the calibration context.

Use of the above statistics in ranking alternative model conceptualizations is also questionable.

As is addressed in chapter 15 of this book, many model defects are invisible to the calibration process. Yet their effect on some model predictions may be profound; the deleterious effects of model defects on these predictions may be amplified through the calibration process. At the same time, as far as other model predictions are concerned, model defects can be happily “calibrated out”. Model parameterization and calibration strategies must therefore be prediction-specific rather than model-specific. The matter is much more nuanced than the above statistics imply.

8. Mathematical Regularization

8.1 General

8.1.1 Formulation of Equations

As in the previous chapter, we start by assuming linear model behaviour. This allows us to develop theory and explore concepts. Adaptation of the theory to nonlinear models is then discussed.

Once again, we start with the equation that describes a linear model \mathbf{Z} running under calibration conditions acting on m parameters encapsulated in the vector \mathbf{k} . Measurements of system state \mathbf{h} (n of these are available so that \mathbf{h} has n elements) are accompanied by noise $\boldsymbol{\epsilon}$ so that

$$\mathbf{h} = \mathbf{Z}\mathbf{k} + \boldsymbol{\epsilon} \quad (8.1.1)$$

In this chapter we assume that \mathbf{k} represents model parameters on the same scale as that to which expert knowledge applies. Hence we assume that the prior probability distribution of \mathbf{k} , including its covariance matrix $\mathbf{C}(\mathbf{k})$, are a direct outcome of expert knowledge as it pertains to the system that is being simulated. We also assume that \mathbf{k} encapsulates system property heterogeneity at a scale that is pertinent to predictions that the model must make. For now, we further assume that the model has no defects; hence any misfit between model outcomes and field measurements is an outcome of measurement noise alone.

Of course, in real-world modelling, none of this is true. Some upscaling is always required in definition of \mathbf{k} ; all models have defects incurred by spatial and temporal discretization, approximations in representation of initial and boundary conditions, and of course through imperfections in their simulation algorithms. Model defects are discussed in a later chapter.

8.1.2 Measurement Noise

For convenience in developing the equations that follow, we assume that an appropriate weight matrix \mathbf{Q} has been built into equation (8.1.1). Ideally this matrix is chosen in such a way as to render measurement noise, as seen by the inversion process, homoscedastic and independent so that its covariance matrix is proportional to the identity matrix \mathbf{I} . An appropriate \mathbf{Q} matrix is

$$\mathbf{Q} = \sigma^2_{\epsilon} \mathbf{C}^{-1}(\boldsymbol{\epsilon}) \quad (8.1.2)$$

Note that we specify the scalar multiplier in the above equation as σ^2_{ϵ} rather than σ^2_r because of the slightly different treatment that this multiplier receives in equations that follow. If all terms of (8.1.1) are multiplied by $\mathbf{Q}^{1/2}$, equation (8.1.2) becomes

$$\mathbf{Q}^{1/2}\mathbf{h} = \mathbf{Q}^{1/2}\mathbf{Z}\mathbf{k} + \mathbf{Q}^{1/2}\boldsymbol{\epsilon} \quad (8.1.3)$$

If we define

$$\mathbf{g} = \mathbf{Q}^{1/2}\mathbf{h} \quad (8.1.4)$$

$$\mathbf{Y} = \mathbf{Q}^{1/2}\mathbf{Z} \quad (8.1.5)$$

$$\boldsymbol{\xi} = \mathbf{Q}^{1/2}\boldsymbol{\epsilon} \quad (8.1.6)$$

equation (8.1.3) becomes

$$\mathbf{g} = \mathbf{Y}\mathbf{k} + \boldsymbol{\xi} \quad (8.1.7)$$

From (3.9.2) and (8.1.2) the covariance matrix of measurement noise ξ as represented in this equation is

$$C(\xi) = \sigma_e^2 C^{-1/2}(\boldsymbol{\varepsilon}) C(\boldsymbol{\varepsilon}) C^{-1/2}(\boldsymbol{\varepsilon}) = \sigma_e^2 \mathbf{I} \quad (8.1.8)$$

The benefits of choosing a weight matrix \mathbf{Q} defined by (8.1.2) are discussed in the context of well-posed inverse problems in the previous chapter. If (8.1.7) constitutes a well-posed inverse problem, an objective function could be formulated whose minimization constitutes solution of that problem. The objective function introduced in the previous chapter is

$$\Phi = (\mathbf{g} - \mathbf{Yk})^t (\mathbf{g} - \mathbf{Yk}) \quad (8.1.9)$$

This is the sum of squared residuals. With substitution of (8.1.4) and (8.1.5) into (8.1.9) this becomes

$$\Phi = (\mathbf{h} - \mathbf{Zk})^t \mathbf{Q} (\mathbf{h} - \mathbf{Zk}) \quad (8.1.10a)$$

That is

$$\Phi = \mathbf{r}^t \mathbf{Q} \mathbf{r} = \sigma_e^2 \mathbf{r}^t C^{-1}(\boldsymbol{\varepsilon}) \mathbf{r} \quad (8.1.10b)$$

where residuals \mathbf{r} are defined as

$$\mathbf{r} = \mathbf{h} - \mathbf{Zk} \quad (8.1.11)$$

Where a problem is ill-posed, the objective function defined above is not actually minimized; see later. Where enough parameters are featured in the inverse problem, it may actually be possible to reduce the objective function to nearly zero, in spite of the presence of measurement noise in \mathbf{h} . For reasons which will become obvious shortly, a parameter set which achieves such a high level of model-to-measurement fit cannot be construed as being an optimal solution to the inverse problem of highly parameterized model calibration. Instead, parameters must be adjusted (in a controlled way) until Φ is reduced to a value that is commensurate with measurement noise.

We can say something about what this Φ value should be. If σ_e^2 is unity, $\boldsymbol{\varepsilon}$ is multinormal, and we replace \mathbf{r} by $\boldsymbol{\varepsilon}$ in equation (8.1.10b), the resulting approximation to Φ possesses a chi-squared distribution with n degrees of freedom, where n is the number of observations comprising the calibration dataset \mathbf{h} . The expected value of Φ is therefore n ; see section 3.2.3.

To simplify the following mathematics, and to maintain conformity with equations that are presented in previous chapters, we will assume that equation (8.1.1) has already been transformed through pre-multiplication of all terms by a user-specified \mathbf{Q} matrix. Hence the \mathbf{Q} matrix will not appear in equations that follow. Nevertheless its presence will be implied in this altered definition of \mathbf{Z} ; this allows us to use \mathbf{Z} (and not \mathbf{Y}) to characterize the action of the model. As stated above, \mathbf{Q} should be calculated using (8.1.2).

8.1.3 Parameter and Predictive Error Variance

The discussion in this section is similar to that of section 7.5.2. However it is repeated here for the sake of completeness of the present chapter. The discussion is also slightly more general than in section 7.5.2, thereby making it more relevant to all regularization methodologies, including mathematical regularization.

Suppose that a vector \mathbf{k} has been found that can be said to “calibrate” the model \mathbf{Z} . Suppose further that the manner of its calculation from the calibration dataset \mathbf{h} can be encapsulated in a matrix \mathbf{G} . Different means of solving an ill-posed inverse problem result in calculation of different \mathbf{G} matrices, this depending on the adopted regularization device. As has been

extensively discussed, there can be no unique solution to an ill-posed inverse problem without regularization of one kind or another. Thus

$$\underline{\mathbf{k}} = \mathbf{G}\mathbf{h} \quad (8.1.12)$$

Substitution of (8.1.1) into (8.1.12) gives

$$\underline{\mathbf{k}} = \mathbf{G}\mathbf{Z}\mathbf{k} + \mathbf{G}\boldsymbol{\varepsilon} \quad (8.1.13a)$$

or

$$\underline{\mathbf{k}} = \mathbf{R}\mathbf{k} + \mathbf{G}\boldsymbol{\varepsilon} \quad (8.1.13b)$$

where

$$\mathbf{R} = \mathbf{G}\mathbf{Z} \quad (8.1.14)$$

\mathbf{R} , of course, is the resolution matrix. If there is no noise associated with the calibration dataset, this matrix describes the relationship between estimated parameters $\underline{\mathbf{k}}$ and real-world parameters \mathbf{k} . Obviously \mathbf{R} is a square matrix. If the inverse problem is well-posed, \mathbf{R} is the identity matrix \mathbf{I} . If the inverse problem is ill-posed then it can be shown that \mathbf{R} is rank-deficient and hence non-invertible. As has also been discussed, each row of the resolution matrix provides factors through which an individual element of $\underline{\mathbf{k}}$ can be calculated from all elements of \mathbf{k} . If \mathbf{R} is not \mathbf{I} (as it cannot be if an inverse problem is ill-posed) then some averaging of the elements of \mathbf{k} thus takes place in calculation of the elements of $\underline{\mathbf{k}}$. $\underline{\mathbf{k}}$ cannot possibly be \mathbf{k} under these circumstances. The best that can be hoped for is that \mathbf{R} somehow approaches \mathbf{I} so that this “blurring” of estimated parameters is minimized. However the more ill-posed is an inverse problem, the less is this possible. It follows that even where measurement noise is zero, $\underline{\mathbf{k}}$ is always in error, with the magnitude of this error increasing with ill-posedness of the inverse problem. This is the cost of seeking parameter uniqueness (and hence the cost of calibrating a model).

Other interesting properties of \mathbf{G} and \mathbf{R} emerge if $\boldsymbol{\varepsilon}$ is zero. These define asymptotic properties of these matrices in more realistic cases where $\boldsymbol{\varepsilon}$ is non-zero. Where $\boldsymbol{\varepsilon}$ is zero a modeller is entitled to seek a perfect fit between model outputs and the calibration dataset \mathbf{h} . Thus

$$\mathbf{h} = \mathbf{Z}\underline{\mathbf{k}} \quad (8.1.15)$$

From (8.1.1) with no noise, (8.1.12) and (8.1.15) we obtain

$$\mathbf{h} = \mathbf{Z}\underline{\mathbf{k}} = \mathbf{Z}\mathbf{k} = \mathbf{Z}\mathbf{G}\mathbf{h} = \mathbf{Z}\mathbf{G}\mathbf{Z}\mathbf{k} \quad (8.1.16)$$

and hence

$$\mathbf{Z} = \mathbf{Z}\mathbf{G}\mathbf{Z} \quad (8.1.17)$$

which, through comparison with (2.9.5), establishes that \mathbf{G} is a generalized inverse of \mathbf{Z} . Premultiplication of both sides of (8.1.17) by \mathbf{G} establishes that the resolution matrix \mathbf{R} is idempotent and hence is a projection operator. If it is symmetric it becomes an orthogonal projection operator. As will be discussed, this is the case where parameter uniqueness is achieved through use of singular value decomposition as a regularization device.

From (8.1.13), the error in the estimated parameter set $\underline{\mathbf{k}}$ in the presence of measurement noise can be calculated as

$$\underline{\mathbf{k}} - \mathbf{k} = -(\mathbf{I} - \mathbf{R})\mathbf{k} + \mathbf{G}\boldsymbol{\varepsilon} \quad (8.1.18)$$

Characterization of parameter error thus requires two terms. One of these terms (the second term on the right of equation 8.1.18) arises from the presence of noise in the measurement

dataset. The other term would arise even if there were no noise in the calibration dataset. This contribution to estimated parameter error arises to the extent that \mathbf{R} is different from \mathbf{I} . It is thus an unavoidable outcome of seeking a unique solution to an ill-posed inverse problem.

Whenever a unique solution is sought to an ill-posed inverse problem, a resolution matrix is implied or, in cases where regularization is mathematical, can be explicitly calculated. As already stated, this resolution matrix \mathbf{R} cannot be \mathbf{I} if an inverse problem is ill-posed. The resolution matrix that emerges from manual regularization (in which an ill-posed inverse problem is turned into a well-posed inverse problem by fixing some parameters and combining others) is discussed in the previous chapter. We note that evaluation of the resolution matrix in contexts of manual regularization is not generally possible as it requires knowledge of the \mathbf{Z} matrix. The intention of manual regularization is to avoid filling this matrix.

An advantage of formalized mathematical regularization over manual regularization is that the \mathbf{G} matrix as it pertains to model parameters \mathbf{k} , and hence the \mathbf{R} matrix, becomes available as the inverse problem is solved. Hence the cost-of-uniqueness term of equation (8.1.18) can be explored if desired. However, the error in estimated parameters $\underline{\mathbf{k}}$ cannot be known. This is because \mathbf{k} , the vector of real-world parameters whose values are being estimated, is not known. \mathbf{k} is featured on the right side of equation (8.1.18); so too is the unknown vector $\mathbf{\epsilon}$ of measurement noise. Hence (8.1.18) cannot be used to calculate $\underline{\mathbf{k}} - \mathbf{k}$.

Nevertheless, while parameter error itself cannot be determined, much can be said about the *propensity* for parameter error. This can be expressed by the covariance matrix of post-calibration parameter error $C(\underline{\mathbf{k}} - \mathbf{k})$. Using (3.9.2), $C(\underline{\mathbf{k}} - \mathbf{k})$ can be derived from (8.1.18) as

$$C(\underline{\mathbf{k}} - \mathbf{k}) = (\mathbf{I} - \mathbf{R})C(\mathbf{k})(\mathbf{I} - \mathbf{R})^t + \mathbf{G}C(\mathbf{\epsilon})\mathbf{G}^t \quad (8.1.19)$$

Two covariance matrices are featured on the right side of (8.1.19); these are $C(\mathbf{k})$ and $C(\mathbf{\epsilon})$. It follows that if we seek to quantify the post-calibration potential for parameter error (or the post-calibration potential for predictive error for which an equation will be derived shortly), then we need to know more than just the statistics of measurement noise; recall that this is all that is required for quantification of parameter and predictive error variance in the well-posed context. In the ill-posed context we also need to encapsulate expert knowledge in a $C(\mathbf{k})$ matrix which quantifies the innate variability of parameters \mathbf{k} . Conceptually this is available from the prior parameter probability distribution, the same probability distribution that is required by Bayes equation. As well as denoting innate parameter variability, off diagonal terms of $C(\mathbf{k})$ denote any correlation (spatial or otherwise) that system properties are likely to possess.

Let s be a prediction of management interest. Let the sensitivity of this prediction to parameters \mathbf{k} be expressed by the vector \mathbf{y} . The true value of this prediction is calculated as

$$s = \mathbf{y}^t \mathbf{k} \quad (8.1.20)$$

Through propagation of covariance (i.e. equation 3.9.2), the variance of this prediction is then given by

$$\sigma^2_s = \mathbf{y}^t C(\mathbf{k}) \mathbf{y} \quad (8.1.21)$$

σ^2_s can be viewed in a number of different ways. It describes the prior uncertainty of the prediction s . (Recall that variance is the square of standard deviation.) That is, it describes the uncertainty of this prediction if made on the basis of expert knowledge alone. σ^2_s can also be viewed as the pre-calibration error variance of the prediction. If a single value were to be sought for this prediction based on expert knowledge alone, then this value would be calculated using prior expected values (i.e. prior means) of all of the parameters encapsulated in \mathbf{k} . According to protocols that are used in this book to linearize a model, these are all zero; see section 6.5.

Hence, without any loss of generality, the prior mean of the prediction s is also zero. σ_s^2 quantifies the potential for error in making this prediction.

The prediction made by the calibrated model is calculated as

$$\underline{s} = \mathbf{y}^t \mathbf{k} \quad (8.1.22)$$

Predictive error is calculated by subtracting (8.1.21) from (8.1.22).

$$\underline{s} - s = \mathbf{y}^t (\underline{\mathbf{k}} - \mathbf{k}) \quad (8.1.23)$$

Using the propagation of variance formula once again, the variance of post-calibration predictive error becomes

$$\sigma_{\underline{s}-s}^2 = \mathbf{y}^t \mathbf{C} (\underline{\mathbf{k}} - \mathbf{k}) \mathbf{y} \quad (8.1.24)$$

which, from (8.1.19), is

$$\sigma_{\underline{s}-s}^2 = \mathbf{y}^t (\mathbf{I} - \mathbf{R}) \mathbf{C} (\mathbf{k}) (\mathbf{I} - \mathbf{R})^t \mathbf{y} + \mathbf{y}^t \mathbf{G} \mathbf{C} (\mathbf{\epsilon}) \mathbf{G}^t \mathbf{y} \quad (8.1.25)$$

Equation (8.1.25) quantifies the potential for error in a prediction \underline{s} when made using a calibrated model.

For a nonlinear model, predictive error variance as calculated using equation (8.1.25) is only approximately correct.

8.1.4 Overview of Mathematical Regularization

Regularization, in its broadest sense, refers to the process through which a unique solution is sought to an ill-posed inverse problem. By definition, many solutions exist to an ill-posed inverse problem. The metric through which one of these can be considered as optimal is provided in section 6.4. The optimal solution to the inverse problem of model calibration is that which endows a model with the ability to make predictions of future system behaviour which are of minimized error variance.

It may be argued that the minimum error variance metric should actually be applied to parameters rather than to predictions, for if estimated parameters are of minimum error variance then minimization of predictive error variance follows automatically. Unfortunately this is not the case when model imperfections are taken into account, for under these circumstances the link between parameter and predictive error variance is broken. Optimality of model calibration may therefore be prediction-specific. This is a subject for a later chapter. In the present chapter, however, it is assumed that models have no defects.

As has already been shown, uniqueness comes at a cost; that cost is a potential for error in estimated parameters. Pivotal to the evaluation of this cost is the resolution matrix. Where regularization is implemented under mathematical control this can be calculated. So too then can the propensity for parameter and predictive error. Post-calibration predictive error analysis thus becomes possible.

We note, however, that this is rarely done. Instead, Bayesian methods that are discussed in subsequent chapters of this book are used to explore parameter and predictive uncertainty. Nevertheless, because the means to explore the potential for parameter and predictive error that is incurred by the quest for uniqueness are provided by the same mathematics that makes this quest possible, mathematical regularization possesses integrity that manual regularization does not.

This is not the only advantage that mathematical regularization holds over manual regularization. Where a model has spatially distributed parameters, all of these parameters can

be retained in the calibration process as elements of the vector \mathbf{k} . This endows the inversion process with maximum flexibility in introducing system property heterogeneity to locations within a model domain at which the calibration dataset suggests that such heterogeneity should exist. Of course the process of heterogeneity emplacement must take place under the control of a regularization strategy that precludes the introduction of spurious heterogeneity, and that provides credibility to patterns that are adopted by emergent heterogeneity. Contrast this with manual regularization wherein heterogeneity must express itself in pre-defined spatial patterns, such as a small number of zones of piecewise constancy whose geometries may or may not be appropriate, and whose number may or may not be adequate to reflect the information content of a calibration dataset.

Whether or not parameterization is spatial, the practicalities of implementing mathematical regularization make it a far easier option than manual regularization. In principle, decisions of which parameters to estimate, which parameters to lump together, and which parameters to fix, never need to be made under a mathematical regularization regime. Instead an approach of “if in doubt, declare it as adjustable” can prevail. If it transpires that the information content of a calibration dataset is insufficient to allow estimation of a particular parameter or group of parameters, then the post-calibration value of the parameter, or group of parameters, will be unchanged from their pre-calibration values. At the same time, numerical instability incurred through attempts to invert non-invertible matrices is precluded as mathematical regularization alters these matrices to enhance their invertibility.

This chapter introduces two broad approaches to mathematical regularization. These are broadly known as “subspace methods” (of which singular value decomposition is the flagship) and Tikhonov. Each has advantages and disadvantages. In practice, it is best to use both of them together.

Subspace methods reduce the inverse problem to that of estimating only as many parameters as are actually estimable on the basis of a calibration dataset. In doing this, they resemble manual regularization. However a significant point of difference with manual regularization is that the inversion process itself determines the number of parameters that are actually estimable. Furthermore, these estimable “parameters” are actually optimally-defined linear combinations of user-specified model parameters. If model parameters are subjected to an appropriate transformation prior to undergoing inversion (see below), not only can a minimum error variance status be guaranteed for estimated parameters, but the process is also guaranteed to be numerically stable as no matrices must be inverted that cannot, in fact, be easily inverted.

In contrast to subspace methods, Tikhonov regularization addresses information deficits in a calibration dataset relative to parameters that require estimation by supplementing the calibration dataset with expert knowledge pertaining to these parameters. A challenge in formulating the inverse problem, particularly in contexts where parameters have spatial connotations and may need to express complex patterns of spatial heterogeneity, is to find the best way to express such expert knowledge in a mathematical way that the inversion process understands. However while challenging, formulation of an appropriate Tikhonov regularization scheme can also be very flexible. Expressions of expert knowledge do not need to be linear. Hence complex, nonlinear inter-parameter relationships can spawn the appearance of sharp edges, anisotropy, proximity of heterogeneity to certain points or lines, and much else besides, where this is hydrogeologically appropriate. While doing this, PEST’s implementation of Tikhonov regularization is able to ensure correct relativity of weighting between expert knowledge on the one hand and field measurements of system behaviour on the other hand. This forestalls numerical difficulties in solution of the inverse problem at the same time as it

ensures that the two sources of information act in concert rather than in opposition to each other.

8.2 Singular Value Decomposition

8.2.1 Estimation of Parameters

Singular value decomposition (SVD) is discussed in section 2.13. The reader is referred to that section for details of the orthonormal matrices that emerge from the singular value decomposition process. It is now applied to estimation of parameters in an inversion context described by equation (8.1.1).

Using singular value decomposition, the \mathbf{Z} matrix of equation (8.1.1) can be decomposed as

$$\mathbf{Z} = \mathbf{USV}^t \quad (8.2.1)$$

Recall from section 2.13 that the orthogonal unit vectors which form the columns of \mathbf{U} span n -dimensional space, where n is the number of rows of \mathbf{Z} . This space includes the output space of \mathbf{Z} (i.e. the output space of the model). The unit vectors which comprise the columns of \mathbf{V} span parameter space. The \mathbf{S} matrix has elements only along its diagonal. These are all non-negative. Their values become smaller as the diagonal is descended. Eventually they either become zero, or the diagonal runs out of elements. The latter occurs if the number of parameters m used by a model exceeds the number of observations n comprising the calibration dataset.

With appropriate partitioning of \mathbf{U} , \mathbf{S} and \mathbf{V} as described in section 2.13, equation (8.2.1) can be written as

$$\mathbf{Z} = \mathbf{U}_1 \mathbf{S}_1 \mathbf{V}_1^t + \mathbf{U}_2 \mathbf{S}_2 \mathbf{V}_2^t \quad (8.2.2)$$

In specifying the partitioning that is embodied in equation (8.2.2), the \mathbf{S}_2 matrix is designed to be either completely zero-valued, or to contain only small singular values along its diagonal. Meanwhile the unit vectors which comprise the columns of \mathbf{V}_1 and \mathbf{V}_2 span two orthogonal subspaces which collectively comprise the whole of parameter space. If (8.2.2) is substituted into (8.1.1) we obtain

$$\mathbf{h} = (\mathbf{U}_1 \mathbf{S}_1 \mathbf{V}_1^t + \mathbf{U}_2 \mathbf{S}_2 \mathbf{V}_2^t) \mathbf{k} + \boldsymbol{\epsilon} \quad (8.2.3)$$

The solution \mathbf{k} of the inverse problem is calculated as (see section 2.13)

$$\mathbf{k} = \mathbf{V}_1 \mathbf{S}_1^{-1} \mathbf{U}_1^t \mathbf{h} \quad (8.2.4)$$

Calculation of \mathbf{k} can proceed without any conceptual or numerical difficulties if partitioning of the \mathbf{S} matrix (and with it the implied partitioning of the \mathbf{U} and \mathbf{V} matrices) is such that \mathbf{S}_1 is square and contains only non-zero diagonal elements. If \mathbf{S}_2 is entirely zero-valued and \mathbf{S}_1 contains no zero-valued singular values, the columns of \mathbf{V}_2 span the null space of the matrix \mathbf{Z} . We will refer to its orthogonal complement (which is spanned by the columns of \mathbf{V}_1) as the “solution space” of \mathbf{Z} . Because \mathbf{V}_1 is the first-occurring matrix on the right side of (8.2.4), the solution \mathbf{k} to the inverse problem is a linear combination of the vectors which define its columns.

Any vector \mathbf{j} which is a linear combination of the columns of \mathbf{V}_2 can be written as

$$\mathbf{j} = \mathbf{V}_2 \mathbf{w} \quad (8.2.5)$$

where \mathbf{w} is an arbitrary vector with as many elements as there are columns in \mathbf{V}_2 . In this equation the multiplier of column i of \mathbf{V}_2 is w_i , the i 'th element of \mathbf{w} . From (8.2.2)

$$\mathbf{Zj} = \mathbf{U}_1 \mathbf{S}_1 \mathbf{V}_1^t \mathbf{V}_2 \mathbf{w} + \mathbf{U}_2 \mathbf{S}_2 \mathbf{V}_2^t \mathbf{V}_2 \mathbf{w} \quad (8.2.6)$$

The first term on the right of equation (8.2.6) is zero because of the orthogonality of \mathbf{V}_1 and \mathbf{V}_2 . The second term is zero if all elements of \mathbf{S}_2 are zero. Hence $\mathbf{Z}\mathbf{j}$ is zero, so that \mathbf{j} is in the null space of \mathbf{Z} .

Normally the partitioning of \mathbf{S} into \mathbf{S}_1 and \mathbf{S}_2 , and with it the partitioning of \mathbf{V} into \mathbf{V}_1 and \mathbf{V}_2 , for the purpose of estimating \mathbf{k} using equation (8.2.4) is not strictly according to the onset of zero diagonal elements of \mathbf{S}_2 . Normally the leading diagonal elements of \mathbf{S}_2 are small but nonzero; the reasons for this will become apparent shortly. Hence \mathbf{j} defined according to (8.2.5) when substituted into (8.2.6) yields a non-zero outcome. Nevertheless, this outcome is small because the first term on the right of (8.2.6) remains zero while the second term of (8.2.6) is small because the diagonal elements of \mathbf{S}_2 are small.

In the discussion that follows, we will continue to refer to the subdivision of parameter space into \mathbf{V}_1 and \mathbf{V}_2 in the way described above as subdivision of this space into the “solution space” and the “null space”. This is in spite of the fact that it is not completely correct to equate the null space with the \mathbf{V}_2 space unless all diagonal elements of \mathbf{S}_2 are strictly zero. Nevertheless, adoption of this protocol makes the following discussion easier.

8.2.2 An Example

Figure 8.1 shows a simple, one-dimensional groundwater model. Steady-state flow of water is from right to left at a rate of 1 flow unit. The head at the left boundary is fixed at zero. The calibration dataset is comprised of the head measured in a single well at the extreme right end of the model domain. Let this head be h . Two model parameters determine the head in this well. These are the resistances within each of the two zones through which groundwater flows. Hence

$$h = (r_1 + r_2) \quad (8.2.7a)$$

or, to write this in matrix form,

$$[h] = [1 \ 1] \begin{bmatrix} r_1 \\ r_2 \end{bmatrix} \quad (8.2.7b)$$

In equation (8.2.7) the \mathbf{Z} matrix of equation (8.1.1) is $[1 \ 1]$, the \mathbf{k} vector is $\begin{bmatrix} r_1 \\ r_2 \end{bmatrix}$, while the \mathbf{h} vector is $[h]$.

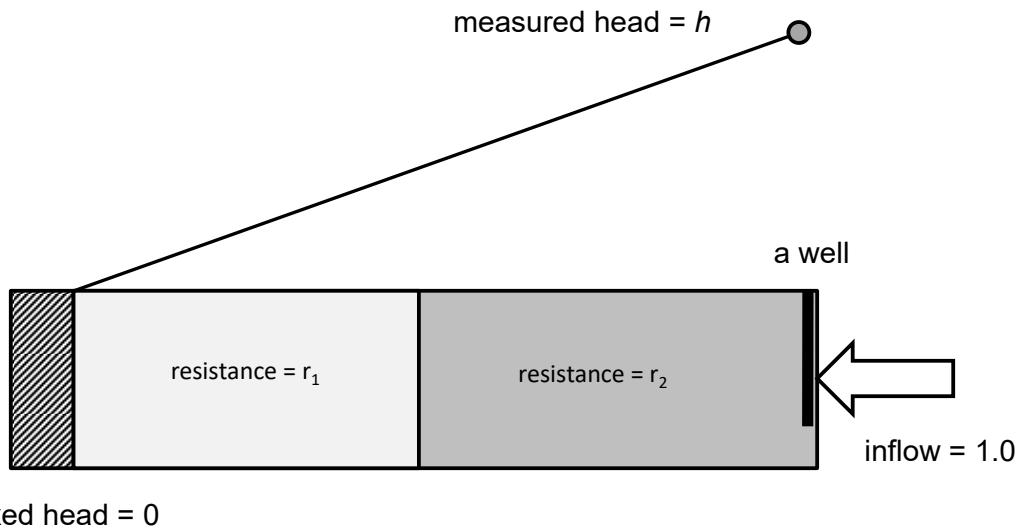


Figure 8.1. A one-dimensional groundwater model domain in which water flows toward a boundary in which the head is fixed at zero.

It can be verified with ease that the vector $\begin{bmatrix} \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} \end{bmatrix}$ is a unit vector. It can be verified with equal ease that

$$\mathbf{Z} \begin{bmatrix} \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} \end{bmatrix} = \mathbf{0} \quad (8.2.8)$$

Hence this vector is in the null space of \mathbf{Z} . It thus comprises the \mathbf{V}_2 matrix of equation (8.2.2). Obviously, in the present example, this matrix has just one column; hence the null space of \mathbf{Z}

is one-dimensional. Meanwhile the unit vector $\begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix}$ comprises the \mathbf{V}_1 matrix; it can also be

readily verified that this is orthogonal to the vector comprising the single column of \mathbf{V}_2 . The single column of the \mathbf{V}_1 matrix comprises the one-dimensional solution space of the \mathbf{Z} matrix. Hence any solution to the ill-posed inverse problem of estimating r_1 and r_2 must be a multiple of this vector. r_1 and r_2 must therefore be estimated as equal.

Intuition readily informs us that the single head which comprises the observation dataset \mathbf{h} allows estimation of only the sum of the resistances r_1 and r_2 that occupy the model domain. Because the inverse problem is ill-posed, there are an infinite number of possible values which r_1 and r_2 can take which, when added together, form this single estimable sum. Uniqueness is achieved by assuming that r_1 and r_2 are equal, for unless there is information to the contrary, this constitutes the “safest” solution to the inverse problem, and hence the solution of minimum error variance.

Differences between r_1 and r_2 are irresolvable on the basis of the calibration dataset \mathbf{h} because they lie in the null space; see equation (8.2.8). For safety, this difference is assigned a value of zero in achieving the minimum error variance solution to the inverse problem that is posed by figure 8.1; as discussed in section (2.13) this minimum error variance solution must have a null space projection of zero. This does not mean, of course, that r_1 and r_2 are necessarily equal. It

only means that the probability of r_1 being greater than r_2 is the same as the probability of r_2 being greater than r_1 . The middle path is chosen in seeking a minimum error variance solution to the inverse problem. Meanwhile the predictive repercussions of r_1 and r_2 being different from each other can be explored later when post-calibration error analysis is undertaken. When undertaking this analysis, r_1 and r_2 can indeed be endowed with different values; however their sum must always be unity so that calibration constraints are respected.

8.2.3 Relationship between Estimated and Real Parameters

With substitution of (8.2.3) into (8.2.4) we obtain

$$\underline{\mathbf{k}} = \mathbf{V}_1 \mathbf{S}^{-1} \mathbf{U}_1 \mathbf{U}_1^t \mathbf{V}_1^t \mathbf{k} + \mathbf{V}_1 \mathbf{S}^{-1} \mathbf{U}_1 \mathbf{U}_2 \mathbf{S}_2 \mathbf{V}_2^t \mathbf{k} + \mathbf{V}_1 \mathbf{S}^{-1} \mathbf{U}_1^t \mathbf{\epsilon} \quad (8.2.9a)$$

Because $\mathbf{U}_1^t \mathbf{U}_1$ is \mathbf{I} and $\mathbf{U}_1^t \mathbf{U}_2$ is $\mathbf{0}$ this becomes

$$\underline{\mathbf{k}} = \mathbf{V}_1 \mathbf{V}_1^t \mathbf{k} + \mathbf{V}_1 \mathbf{S}^{-1} \mathbf{U}_1^t \mathbf{\epsilon} \quad (8.2.9b)$$

Equation (8.2.9b) is in the form of equation (8.1.13b). Hence, for singular value decomposition as a regularization device

$$\mathbf{R} = \mathbf{V}_1 \mathbf{V}_1^t \quad (8.2.10)$$

$$\mathbf{G} = \mathbf{V}_1 \mathbf{S}^{-1} \mathbf{U}_1^t \quad (8.2.11)$$

As can be readily verified from its symmetry and idempotency, \mathbf{R} is an orthogonal projection operator. The relationship between the estimated parameter set $\underline{\mathbf{k}}$ and the true parameter set \mathbf{k} is thus that $\underline{\mathbf{k}}$ is the projection of \mathbf{k} into the solution space of the model matrix \mathbf{Z} . This is the subspace of parameter space spanned by unit vectors comprising the columns of \mathbf{V}_1 . The relationship is illustrated in figure 8.2. This is a repeat of figure 2.3 with \mathbf{k} substituted for \mathbf{x} which features in the previous figure. For the three-dimensional parameter space schematized in figure 8.2, \mathbf{V} is given by

$$\mathbf{V} = [\mathbf{v}_1 \ \mathbf{v}_2 \ \mathbf{v}_3] \quad (8.2.12)$$

The solution space of \mathbf{Z} is spanned by \mathbf{v}_1 and \mathbf{v}_2 while the null space of \mathbf{Z} is spanned by \mathbf{v}_3 .

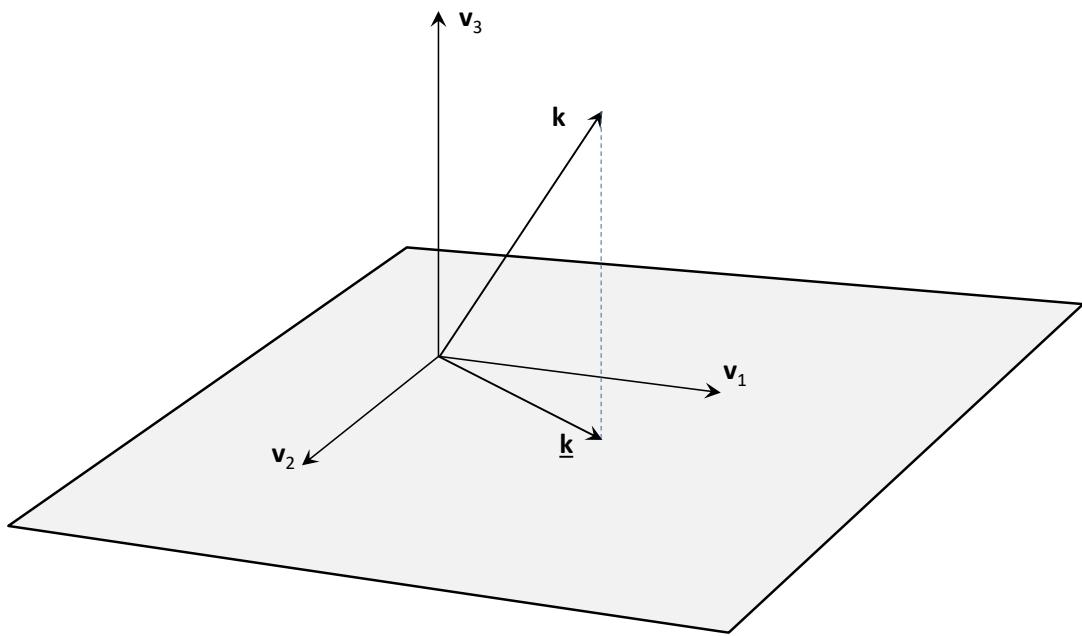


Figure 8.2. The relationship between estimated and true parameters is that of orthogonal projection onto a limited dimensional solution space.

From (8.2.9b) and (2.13.14) parameter error is calculated as

$$\underline{k} - k = -(\mathbf{I} - \mathbf{V}_1 \mathbf{V}_1^t)k + \mathbf{V}_1 \mathbf{S}^{-1} \mathbf{U}_1^t \boldsymbol{\epsilon} = -\mathbf{V}_2 \mathbf{V}_2^t k + \mathbf{V}_1 \mathbf{S}^{-1} \mathbf{U}_1^t \boldsymbol{\epsilon} \quad (8.2.13)$$

It follows that the error in the estimated parameter set \underline{k} is comprised of two orthogonal terms. The first is the projection of the true parameter set onto the null space of \mathbf{Z} ; this is the first term on the right of equation (8.2.13). The second contribution to parameter error lies in the solution space of \mathbf{Z} . This is the error that is incurred in estimating the projection of the true parameter vector k onto this space because of noise $\boldsymbol{\epsilon}$ in the calibration dataset \mathbf{h} . The situation is depicted in figure 8.3.

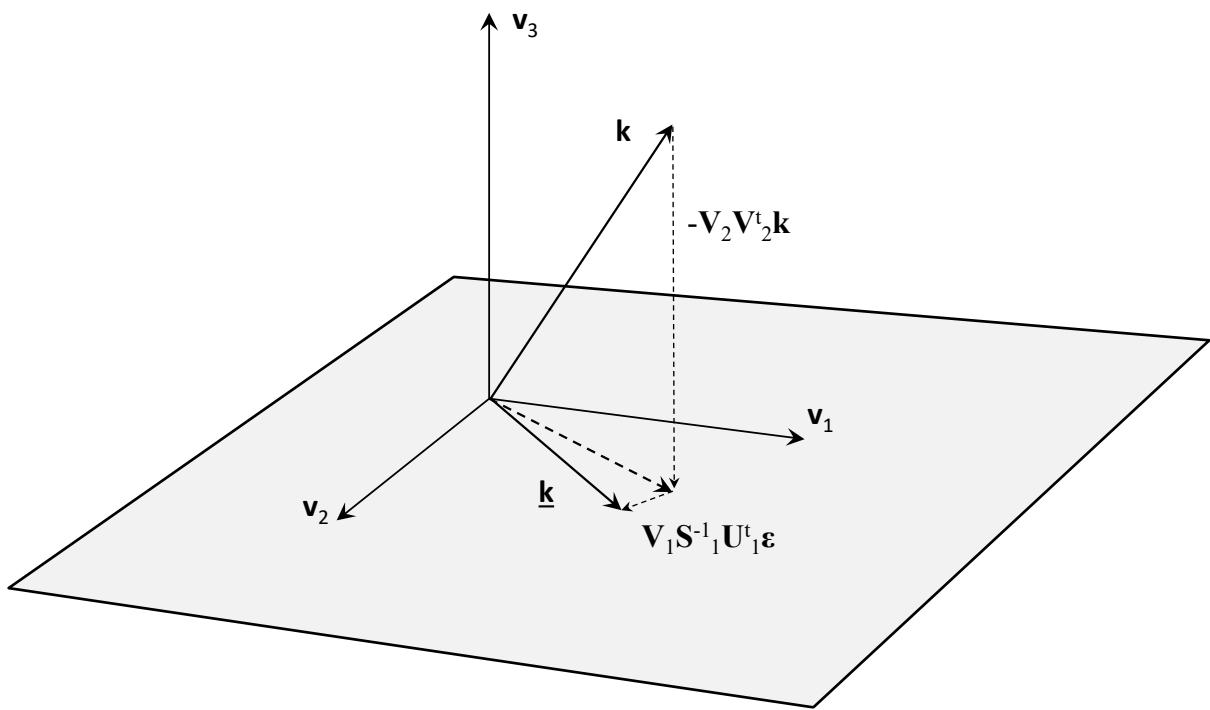


Figure 8.3. The two components of error in the estimated parameter set \underline{k} .

The decision to exclude any null space components of \mathbf{k} from $\underline{\mathbf{k}}$ is fundamental to seeking a solution to the inverse problem which has minimum potential for error, and hence is of minimum error variance. We do not know \mathbf{k} . We cannot know its null space projection because \mathbf{h} does not provide this information; this is because model outputs can fit the calibration dataset \mathbf{h} without parameters venturing into the null space at all. Meanwhile our estimate of the solution space component of \mathbf{k} from \mathbf{h} is somewhat in error because of the presence of measurement noise $\boldsymbol{\epsilon}$ in the calibration dataset. We refer to the estimated solution space component of \mathbf{k} as $\underline{\mathbf{k}}$; we deem our model to be calibrated once we have calculated this. At the same time we exclude from $\underline{\mathbf{k}}$ anything that has a non-zero projection onto the calibration null space, for it is just as likely that any null space component that we may include in $\underline{\mathbf{k}}$ would take us further from \mathbf{k} rather than closer to it. So we stay on the solution space hyperplane depicted in figure 8.3. This is the safest option – the minimum error variance option. It does not give us “the truth” as far as estimation of \mathbf{k} is concerned. However it gives us maximum protection from inadvertent falsehood.

8.2.4 Kahunen-Loève Transformation

8.2.4.1 The Need for Kahunen-Loève Transformation

Where there is no noise in a calibration dataset (and hence $\boldsymbol{\epsilon}$ of equation (8.1.1) is $\mathbf{0}$), partitioning of \mathbf{S} into \mathbf{S}_1 and \mathbf{S}_2 can take place exactly at the point where \mathbf{S}_2 is zero. (The reason why this should not happen if $\boldsymbol{\epsilon}$ is non-zero will be provided shortly.) In that case, as discussed in section 2.13, equation (8.2.4) provides the minimum norm (i.e. minimum magnitude) solution for $\underline{\mathbf{k}}$; calculation of $\underline{\mathbf{k}}$ from \mathbf{h} takes place using the Moore-Penrose pseudoinverse of \mathbf{Z} .

If the prior probability distribution of \mathbf{k} is multiGaussian with mean $\mathbf{0}$ (recall that a mean of $\mathbf{0}$ is a design specification that we have adopted for model linearization, and have formulated the definition of \mathbf{k} accordingly), and if the elements of \mathbf{k} are statistically independent and have the

same standard deviations, this solution is also that of highest probability density function value according to the prior probability distribution of \mathbf{k} . Under these conditions the exponent of equation (3.6.1) coincides with the equation for the L_2 norm of vector \mathbf{k} ; hence the estimated parameter set $\underline{\mathbf{k}}$ deviates from the prior mean parameter set of $\mathbf{0}$ to the smallest extent required to satisfy the condition that model outputs fit field measurements. This is obviously a good outcome, for it guarantees a high level of credibility for the estimated parameter set. It does not guarantee that this parameter set is right; it does guarantee however that it is minimally wrong.

In most cases of real-world modelling interest, the prior probability distribution of \mathbf{k} is not such that

$$C(\mathbf{k}) = \sigma^2_{\mathbf{k}} \mathbf{I} \quad (8.2.14)$$

as is required for minimum norm to coincide with maximum pre-calibration probability. However a parameter set \mathbf{m} can be calculated from a model's parameter set \mathbf{k} that has these properties; see section 3.10.1. As usual, let $C(\mathbf{k})$ denote the prior covariance matrix of \mathbf{k} . Let the eigenvector matrix \mathbf{E} and the diagonal eigenvalue matrix \mathbf{F} be obtained from $C(\mathbf{k})$ through singular value decomposition as follows

$$C(\mathbf{k}) = \mathbf{E} \mathbf{F} \mathbf{E}^t \quad (8.2.15)$$

Then the parameter set \mathbf{m} defined as

$$\mathbf{m} = \mathbf{F}^{-1/2} \mathbf{E}^t \mathbf{k} \quad (8.2.16)$$

has the desired property that its covariance matrix is \mathbf{I} .

Minimization of the L_2 norm of \mathbf{m} requires minimization of $\mathbf{m}^t \mathbf{m}$. From (8.2.16)

$$\mathbf{m}^t \mathbf{m} = \mathbf{k}^t \mathbf{E} \mathbf{F}^{-1/2} \mathbf{F}^{-1/2} \mathbf{E}^t \mathbf{k} = \mathbf{k}^t C^{-1}(\mathbf{k}) \mathbf{k} \quad (8.2.17)$$

The Mahalanobis distance from \mathbf{k} to its prior mean is therefore minimized.

From (8.2.16), \mathbf{k} can be back-calculated from \mathbf{m} as

$$\mathbf{k} = \mathbf{E} \mathbf{F}^{1/2} \mathbf{m} \quad (8.2.18)$$

It is apparent that use of singular value decomposition to solve the inverse problem of model calibration requires that the prior probability distribution of model parameters be taken into account if the solution to that inverse problem is to be of minimum error variance; this is done by subjecting model parameters to Kahunen-Loëve transformation before actually estimating them. Accommodation of the prior parameter probability distribution of parameters thus becomes an indispensable feature of an appropriately-formulated inverse problem.

In later sections of this chapter, prior knowledge of parameters will be introduced to the parameter estimation process through Tikhonov regularization, as this provides more flexible ways to express expert knowledge; however singular value decomposition will still play a strategic role in solving the inverse problem because of its ability to calculate, and then invert, a matrix that may approach singularity.

8.2.4.2 Repercussions of not Transforming

Suppose that we solve an ill-posed inverse problem for $\underline{\mathbf{k}}$ and that we do not pre-transform \mathbf{k} to obtain an \mathbf{m} to which equation (8.2.14) applies. From (8.2.9b) without measurement noise, the estimated $\underline{\mathbf{k}}$ is related to the true \mathbf{k} through the equation

$$\underline{\mathbf{k}} = \mathbf{V}_1 \mathbf{V}_1^t \mathbf{k} \quad (8.2.19)$$

Transformation and back-transformation between \mathbf{m} -space and \mathbf{k} -space is available through

equations (8.2.16) and (8.2.18). Through use of these two equations the relationship between estimated \mathbf{m} (i.e. $\underline{\mathbf{m}}$) and real \mathbf{m} parameters when singular value decomposition is undertaken in \mathbf{k} -space is

$$\underline{\mathbf{m}} = \mathbf{F}^{-1/2} \mathbf{E}^t \mathbf{V}_1 \mathbf{V}_1^t \mathbf{E} \mathbf{F}^{1/2} \mathbf{m} = \mathbf{R}' \mathbf{m} \quad (8.2.20)$$

As can be readily verified, the resolution matrix \mathbf{R}' of (8.2.20) is idempotent. Hence it is a projection operator. However, in general, it is not symmetric. Therefore it is not an orthogonal projection operator. What this means is that $\underline{\mathbf{m}}$ may include some \mathbf{m} -space null space components. The presence of these null space components invalidates claims to minimum error variance status of $\underline{\mathbf{m}}$. Hence it may suffer history-match-induced bias. Some (perhaps all) management-pertinent predictions made by the model will inherit this bias. Watson et al (2013) discuss this matter further, and show an example of calibration bias incurred in this way for a simple model.

The fact that singular value decomposition conducted without Kahunen-Loëve transformation can lead to inverse problem solution bias through unwitting entrainment of null space components in the calibrated parameter field can be further demonstrated as follows. Using (8.2.18), equation (8.1.1), which poses the inverse problem which is the subject of this chapter, can be expressed in \mathbf{m} -space as

$$\mathbf{h} = \mathbf{Z} \mathbf{E} \mathbf{F}^{1/2} \mathbf{m} + \boldsymbol{\varepsilon} = \mathbf{W} \mathbf{m} + \boldsymbol{\varepsilon} \quad (8.2.21)$$

where \mathbf{W} , the model matrix in \mathbf{m} -space, is defined through the above equation as

$$\mathbf{W} = \mathbf{Z} \mathbf{E} \mathbf{F}^{1/2} \quad (8.2.22)$$

If (8.2.22) is post-multiplied by $\mathbf{F}^{-1/2}$ and then by \mathbf{E}^t it follows that

$$\mathbf{Z} = \mathbf{W} \mathbf{F}^{-1/2} \mathbf{E}^t \quad (8.2.23)$$

Suppose that a vector \mathbf{m} is in the null space of \mathbf{W} . Then

$$\mathbf{W} \mathbf{m} = \mathbf{0} \quad (8.2.24)$$

If, in this equation, (8.2.22) is substituted for \mathbf{W} and (8.2.16) is substituted for \mathbf{m} , it immediately follows that

$$\mathbf{Z} \mathbf{E} \mathbf{F}^{1/2} \mathbf{F}^{-1/2} \mathbf{E}^t \mathbf{k} = \mathbf{Z} \mathbf{k} = \mathbf{0} \quad (8.2.25)$$

Hence a vector in \mathbf{m} -space which is in the null space of \mathbf{W} transforms to a vector in \mathbf{k} -space which is in the null space of \mathbf{Z} . The reverse is also true. Suppose that \mathbf{k} is in the null space of \mathbf{Z} so that (8.2.25) holds. If (8.2.23) is substituted for \mathbf{Z} and (8.2.18) is substituted for \mathbf{k} in (8.2.25) then

$$\mathbf{W} \mathbf{F}^{-1/2} \mathbf{E}^t \mathbf{E} \mathbf{F}^{1/2} \mathbf{m} = \mathbf{W} \mathbf{m} = \mathbf{0} \quad (8.2.26)$$

It follows that the null spaces of \mathbf{Z} and \mathbf{W} are directly transferrable between \mathbf{k} -space and \mathbf{m} -space. However it was demonstrated above that the orthogonal complement to the \mathbf{Z} matrix null space in \mathbf{k} -space does not transform to the orthogonal complement of the \mathbf{W} matrix null space in \mathbf{m} -space. Conversely, if a vector in \mathbf{m} -space is orthogonal to the null space of \mathbf{W} , it does not follow that the \mathbf{k} -space equivalent of this vector will not contain a null space component of \mathbf{Z} .

8.2.4.3 An Alternative Transformation

It was shown above that use of equation (8.2.16) can yield a parameter set \mathbf{m} whose covariance matrix is the identity matrix. Minimum norm of \mathbf{m} is therefore equivalent to maximum

likelihood (and minimum error variance) of \mathbf{m} . An alternative transformation that can yield the same objectives as far as pursuit of calibration optimality is concerned is

$$\mathbf{j} = \mathbf{E}\mathbf{F}^{1/2}\mathbf{E}^t\mathbf{k} \quad (8.2.27a)$$

That is

$$\mathbf{j} = \mathbf{C}^{-1/2}(\mathbf{k})\mathbf{k} \quad (8.2.27b)$$

Back transformation yields the following equations

$$\mathbf{k} = \mathbf{E}\mathbf{F}^{1/2}\mathbf{E}^t\mathbf{j} = \mathbf{C}^{1/2}(\mathbf{k})\mathbf{j} \quad (8.2.28)$$

We have seen this transformation before; see, for example equation (3.13.6). Where parameters have a spatial connotation (for example where they are hosted by pilot points), use of this transformation allows transformed parameters \mathbf{j} to occupy the same locations as untransformed parameters \mathbf{k} despite the fact that they are iids.

Similar mathematics to that presented above shows that minimization of the norm of \mathbf{j} leads to minimization of $\mathbf{k}^t\mathbf{C}^{-1}(\mathbf{k})\mathbf{k}$. It follows that \mathbf{j} -space is an effective parameter space in which to apply singular value decomposition. Additionally, it is easy to show that singular value decomposition performed in \mathbf{k} -space yields a \mathbf{j} -space resolution matrix that is idempotent but non-orthogonal showing, once again, the requirement for parameter transformation before undertaking singular value decomposition in order to achieve optimality of model calibration.

8.2.5 Predictive Error Variance

8.2.5.1 Expressing Predictive Error Variance

From equation (8.2.13) the covariance matrix of parameter error can be obtained using the propagation of variance relationship (3.9.2) as

$$\mathbf{C}(\underline{\mathbf{k}} - \mathbf{k}) = \mathbf{V}_2\mathbf{V}_2^t\mathbf{C}(\mathbf{k})\mathbf{V}_2\mathbf{V}_2^t + \mathbf{V}_1\mathbf{S}_1^{-1}\mathbf{U}_1^t\mathbf{C}(\boldsymbol{\varepsilon})\mathbf{U}_1\mathbf{S}_1^{-1}\mathbf{V}_1^t \quad (8.2.29)$$

Suppose that

$$\mathbf{C}(\mathbf{k}) = \sigma_k^2 \mathbf{I} \quad (8.2.30)$$

$$\mathbf{C}(\boldsymbol{\varepsilon}) = \sigma_\varepsilon^2 \mathbf{I} \quad (8.2.31)$$

Ideally (8.2.30) should have been achieved through pre-inversion Kahunen-Loëve parameter transformation while (8.2.31) should have been achieved through pre-inversion weight matrix observation transformation, both of which have been discussed above. Equation (8.2.29) then becomes

$$\mathbf{C}(\underline{\mathbf{k}} - \mathbf{k}) = \sigma_k^2 \mathbf{V}_2\mathbf{V}_2^t + \sigma_\varepsilon^2 \mathbf{V}_1\mathbf{S}_1^{-1}\mathbf{V}_1^t \quad (8.2.32)$$

Let s be a prediction of interest whose sensitivity to parameters \mathbf{k} is expressed by the vector \mathbf{y} . The error in this prediction, as made by the calibrated model, is given by

$$\underline{s} - s = \mathbf{y}^t(\underline{\mathbf{k}} - \mathbf{k}) \quad (8.2.33)$$

so that, through propagation of variance,

$$\sigma_{\underline{s}-s}^2 = \sigma_k^2 \mathbf{y}^t \mathbf{V}_2\mathbf{V}_2^t \mathbf{y} + \sigma_\varepsilon^2 \mathbf{y}^t \mathbf{V}_1\mathbf{S}_1^{-1}\mathbf{V}_1^t \mathbf{y} \quad (8.2.34)$$

This, of course, is equation (8.1.25) formulated for the specific conditions wherein regularization is undertaken using singular value decomposition and the covariance matrices of prior parameter uncertainty and measurement error are expressed by (8.2.30) and (8.2.31).

8.2.5.2 Optimum Truncation Point

Equation (8.2.34) can be used to design optimal partitioning of the \mathbf{U} , \mathbf{S} and \mathbf{V} matrices to define \mathbf{V}_1 , \mathbf{S}_1 and \mathbf{U}_1 for use in equation (8.2.4). Such partitioning is often referred to as “truncation”; singular value decomposition is often referred to as “truncated singular value decomposition”.

It will be recalled that singular values are arranged in the \mathbf{S} matrix from highest to lowest down its diagonal. Obviously, truncation should occur before singular values become zero, for if this does not occur then equation (8.2.4) is asked to invert a non-invertible matrix. However, as will now be shown, the optimum truncation point should occur before this, where singular values become small rather than zero. This prevents so-called “over-fitting” to the calibration dataset, and the consequential potential for parameter and predictive error that over-fitting can incur.

Suppose that we truncate \mathbf{S} after w singular values. Recalling that the number of parameters featured in the inverse problem is m , equation (8.2.34) can be written as

$$\sigma_{s-s}^2 = \sigma_k^2 \sum_{i=w+1}^m (\mathbf{y}^t \mathbf{v}_i)^2 + \sigma_\varepsilon^2 \sum_{i=1}^w s_i^{-2} (\mathbf{y}^t \mathbf{v}_i)^2 \quad (8.2.35)$$

where \mathbf{v}_i are the unit vectors comprising the columns of \mathbf{V} . All terms in both of the summations featured on the right side of equation (8.2.35) are positive. Therefore the first summation decreases as w increases and the second summation increases as w increases. Where there is no truncation at all and w is equal to 0, then

$$\sigma_{s-s}^2 = \sigma_k^2 \mathbf{y}^t \mathbf{V} \mathbf{V}^t \mathbf{y} = \sigma_k^2 \mathbf{y}^t \mathbf{y} = \mathbf{y}^t \mathbf{C}(\mathbf{k}) \mathbf{y} \quad (8.2.36)$$

When w is zero the null space is equivalent to the entirety of parameter space as there is no calibration. Equation (8.2.36) states that the predictive error variance is then equal to the prior uncertainty of the prediction which, of course, depends only on $\mathbf{C}(\mathbf{k})$.

As w increases from 0, the first term of equations (8.2.34) and (8.2.35) falls monotonically as the number of terms in the summation falls. Eventually it becomes zero when w is equal to m , the number of parameters featured in the inverse problem; at this point the dimensions of the null space have shrunk to zero. This first term of equations (8.2.34) and (8.2.35) will be referred to as the “null space term”. It quantifies the cost of uniqueness that is paid in seeking a solution to an ill-posed inverse problem. For singular value decomposition, this cost is paid by seeking a value only for the projection of the real-world parameter set \mathbf{k} onto the solution space of \mathbf{Z} and thereby foregoing entry into the null space of \mathbf{Z} .

The behaviour of the second term of equations (8.2.34) and (8.2.35) is opposite to that of the first; it rises monotonically as w increases from zero. It is zero when w is zero because the summation has no terms; then it rises with w as the number of terms in the summation increases. The rate of increase of the summation becomes very rapid as singular values become small because of the presence of s_i^{-2} in this term. When singular values become zero, this term becomes infinite. It is referred to herein as the “solution space term”. It quantifies the effect of measurement noise on the error variance of a prediction. Note that a prediction can also be the value of a single parameter; this occurs when all elements in \mathbf{y} are zero except for the element corresponding to the pertinent parameter, which is set to one.

The solution and null space terms of equations (8.2.34) and (8.2.35) sum to the total predictive error variance.

The situation is depicted in figure 8.4. As the truncation point increases from zero, predictive error variance initially falls. As the truncation point is increased further to include more

singular values, the dimensions of the null space decrease while the dimensions of the solution space increase. This provides more receptacles for information contained within the calibration dataset; these receptacles are the coefficients by which individual \mathbf{v}_i are multiplied before vector addition to form the overall \mathbf{k} vector calculated through equation (8.2.4). However as the null space term of equations (8.2.34) and (8.2.35) falls, the solution space term rises. Eventually the cost of achieving a better fit with the calibration dataset outweighs the benefit of reducing the dimensionality of the null space. The total predictive error variance curve, which has hitherto fallen with increasing number of singular values, now starts to rise as measurement noise is more greatly amplified because of the presence of s_i^{-2} in the solution space term. “Over-fitting” to the calibration dataset is now occurring. The optimal singular value truncation point corresponds to the minimum of the predictive error variance curve.

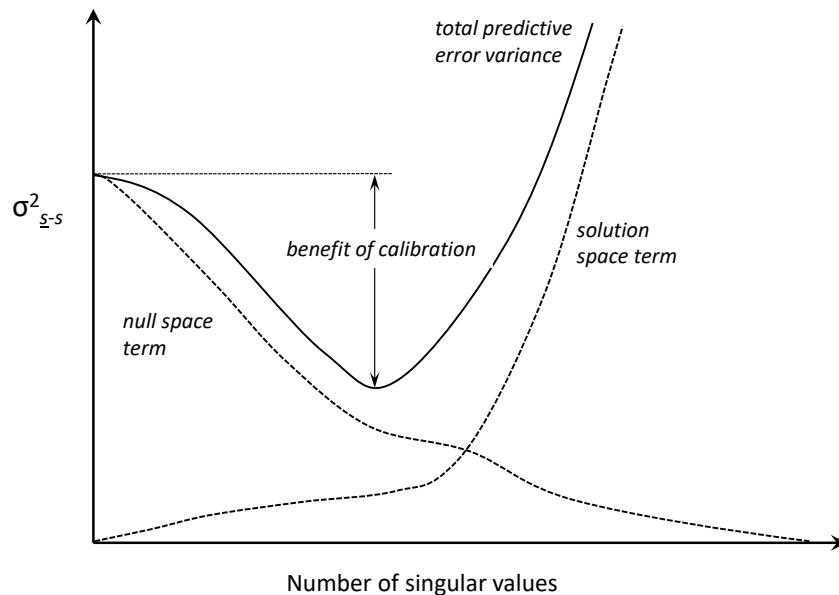


Figure 8.4. Predictive error variance as a function of number of singular values used in the inversion process. Note that as the number of singular values increases, the fit between model outputs and the calibration dataset improves. The horizontal axis of the above figure could therefore be alternatively labelled as “goodness of fit”.

Equations (8.2.34) and (8.2.35) can be used to calculate the benefits of model calibration as these pertain to a specific prediction. This benefit is the reduction in the error variance of that prediction below its pre-calibration level that is accrued through the calibration process. Recall that the pre-calibration error variance of a prediction is equal to the variance of its pre-calibration uncertainty. Maximum predictive benefit is gained where singular value truncation accords with the minimum of the predictive error variance curve. See figure 8.4. It will be shown later in this book that the minimized error variance of a prediction cannot fall below the posterior uncertainty of that prediction as calculated through direct application of Bayes equation; in fact it will probably be a little higher.

Plots such as that depicted in figure 8.4 reveal the components of predictive error variance at any number of singular values, including the optimum number of singular values. A modeller is thus able to ascertain the relative contributions made to predictive error variance by the solution and null space terms of (8.2.34) and (8.2.35). If minimized predictive error variance (and hence posterior predictive uncertainty for which it can be considered as an approximate surrogate) is dominated by the null space term, then it is the deficit of information in the

calibration dataset that is responsible for this error variance. On the other hand, if minimized predictive error variance is dominated by the solution space term then, to the extent that parameters to which a prediction is sensitive are informed by the calibration process, estimates of their values are contaminated by noise associated with the calibration dataset.

The nature of the curve that is obtained by plotting predictive error variance against singular value truncation point is prediction-specific. However the location of the minimum is unlikely to vary greatly from prediction to prediction. Equation (8.2.34) indicates that if a prediction is sensitive to predominantly null space components of \mathbf{Z} then its error variance (and hence uncertainty) will be reduced very little from its pre-calibration level through the calibration process; in this case the calibration dataset contains little information pertaining to parameters to which the prediction is most sensitive. Predictions which depend on system property detail, predictions of system state which are very different from those for which measurements are available in the calibration dataset, and predictions of system response to a stress regime which is very different from that which prevailed when the calibration dataset was acquired, are all likely to fall into this category. This is something that those who are unfamiliar with highly parameterized inversion may not always appreciate. Just because a model has been “calibrated”, it does not follow that its predictive abilities have thereby been enhanced, for this depends entirely on the prediction.

It should be especially noted that failure to represent null space parameterization in a model (as is often done as a manual regularization device to achieve inverse problem well-posedness) may prevent a modeller from ever realizing the inadequacies of the calibration process in lowering the uncertainties of critical model predictions. If null space parameterization is eradicated from a model, the first term of equation (8.2.34) is eliminated; predictive error variance may therefore be grossly underestimated.

It is of interest to note that calculation of predictive error variance using equation (8.2.34) requires only sensitivities of model outputs to model parameters. These sensitivities comprise the elements of the vector \mathbf{y} and the elements of the matrix \mathbf{Z} . The latter is subjected to singular value decomposition to yield the \mathbf{S} and \mathbf{V} matrices featured in this equation. Values of parameters \mathbf{k} , values of measurements \mathbf{h} , and the value of the prediction s do not feature in this equation. This provides some interesting possibilities. For example if the error variance of a prediction can be calculated using the current calibration dataset through equation (8.2.34), then its error variance can also be calculated using an expanded dataset. The values of measurements comprising that expanded dataset are not required. All that is required is that the model calculate the sensitivities of model outputs that correspond to these measurements and supplement the \mathbf{Z} matrix accordingly. The ability of different data acquisition strategies to lower the error variances of predictions of interest can thereby be explored. For a nonlinear model, sensitivities will depend on parameter values, so data worth analysis using equation (8.2.34) can only be approximate. Nevertheless it may be informative. This is further discussed in the following chapter where a Bayesian alternative to linear data worth analysis is described.

Equations (8.2.34) and (8.2.35) were derived under the assumption that $C(\mathbf{k})$ and $C(\boldsymbol{\varepsilon})$ are specified by (8.2.30) and (8.2.31). In the more general case where $C(\mathbf{k})$ and $C(\boldsymbol{\varepsilon})$ are not described by these equations, it follows from (8.2.29) that

$$\sigma_{s-s}^2 = \mathbf{y}^t \mathbf{V}_2 \mathbf{V}_2^t C(\mathbf{k}) \mathbf{V}_2 \mathbf{V}_2^t \mathbf{y} + \mathbf{y}^t \mathbf{V}_1 \mathbf{S}^{-1} \mathbf{U}^t \mathbf{C}(\boldsymbol{\varepsilon}) \mathbf{U}_1 \mathbf{S}^{-1} \mathbf{V}_1^t \mathbf{y} \quad (8.2.37)$$

If σ_{s-s}^2 is plotted against number of pre-truncation singular values, monotonicity of the first and second terms of equation (8.2.37) cannot be guaranteed where (8.2.30) and (8.2.31) do not apply. Furthermore, the minimized predictive error variance will be greater than it would have

been if Kahunen-Loëve and weight matrix transformation had first been undertaken in order to achieve the conditions expressed by (8.2.30) and (8.2.31).

8.2.5.3 Optimizing Truncation without a Prediction

The assistance that equation (8.2.34) can provide in identifying the point of optimal singular value truncation was discussed above. It was shown that the best truncation point is that at which the benefits of further extraction of information from a calibration dataset are balanced by the propensity for parameter and predictive error to increase as a better fit is sought with a calibration dataset. This occurs because the values estimated for some combinations of parameters (these combinations being the vectors \mathbf{v}_i of the matrix \mathbf{V}) become more reflective of measurement noise that accompanies \mathbf{h} than of information that is resident in \mathbf{h} itself, as singular values s_i associated with respective \mathbf{v}_i decrease in value.

The exact shape of the curve that emerges from application of (8.2.34) will vary from prediction to prediction. Some predictions will have their error variances reduced very little through the calibration process. The opposite is the case for others. Also, the exact location of the minimum of the total error variance curve may vary slightly from prediction to prediction, this depending on alignment of the vector \mathbf{y} with the vectors \mathbf{v}_i comprising the columns of \mathbf{V} .

So which prediction should a modeller use when establishing the optimum singular value truncation point for calibration of a model? The answer is “probably none”. Generally, predictive sensitivities are unavailable when a model is being calibrated.

A more practical means of securing truncation point optimality is to examine each term of the two summations that are made explicit when equation (8.2.34) is re-written as (8.2.35). Furthermore, in each term of these summations, the vector \mathbf{y} should be replaced by the respective \mathbf{v}_i vector. The prediction is thus different for each i in the overall summation. For each term, the “true” value of the prediction thus becomes $\sigma_k^2 \mathbf{v}_i^t \mathbf{v}_i$, which is σ_k^2 times the magnitude of \mathbf{v}_i ; this is, of course, equal to σ_k^2 . Starting at the first singular value, the solution and null space estimates for the error variance of the estimated value of $\sigma_k^2 \mathbf{v}_i^t \mathbf{v}_i$ should be compared. The solution space error variance is $\sigma_\epsilon^2 s_i^{-2} (\mathbf{v}_i^t \mathbf{v}_i)^2$ (which is, of course, equal to $\sigma_\epsilon^2 s_i^{-2}$), while the null space error variance is $\sigma_k^2 (\mathbf{v}_i^t \mathbf{v}_i)^2$ (which is, of course, equal to σ_k^2). For low values of i , where s_i is likely to be large, the null space error variance will normally be higher than the solution space error variance. The term should therefore be relegated to the solution space. However eventually an i will be encountered where the opposite occurs. This marks the i at which singular value truncation should occur. If σ_ϵ^2 is equal to σ_k^2 this occurs where s_i becomes less than 1.0.

8.2.6 Well-Posedness and Flow of Information

8.2.6.1 Already Well-Posed

Suppose that the inverse problem described by equation (8.1.1) is well-posed. Let us also assume that weight matrix transformation of the problem has taken place so that equation (8.2.31) is satisfied. Then, based on theory presented in the previous chapter, $\underline{\mathbf{k}}$ can be estimated as

$$\underline{\mathbf{k}} = (\mathbf{Z}^t \mathbf{Z})^{-1} \mathbf{Z}^t \mathbf{h} \quad (8.2.38)$$

If (8.2.1) is substituted for \mathbf{Z} in this equation, (8.2.4) results. It is therefore apparent that singular value decomposition and the Gauss-Newton method described in the previous chapter are equivalent when an inverse problem is well-posed.

From (8.2.1)

$$\mathbf{Z}^t \mathbf{Z} = \mathbf{V} \mathbf{S}^2 \mathbf{V}^t \quad (8.2.39)$$

The singular values of $\mathbf{Z}^t \mathbf{Z}$ are thus the squares of the singular values of \mathbf{Z} . If (8.2.39) is post-multiplied by \mathbf{V} we obtain

$$\mathbf{Z}^t \mathbf{Z} \mathbf{V} = \mathbf{V} \mathbf{S}^2 \quad (8.2.40)$$

From (2.12.1) it is apparent that the vectors \mathbf{v}_i are the eigenvectors of $\mathbf{Z}^t \mathbf{Z}$ while the scalars s_i^2 comprising the diagonals of \mathbf{S}^2 are its eigenvalues. The eigenvalues of $\mathbf{Z}^t \mathbf{Z}$ are therefore the squares of the singular values of \mathbf{Z} .

It is easily verified that

$$(\mathbf{Z}^t \mathbf{Z})^{-1} = \mathbf{V} \mathbf{S}^{-2} \mathbf{V}^t \quad (8.2.41)$$

and that

$$(\mathbf{Z}^t \mathbf{Z})^{-1} \mathbf{V} = \mathbf{V} \mathbf{S}^{-2} \quad (8.2.42)$$

Recall from the previous chapter that, after multiplication by the reference variance, $(\mathbf{Z}^t \mathbf{Z})^{-1}$ is the post-calibration covariance matrix of parameter error. As is discussed in section 7.3, much can be learned about the well-posedness or otherwise of an inverse problem through examining the eigenvectors and eigenvalues of this matrix. In particular, if the ratio of largest to smallest eigenvalue of this matrix is very large, then the inverse problem is verging on ill-posedness and should be reformulated with stronger manual regularization. Unfortunately, however, if the inverse problem is entirely ill-posed, then the $\mathbf{Z}^t \mathbf{Z}$ matrix cannot be inverted at all to obtain the post-calibration parameter error covariance matrix. The parameter combinations that feature in the eigenvector corresponding to the highest eigenvalue of this matrix are not therefore available for inspection in order for one or a number of salient parameters to be eliminated in order to improve manual regularization. This illustrates an important flaw in manually-regularized inversion that was discussed in the previous chapter; the ill-posedness, or incipient ill-posedness, of an inverse problem is not apparent until after (possibly failed) attempts to solve it.

Singular value decomposition allows a modeller to inspect the would-be eigenvalues of a post-calibration parameter error covariance matrix before, rather than after, this matrix is obtained. These are simply the inverse of the squares of the singular values of \mathbf{Z} . If some singular values are zero, or are very low relative to the highest singular value, then ill-posedness, or incipient ill-posedness, of the inverse problem is thereby identified. The eigenvectors corresponding to would-be troublesome eigenvalues are also available in the columns of \mathbf{V} .

8.2.6.2 Well-Posed through SVD

From (8.1.1) and (8.2.2)

$$\mathbf{h} = \mathbf{Z} \mathbf{k} + \boldsymbol{\varepsilon} = \mathbf{U}_1 \mathbf{S}_1 \mathbf{V}^t_1 \mathbf{k} + \mathbf{U}_2 \mathbf{S}_2 \mathbf{V}^t_2 \mathbf{k} + \boldsymbol{\varepsilon} \quad (8.2.43)$$

Adopting the same symbolism as in section 7.1, equation (8.2.43) can be written as

$$\mathbf{h} = \mathbf{X} \mathbf{p} + \boldsymbol{\eta} + \boldsymbol{\varepsilon} \quad (8.2.44)$$

where

$$\mathbf{X} = \mathbf{U}_1 \mathbf{S}_1 \quad (8.2.45)$$

$$\mathbf{p} = \mathbf{V}^t_1 \mathbf{k} \quad (8.2.46)$$

$$\underline{\eta} = \mathbf{U}_2 \mathbf{S}_2 \mathbf{V}^t_2 \mathbf{k} \quad (8.2.47)$$

With appropriate singular value truncation, equation (8.2.44) constitutes a well-posed inverse problem. Meanwhile $\underline{\eta}$ is the structural noise incurred through the simplification that is required to achieve well-posedness of the inverse problem. The fewer non-zero singular values that feature in \mathbf{S}_2 , the smaller is $\underline{\eta}$. The elements of $\underline{\mathbf{p}}$ are the scalar projections of the unknown parameter vector \mathbf{k} onto each of the unit vectors that comprise the columns of \mathbf{V}_1 ; recall that the columns of \mathbf{V}_1 span the solution space of \mathbf{Z} . Once $\underline{\mathbf{p}}$ has been estimated, the solution vector in \mathbf{k} -space can be recovered from the solution vector in \mathbf{p} -space by multiplying each of these scalar projections by the corresponding unit vector onto which it is projected. Thus

$$\underline{\mathbf{k}} = \mathbf{V}_1 \underline{\mathbf{p}} \quad (8.2.48)$$

A solution to this well-posed inverse problem can be obtained, according to methods discussed in the previous chapter, as

$$\underline{\mathbf{p}} = (\mathbf{X}^t \mathbf{X})^{-1} \mathbf{X}^t \mathbf{h} = (\mathbf{S}_1 \mathbf{U}^t_1 \mathbf{U}_1 \mathbf{S}_1)^{-1} \mathbf{S}_1 \mathbf{U}^t_1 \mathbf{h} = \mathbf{S}^{-1} \mathbf{U}^t_1 \mathbf{h} \quad (8.2.49)$$

which, combined with (8.2.48), leads to (8.2.4).

So singular value decomposition has in common with manual regularization the fact that it achieves inverse problem solution uniqueness through parameter simplification. The simplification strategy adopted by singular value decomposition is the estimation of strategic combinations of parameters in place of individual parameters of the original inverse problem. While estimation of parameter combinations sometimes constitutes a simplification strategy sought through manual regularization (where, for example, a modeller may link parameters in constant ratios or combine them to form homogenous zones), parameter combinations defined by singular value decomposition are optimal. Provided it is preceded by Kahunen-Loëve transformation of parameters, the outcome of regularized inversion achieved through singular value decomposition is a solution of minimum error variance to the inverse problem. This occurs because simplification is carried by the parameter estimation process itself which, through its ability to separate combinations of parameters which are estimable (columns of the \mathbf{V}_1 matrix) from combinations of parameters which are inestimable (columns of the \mathbf{V}_2 matrix), is able to isolate the latter so that only the former are estimated.

8.2.6.3 Flow of Information

If both sides of equation (8.2.4) are multiplied by \mathbf{V}^t_1 the following equation is obtained

$$\mathbf{S}^{-1} \mathbf{U}^t_1 \mathbf{h} = \mathbf{V}^t_1 \underline{\mathbf{k}} \quad (8.2.50)$$

If there is no noise in the calibration dataset $\mathbf{V}^t_1 \underline{\mathbf{k}}$ can be replaced by $\mathbf{V}^t_1 \mathbf{k}$ in the above equation. (Under these circumstances, the projection of the real parameter set \mathbf{k} onto the solution space is $\underline{\mathbf{k}}$.) Because \mathbf{S}^{-1} is a diagonal matrix, (8.2.50) can be re-written as a series of equations as follows

$$s_i^{-1} \mathbf{u}^t_i \mathbf{h} = \mathbf{v}^t_i \mathbf{k} \quad (8.2.51)$$

$\mathbf{u}^t_i \mathbf{h}$ is the scalar projection of the calibration dataset onto a single column of the \mathbf{U}_1 matrix. (The columns of \mathbf{U}_1 span part or all of the range space of \mathbf{Z} .) $\mathbf{v}^t_i \mathbf{k}$ is the scalar projection of the parameter set \mathbf{k} into the i 'th column of the \mathbf{V}_1 matrix. (The columns of \mathbf{V}_1 span the solution space of \mathbf{Z} .) Each of these scalar projections is, of course, a scalar. There is one scalar that links the scalar projection on the left side of (8.2.51) to that on its right side. This scalar is s_i^{-1} . This, together with the orthonormality of \mathbf{U} and \mathbf{V} , implies that $\mathbf{u}^t_i \mathbf{h}$ is uniquely and entirely informative of $\mathbf{v}^t_i \mathbf{k}$. \mathbf{u}_i is thus the “characteristic” or “signal” in the field-measured dataset \mathbf{h}

that speaks solely and uniquely about the “characteristic” or “signal” that is \mathbf{v}_i in the model’s parameter set. To put it another way, (8.2.51) identifies orthogonal combinations of observations (i.e. “signals” in the calibration dataset) that are uniquely and entirely informative of orthogonal combinations of parameters (i.e. “characteristics” of the model’s parameter field).

The \mathbf{u}_i defined through singular value decomposition of \mathbf{Z} thus constitute individual and separate sources of information in the observation dataset. Regardless of the size of this dataset (i.e. regardless of the number of elements of \mathbf{h}), there are only as many of these sources of information available to the parameter estimation process as there are pre-truncation singular values employed in that process. \mathbf{v}_i constitutes the receptacle for the information that is \mathbf{u}_i .

Where observations and parameters have spatial or temporal associations, it is sometimes interesting to plot the elements of high index \mathbf{u}_i and \mathbf{v}_i vectors in space and time. The spatial/temporal distribution of measurement data that constitutes a single piece of information on the one hand, and the spatial/temporal distribution of properties of an environmental system that constitutes the receptacle for that information on the other hand, then become apparent.

After model parameters have been estimated, residuals can be calculated as

$$\mathbf{r} = \mathbf{h} - \mathbf{Zk} = \mathbf{h} - \mathbf{ZV}_1\mathbf{S}_1^{-1}\mathbf{U}_1^t\mathbf{h} \quad (8.2.52)$$

which, after substitution of \mathbf{USV}^t for \mathbf{Z} , becomes

$$\mathbf{r} = (\mathbf{I} - \mathbf{U}_1\mathbf{U}_1^t)\mathbf{h} \quad (8.2.53)$$

Direct multiplication shows that

$$\mathbf{r}^t \mathbf{Zk} = \mathbf{0} \quad (8.2.54)$$

Residuals are thus orthogonal to outputs generated by the calibrated model. It does not necessarily follow, however, that the measurement dataset \mathbf{h} has thereby been emptied of all of its information. To the extent that singular value truncation takes place prior to singular values becoming zero, the residuals vector \mathbf{r} may have a non-zero projection onto some of the \mathbf{u}_i vectors which comprise the columns of \mathbf{U}_2 . However if singular value truncation is optimal, the information that is associated with these \mathbf{u}_i vectors cannot be assimilated into the model’s parameter set because noise within the measurement dataset that is also projected onto these \mathbf{u}_i has the potential to worsen, rather than improve, estimates of parameter combinations to which that information is directed.

8.2.7 Strengths and Weaknesses of Singular Value Decomposition

As a method for solving an ill-posed inverse problem, singular value decomposition has many attractive features. Its numerical stability in the face of problem ill-posedness is an obvious advantage. Another nice feature of singular value decomposition is the conceptual insights that it offers into what can be achieved, and what cannot be achieved, in solving an ill-posed inverse problem. It also illuminates how information flows from measurements to parameters. In doing so, it confirms our previous conjecture that parameters can be considered as hosts for information that is harvested from a calibration dataset.

A further numerical benefit that singular value decomposition offers is illustrated by equations (8.2.44) to (8.2.47). Once a singular value truncation point has been selected, a set of parameters \mathbf{p} can be defined through (8.2.46). In PEST parlance, these are so-called “super parameters”. In many modelling contexts, the number of parameters featured in the \mathbf{p} vector is much smaller than those that are featured in the \mathbf{k} vector. For a given \mathbf{p} , a corresponding \mathbf{k} can

be calculated as

$$\mathbf{k} = \mathbf{V}_1 \mathbf{p} \quad (8.2.55)$$

In a parameter estimation setting where the Jacobian matrix must be computed using finite parameter differences, a \mathbf{Z} matrix can first be computed by undertaking m model runs, where m is the number of elements in \mathbf{k} . A smaller set of parameters \mathbf{p} can then be defined using (8.2.46). This smaller set of parameters can then form the basis for further iterations of the nonlinear parameter estimation process using equation (7.1.6) in which a matrix \mathbf{X} acts on parameters \mathbf{p} . Finite differences of \mathbf{p} -parameters and corresponding model outputs can be used to compute a Jacobian matrix to represent \mathbf{X} ; for every \mathbf{p} vector used in finite differencing to fill the Jacobian matrix, or in testing updated parameters calculated using different Marquardt lambdas, a corresponding \mathbf{k} vector is calculated using (8.2.55) prior to running the model. Actually, there is no need for the singular value truncation process on which basis \mathbf{p} is calculated to be such as to guarantee a well-posed inverse problem; singular value decomposition of \mathbf{X} can be used to estimate \mathbf{p} so that numerical stability is ensured. Once the inverse problem has been solved, and \mathbf{p} has been estimated, a corresponding \mathbf{k} is then calculated using (8.2.55). This is the basis of the “SVD-assist” methodology implemented in PEST; see Tonkin and Doherty (2005) for further details. (Note that it is recommended practice for Tikhonov regularization to be included in “SVD-assisted” solution of an ill-posed inverse problem. Tikhonov regularization will be discussed shortly. SVD-assist is further discussed in chapter 9.)

One weakness of singular value decomposition as a parameter estimation device in the environmental modelling context is the rigidity of the receptacles that it provides for expert knowledge. As has already been discussed, use of singular value decomposition to attain a minimum error variance solution to an ill-posed inverse problem requires that model parameters be subjected to Kahunen-Loëve transformation prior to commencement of the inversion process. In many environmental contexts it may be difficult or impossible to provide such a transformation. In the discussion above, Kahunen-Loëve transformation is based on the prior covariance matrix $\mathbf{C}(\mathbf{k})$. However the complexity and heterogeneity of natural systems (for example geological media) is such that a $\mathbf{C}(\mathbf{k})$ matrix may be very difficult to formulate at best, and may provide an inappropriately simple descriptor of prediction-relevant heterogeneity at worst. This will almost certainly be the case if $\mathbf{C}(\mathbf{k})$ is formulated under the notion that parameter correlation lengths are stationary throughout a model domain. However, the subsurface is much more complex that can be described even by nonstationary geostatistics. Furthermore, as discussed in chapter 5, uncertainties in geostatistical hyperparameters make $\mathbf{C}(\mathbf{k})$ uncertain. A parameter transformation that is based on a simplistic $\mathbf{C}(\mathbf{k})$ cannot therefore provide a totally reliable basis for calibration optimality.

Suppose, nevertheless, that a $\mathbf{C}(\mathbf{k})$ matrix that encapsulates expert knowledge of system property variability and correlation is indeed formulated. Singular value decomposition can then be used to estimate coefficients of its eigencomponents; see section 8.2.4. While this may provide the benefit that the prior covariance matrix of estimated, transformed parameters is \mathbf{I} , the inversion process is then locked into the use of these eigencomponents as the sole expression of system property variability and heterogeneity. This suppresses the ability of the inversion process to reveal the presence of system property heterogeneity of which a modeller may have been previously unaware. In doing so, it annuls one of the benefits of adopting a highly parameterized approach to inversion in the first place, this being an ability to formulate the inverse problem in such a way as to provide the inversion process with as much freedom as it needs to respond to information contained in the calibration dataset in ways that are

flexible and informative to a modeller.

Another practical problem with the use of singular value decomposition in real-world model calibration is the prevention of over-fitting. While the theory of singular value decomposition provides elegant concepts through which over-fitting can be defined and avoided (see figure 8.4), in many environmental contexts neither $C(\mathbf{k})$ nor $C(\boldsymbol{\varepsilon})$ is known. Furthermore model-to-measurement misfit is likely to be dominated by structural noise arising from model defects rather than by noise associated with a measurement dataset. Because of this, a modeller often learns through the calibration process itself the extent of parameter variability and the magnitude of model-to-measurement misfit that he/she is prepared to tolerate. In fact, these are often subjectively traded off against each other until a modeller finds a level of parametric heterogeneity, and a corresponding level of model-to-measurement misfit with which he/she feels comfortable. Where singular value decomposition is a modeller's only means of regularization, the only control that he/she can exert on the level of model-to-measurement fit attained through the inversion process is the number of singular values at which truncation takes place; either this, or the ratio of highest to lowest pre-truncation singular value can be used to define the truncation point. One or the other of these must be provided by the modeller prior to commencement of the inversion process. Unfortunately, the relationship between either of these and the final objective function attained through the inversion process is often unclear, and must be ascertained through trial and error. This can be a computer-intensive and time-consuming business.

Fortunately, as will now be discussed, Tikhonov regularization as implemented by PEST, offers a modeller more direct control over the level of model-to-measurement fit that is sought through an inversion process.

8.3 Tikhonov Regularization

8.3.1 Concepts

Tikhonov regularization achieves well-posedness of an inverse problem by supplementing the information content of a calibration dataset with expert knowledge. Conceptually the idea is very simple, and has already been expressed through equation (7.2.38) wherein prior information pertaining to all parameters is added to a calibration dataset. In equation (7.2.38) prior information is weighted in inverse proportion to the prior parameter covariance matrix. Hence its influence on the outcome of the inversion process is proportional to the strength of expert knowledge in comparison to the credibility of the measurement dataset as expressed through weights assigned to various components of that dataset; ideally weights should vary in inverse proportion to measurement noise. As was stated earlier, where the prior probability distributions of parameters on the one hand and measurement noise on the other hand are both multiGaussian, the outcome of the inversion process formulated as (7.2.38) is a set of parameters that correspond to the mean of the posterior parameter probability distribution. This probability distribution is itself multiGaussian if the model is linear.

In practice, it is rarely possible to balance prior knowledge against measurement noise in this way. As has been stated, the level of "measurement noise" (which in most real-world environmental modelling contexts is predominantly structural noise) is rarely known before inversion is attempted; in fact, knowledge of its magnitude and characteristics is often forthcoming from the inversion process itself. The same applies to expert knowledge, for the uncertainties of system properties and the characteristics of their spatial variability are often very difficult to specify. Furthermore, flexible methods of mathematically encapsulating expert

knowledge that extend beyond a simple prior covariance matrix are often required by an inversion process if the parameter field that emerges from that process is to reflect the complex dispositions of hydraulic property heterogeneity that prevail at a particular study site. At the same time, flexibility is required in balancing expert knowledge, regardless of how it is mathematically expressed, against field measurements so that over-fitting is prevented while, at the same time, granting freedom to the inversion process to insert hydraulic property heterogeneity into a model domain where this is required for model outputs to match field measurements.

Let the calibration dataset \mathbf{h} be supplemented by expert knowledge \mathbf{w} . Equation (8.1.1) then becomes

$$\begin{bmatrix} \mathbf{h} \\ \mathbf{w} \end{bmatrix} = \begin{bmatrix} \mathbf{Z} \\ \mathbf{Z}_w \end{bmatrix} \mathbf{k} + \begin{bmatrix} \boldsymbol{\varepsilon} \\ \boldsymbol{\omega} \end{bmatrix} \quad (8.3.1a)$$

In equation (8.3.1a) \mathbf{Z}_w is a “regularization model” which operates on \mathbf{k} to produce a set of regularization outcomes. The vector \mathbf{w} contains the “observed values” of these outcomes. $\boldsymbol{\omega}$ is the “noise” associated with \mathbf{w} . Its inverse is therefore a measure of the strength of expert belief in \mathbf{w} . As such it expresses the extent to which failure of the inversion process to produce a set of parameters \mathbf{k} which respects \mathbf{w} will be tolerated.

Equation (8.3.1a) can be written as two separate matrix equations. The first of these is the same as (8.1.1). The second of these equations expresses regularization constraints on parameter values. The equations are

$$\mathbf{h} = \mathbf{Z}\mathbf{k} + \boldsymbol{\varepsilon} \quad (8.3.1b)$$

$$\mathbf{w} = \mathbf{Z}_w\mathbf{k} + \boldsymbol{\omega} \quad (8.3.1c)$$

The nature of expert knowledge that is expressed by \mathbf{Z}_w varies with context. It can be as simple as the observation that each parameter \mathbf{k} is equal to its prior mean value, as was done for equation (7.2.38). Alternatively, it can be considerably more complex than this. In spatial parameterization contexts to which Tikhonov regularization is commonly applied, \mathbf{Z}_w may express a homogeneity condition through which it is “observed” that the value of each parameter is equal to that of its neighbour, or possibly to those of many of its neighbours. Weights assigned to these observations may decrease with increasing parameter separation; see, for example, Doherty (2003). In other cases, a parameter smoothness condition may be expressed whereby differences of differences (i.e. second spatial derivatives) of parameter values are observed to be zero; this type of regularization is sometimes employed in geophysical data interpretation where a parameter is associated with every grid cell of a spatial model domain. If spatial heterogeneity arises in solving the inverse problem, it therefore arises in a smooth, “spread out” way. In other spatial parameterization contexts, however, this may be an inappropriate way for parameter heterogeneity to arise. In such cases it may be desired that spatial heterogeneity, should it need to arise, be confined to as small an area as possible, and perhaps have reasonably sharp boundaries with neighbouring areas where homogeneity prevails. Alternatively, or at the same time, it may be desired that heterogeneity be aligned in a certain direction, or that it arise as close to a certain known geological feature as possible. All of these conditions can be formulated through appropriate Tikhonov constraints imposed on parameter values estimated through the inversion process. These constraints may not be linear as implied by the \mathbf{Z}_w matrix of equation (8.3.1). But then neither is a numerical model generally linear. These types of complex nonlinear constraints are then accommodated in the same way as the nonlinear behaviour of the model, that is through computing an appropriate Jacobian matrix which is updated during every iteration of the nonlinear parameter estimation process.

A fundamental difference between “regularization observations” encapsulated in \mathbf{w} and observations of system state encapsulated in \mathbf{h} is that the former are expected to be violated while the latter are not. In a spatial model where Tikhonov regularization embodies a parameterization smoothness constraint, it is not expected that the solution to the inverse problem will express no system property heterogeneity whatsoever. Tikhonov constraints are formulated in full knowledge of the fact that heterogeneity will arise, and hence that these constraints will be violated; however the importance of their formulation is that they give the modeller control over *how* heterogeneity arises, and hence *how* regularization constraints on parameters are optimally violated. At the same time, a modeller does not want heterogeneity to arise unless its existence is supported by a measurement dataset; nor does he/she want overfitting of model outputs to field measurements to precipitate the emergence of spurious heterogeneity. For both of these reasons, it is important that the parameter estimation process be given control over the relative commitment of that process to respecting measurements represented by \mathbf{h} on the one hand, and regularization constraints encapsulated in \mathbf{w} on the other hand.

Of course it is not necessary that Tikhonov regularization be confined to spatial parameterization schemes such as are employed by groundwater and subsurface reservoir models and models used in geophysical data interpretation. Tikhonov regularization can also play an important role in calibration of surface water and land use models. In doing so they may form the basis for parameter regionalization. Thus parameters which pertain to complex physical or chemical processes in neighbouring watersheds may be “observed” to have similar values if land use and soil conditions are the same in those watersheds; however they may be allowed to differ from each other if a good fit with streamflows measured at respective watershed gauging stations deems this necessary.

Because regularized inversion places no theoretical limits on the number of parameters that can be accommodated in an inversion process, calibration of multiple submodels associated with multiple watersheds feeding multiple stream gauges can all be done at once. This allows the modeller to “suggest”, through Tikhonov constraints, that certain ordering relationships prevail among parameters, these reflecting the physical/chemical characteristics of the watersheds that feature in the inversion process. Alternatively, certain relationships may be posited between watershed parameters and independently measurable characteristics of these watersheds; perhaps these relationships may be encapsulated in a regression equation. The parameters governing the regression equation can then be estimated at the same time as parameters for the individual watershed models. Regression and watershed model parameters can then be linked through Tikhonov constraints which suggest maximum respect by watershed parameters for regional regression relationships, while tolerating departures from these relationships where these are necessary to fit local streamflow data.

In short, there is no limit to the flexibility that a modeller can exercise in formulating a set of Tikhonov constraints that comprise expressions of expert knowledge that are most pertinent to a particular study site.

Actually, the Tikhonov regularization process can be viewed in a number of ways. One of these has just been described. Another is to view Tikhonov constraints as encapsulating a default parameter condition. This is the condition that will prevail unless there is information to the contrary contained in a calibration dataset. They thus express a condition which is deemed to be of minimum parameter error variance where the only available information pertaining to parameters is expert knowledge. The status of minimum prior parameter error variance does not make this condition right; it only minimizes its potential for wrongness. The formulation

of Tikhonov constraints which instruct the inversion process to respect this prior parameter condition is then “safe” or “conservative”. This is because it instructs the inversion process to express no departures from this preferred condition that cannot be supported by data. What is equally important is that the manner in which these (possibly nonlinear) Tikhonov constraints are formulated and weighted must also specify *how* departures from this prior parameter condition should arise if, indeed, they need to arise in order for pertinent model outputs to fit historical measurements of system states and fluxes.

8.3.2 Formulation of Equations

Ideally, with \mathbf{h} supplemented by \mathbf{w} , the inverse problem of equation (8.3.1) is well-posed and can thus be solved using the Gauss-Newton method.

Let the weight matrix used for measurements be designated as \mathbf{Q} ; normally this matrix is diagonal. As usual, measurement weighting should reflect the relative noise contents of different components of an observation dataset. Where the overall objective function has a number of components, an adopted weighting strategy may also ensure that these different components have roughly the same visibility in the initial objective function. More will be said on this topic later in this book.

In the case of regularization weights, relative weighting between different constraints should reflect the comparative tolerance of a modeller to their violation. For some or all of these constraints, the weight matrix may be non-diagonal. It may, for example, comprise the inverse of a $\mathbf{C}(\mathbf{k})$ matrix which expresses correlation between parameters arising from spatial proximity or physical/chemical property co-dependence on a more fundamental attribute of the medium to which they pertain. We will denote this matrix as $\mu^2 \mathbf{Q}_r$. The factor μ^2 is made explicit to denote the fact that the regularized inversion process is allowed to adjust this factor. This factor is referred to as the “regularization weight multiplier” in the discussion that follows.

In accordance with equation (7.2.14), the solution to the presumably well-posed inverse problem described by equation (8.3.1) is

$$\underline{\mathbf{k}} = \left\{ [\mathbf{Z}^t \quad \mathbf{Z}_w^t] \begin{bmatrix} \mathbf{Q} & \mathbf{0} \\ \mathbf{0} & \mu^2 \mathbf{Q}_r \end{bmatrix} [\mathbf{Z} \quad \mathbf{Z}_w] \right\}^{-1} [\mathbf{Z}^t \quad \mathbf{Z}_w^t] \begin{bmatrix} \mathbf{Q}^t & \mathbf{0} \\ \mathbf{0} & \mu^2 \mathbf{Q}_r \end{bmatrix} [\mathbf{h}] \quad (8.3.2a)$$

which, when all the matrix multiplications are carried out, becomes

$$\underline{\mathbf{k}} = [\mathbf{Z}' \mathbf{Q} \mathbf{Z} + \mu^2 \mathbf{Z}_w^t \mathbf{Q}_r \mathbf{Z}_w]^{-1} [\mathbf{Z}' \mathbf{Q} \mathbf{h} + \mu^2 \mathbf{Z}_w^t \mathbf{Q}_r \mathbf{w}] \quad (8.3.2b)$$

Let us assume that \mathbf{w} is equal to $\mathbf{0}$. This does not have to be the case, but is often so in practice. This follows from the role of \mathbf{w} as expressing a preferred parameter condition and the fact that, for a linearized model, prior parameter expected values are zero; see section 6.5. In this case equation (8.3.2b) becomes

$$\underline{\mathbf{k}} = [\mathbf{Z}' \mathbf{Q} \mathbf{Z} + \mu^2 \mathbf{Z}_w^t \mathbf{Q}_r \mathbf{Z}_w]^{-1} \mathbf{Z}' \mathbf{Q} \mathbf{h} \quad (8.3.3)$$

It will be recalled from the previous chapter that equations (8.3.2) and (8.3.3) constitute solutions to an objective function minimization problem. In the present case the objective function has two components, one pertaining to model-to-measurement misfit and the other pertaining to respect for Tikhonov constraints. In PEST parlance, these are referred to as the measurement and regularization objective functions respectively; they are designated as Φ_m and Φ_r and defined as follows

$$\Phi_m = (\mathbf{h} - \mathbf{Z}\underline{\mathbf{k}})^t \mathbf{Q} (\mathbf{h} - \mathbf{Z}\underline{\mathbf{k}}) \quad (8.3.4)$$

$$\Phi_r = (\mathbf{w} - \mathbf{Z}_w \underline{\mathbf{k}})^t \mathbf{Q}_r (\mathbf{w} - \mathbf{Z}_w \underline{\mathbf{k}}) \quad (8.3.5)$$

If \mathbf{w} is $\mathbf{0}$, (8.3.5) becomes

$$\Phi_r = \mathbf{k}^t \mathbf{Z}_w^t \mathbf{Q}_r \mathbf{Z}_w \mathbf{k} \quad (8.3.6)$$

The total objective function Φ_t which is minimized through solution of (8.3.2) or (8.3.3) is

$$\Phi_t = \Phi_m + \mu^2 \Phi_r \quad (8.3.7)$$

8.3.3 Optimization of Regularization Weight Factor

Equations that have just been presented are similar to those that were presented in the previous chapter for solution of a well-posed inverse problem where well-posedness is sought by supplementing the calibration dataset with expert knowledge; see equations (7.2.38). Fundamental to robust implementation of Tikhonov regularization, however, is that μ^2 be adjusted as part of the inversion process together with model parameters whose estimates are sought through that process. There are a number of different ways in which this can be done. A numerically costly, but very instructive, option is to solve for \mathbf{k} using (8.3.2) or (8.3.3) for many different values of μ^2 , starting with very high values of μ^2 . For (8.3.3) the solution to the inverse problem with μ^2 set to a very high value is obviously that \mathbf{k} is equal to $\mathbf{0}$; for this value of \mathbf{k} the regularization objective function defined by (8.3.5) is zero. Meanwhile the measurement objective function is probably seriously non-zero as model outputs under these conditions are unlikely to provide a good match with the calibration dataset.

As μ^2 is lowered, the measurement dataset \mathbf{h} starts to have some impact on estimated \mathbf{k} . The regularization objective function starts to rise as \mathbf{k} departs from its preferred condition and the measurement objective function starts to fall. Normally this process continues until a point is reached where the regularization objective function rises very quickly for only a small decrease in the measurement objective function. Figure 8.5 shows a real-world example. (The numbers on this graph are case-specific.)

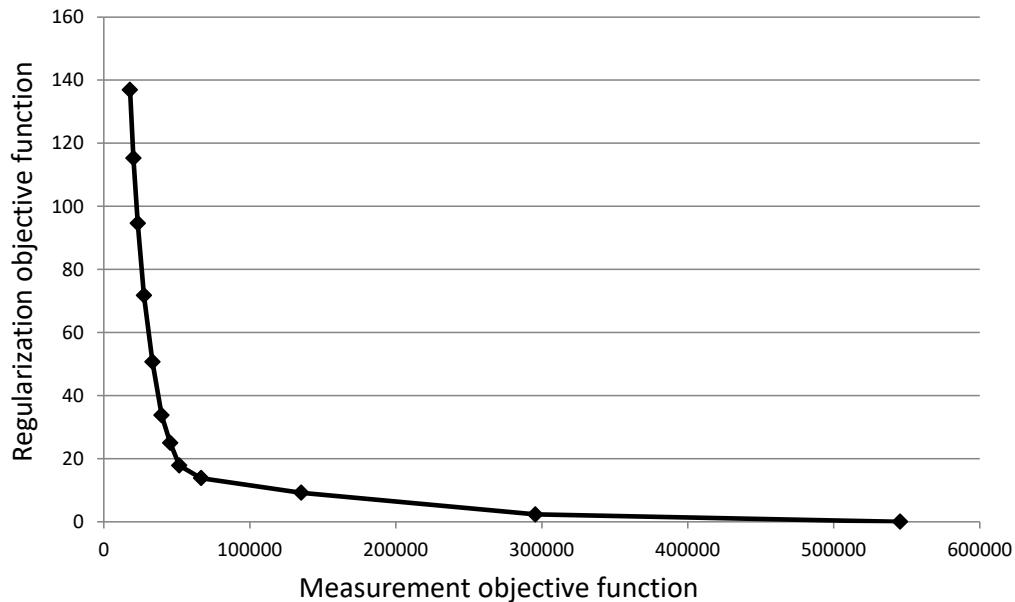


Figure 8.5. Trade-off between measurement and regularization objective functions.

The curve shown in figure 8.5 is often referred to as the “L-curve”. Intuitively, the optimum value of μ^2 is that at which the curve deflects sharply upwards. The requirement that parameters deviate markedly and pervasively from their preferred condition at this point for very little improvement in model-to-measurement fit suggests that they are responding to measurement

or structural noise that is associated with the measurement dataset rather than to real information that is contained within that dataset.

PEST is able to compute an L-curve when it is run in Pareto mode. The term “Pareto” is often used to describe a numerical process wherein two or more competing objective functions are traded off against each other.

Another option in seeking an optimal value of μ^2 is that suggested by de Groot-Hedlin and Constable (1990) and Doherty (2003); this is also implemented in PEST. Instead of directly seeking a value for μ^2 , a modeller sets a desired value for Φ_m , the measurement objective function. This value is referred to herein as the “target measurement objective function” and is designated as Φ_m^t . The Tikhonov inversion process then becomes a constrained minimization process formulated thus:

Minimize the regularization objective function Φ_r subject to the constraint that Φ_m is equal to a user-specified target measurement objective function Φ_m^t .

The inversion process then determines μ^2 itself. If a modeller supplies a value for Φ_m^t that is too low to achieve, then μ^2 should not be reduced any lower than a value that maintains numerical stability of the inversion process while Φ_m is lowered as much as possible.

The Tikhonov-regularized inversion process, formulated as above, is thus asked to seek a certain level of fit with a measurement dataset. At the same time, it is informed that it must not seek a level of fit which is any better than this. The measurement objective function is not therefore minimized. In the previous chapter, it was minimization of the objective function that established uniqueness of the inverse problem. Things are different in the present chapter. Presumably many sets of parameters are compatible with a target measurement objective function that is elevated above its minimum. The Tikhonov-regularized inversion process maintains uniqueness by insisting that the set of parameters that is deemed to solve the inverse problem is that which requires least deviation from the default parameter condition for which the regularization objective function Φ_r is zero. If the default parameter condition is one of minimized parameter error variance from an expert knowledge point of view, then the solution to the inverse problem can lay partial claim to a status of minimized post-calibration parameter error variance due to the fact that in fitting the calibration dataset it pays maximum respect to expert knowledge. It is important to note, however, that this can constitute only part of the basis for such a claim. The other important factor that complements proximity to the default parameter condition is the manner in which deviations from that condition have arisen in order to attain the desired level of fit. It is the design details of Φ_r , and the \mathbf{Q}_r matrix used in formulation of Φ_r , that achieves this.

So to what value should a modeller set Φ_m^t , the target measurement objective function? If measurement noise associated with the calibration dataset is the only determinant of model-to-measurement misfit, and if this noise is normally distributed, and if Φ_m^t is calculated from residuals \mathbf{r} as $\mathbf{r}^T \mathbf{C}^{-1}(\mathbf{e}) \mathbf{r}$, then setting Φ_m^t equal to n , the number of observations comprising the calibration dataset, seems appropriate. This is the expected value of a chi-squared distribution with n degrees of freedom. In practice however (as has been stated many times herein), it is not measurement noise that determines model-to-measurement misfit. In most cases, model-to-measurement misfit is dominated by so-called structural noise which reflects the fact that an environmental model is not a perfect simulator of complex, real-world, environmental processes. It is unlikely in the extreme that structural noise will be homoscedastic, independent, or respect a multiGaussian distribution. Furthermore its magnitude must generally be discovered during the calibration process itself. Selection of an appropriate Φ_m^t thus becomes

subjective.

One option for the setting of Φ_m^t is to first estimate parameters with Φ_m^t set unachievably low. This informs the modeller just how good a fit with the calibration dataset it is possible to obtain. Signs of over-fitting will probably be plainly visible in the calibrated parameter field which achieves this level of fit. In the ensuing inversion process, Φ_m^t can then be set 5% to 10% higher than this minimum. It is normally found that with only a slight diminution of model-to-measurement fit, a parameter field can be obtained that deviates from its Tikhonov-defined default condition by a readily-acceptable amount. This is suggested by L-curves such as that plotted in figure 8.5.

8.3.4 Relationship between Estimated and Real Parameters

With substitution of (8.1.1) into (8.3.3) we obtain

$$\underline{\mathbf{k}} = [\mathbf{Z}^t \mathbf{Q} \mathbf{Z} + \mu^2 \mathbf{Z}_w^t \mathbf{Q}_r \mathbf{Z}_w]^{-1} \mathbf{Z}^t \mathbf{Q} \underline{\mathbf{k}} + [\mathbf{Z}^t \mathbf{Q} \mathbf{Z} + \mu^2 \mathbf{Z}_w^t \mathbf{Q}_r \mathbf{Z}_w]^{-1} \mathbf{Z}^t \mathbf{Q} \boldsymbol{\epsilon} \quad (8.3.8)$$

From this equation the **G** and **R** matrices that are associated with Tikhonov regularization can be inferred as follows

$$\mathbf{G} = [\mathbf{Z}^t \mathbf{Q} \mathbf{Z} + \mu^2 \mathbf{Z}_w^t \mathbf{Q}_r \mathbf{Z}_w]^{-1} \mathbf{Z}^t \mathbf{Q} \quad (8.3.9)$$

$$\mathbf{R} = [\mathbf{Z}^t \mathbf{Q} \mathbf{Z} + \mu^2 \mathbf{Z}_w^t \mathbf{Q}_r \mathbf{Z}_w]^{-1} \mathbf{Z}^t \mathbf{Q} \mathbf{Z} \quad (8.3.10)$$

With these matrices defined, the covariance matrix of parameter error can be calculated using equation (8.1.19). Predictive error variance can be calculated using equation (8.1.25).

If, for a given prediction made by a calibrated, Tikhonov-regularized model, the first, second and summed terms of equation (8.1.25) are plotted against $1/\mu^2$, curves similar to those depicted in figure 8.4 result. A low value for $1/\mu^2$ (i.e. a high value for μ^2) signifies no calibration at all. A high value for $1/\mu^2$ (i.e. a low value for μ^2) signifies over-fitting. An optimum value for μ^2 lies between the two. For this value of μ^2 predictive error variance is minimized. See Moore and Doherty (2005) for further details.

Suppose that $\mathbf{C}(\boldsymbol{\epsilon})$ is the identity matrix **I**. (If necessary, this can be achieved through the transformations listed in equations 8.1.2 to 8.1.6). The weight matrix **Q** can then be chosen as **I**, so that equation (8.3.3) becomes

$$\underline{\mathbf{k}} = [\mathbf{Z}^t \mathbf{Z} + \mu^2 \mathbf{Z}_w^t \mathbf{Q}_r \mathbf{Z}_w]^{-1} \mathbf{Z}^t \mathbf{h} \quad (8.3.11)$$

Now let us further suppose that Tikhonov regularization is formulated to constrain all parameters to be individually equal to their prior expected values, and that the weighting matrix that is associated with these constraints is $\mathbf{C}^{-1}(\underline{\mathbf{k}})$, the inverse of the prior parameter covariance matrix. The above equation then becomes

$$\underline{\mathbf{k}} = [\mathbf{Z}^t \mathbf{Z} + \mu^2 \mathbf{C}^{-1}(\underline{\mathbf{k}})]^{-1} \mathbf{Z}^t \mathbf{h} \quad (8.3.12)$$

Now, instead of solving for $\underline{\mathbf{k}}$, let us assume that we are solving for $\underline{\mathbf{m}}$, a Kahunen-Loëve transformed parameter set for which $\mathbf{C}(\underline{\mathbf{m}}) = \mathbf{I}$. $\underline{\mathbf{k}}$ can then be calculated from $\underline{\mathbf{m}}$ using (8.2.18). \mathbf{Z} in (8.3.12) must therefore be replaced by \mathbf{W} defined by equation (8.2.22). Because $\mathbf{C}(\underline{\mathbf{m}})$ is **I**, then so too is its inverse $\mathbf{C}^{-1}(\underline{\mathbf{m}})$. Equation (8.3.12) then becomes

$$\underline{\mathbf{m}} = [\mathbf{W}^t \mathbf{W} + \mu^2 \mathbf{I}]^{-1} \mathbf{W}^t \mathbf{h} \quad (8.3.13)$$

Let \mathbf{W}^+ denote the Moore-Penrose pseudoinverse of \mathbf{W} . As is discussed in section 2.9, this is the unique generalized inverse of \mathbf{W} that leads to a minimum norm solution to the noise-free inverse problem; this minimum norm solution is the solution of minimum error variance if

$C(\mathbf{m})$ is \mathbf{I} and \mathbf{m} is multinormally distributed. It is also the solution to the noise-free inverse problem achieved through singular value decomposition. As previously discussed in section 2.9, Albert (1972) shows that

$$\mathbf{W}^+ = \lim_{\mu^2 \rightarrow 0} (\mathbf{W}^t \mathbf{W} + \mu^2 \mathbf{I})^{-1} \mathbf{W}^t = \lim_{\mu^2 \rightarrow 0} \mathbf{W}^t (\mathbf{W} \mathbf{W}^t + \mu^2 \mathbf{I})^{-1} \quad (8.3.14)$$

If the first of these equalities is compared to (8.3.13) it is apparent that, like singular value decomposition acting on a Kahanen-Loève transformed parameter set, a Tikhonov regularized solution to the inverse problem approaches the Moore-Penrose pseudoinverse as measurement noise approaches zero.

So is one of these forms of regularization just as good as the other? This has already been discussed to some extent. In fact singular value decomposition and Tikhonov regularization complement each other. The attraction of Tikhonov regularization is the flexibility that it offers in expressing expert knowledge. A \mathbf{Z}_r matrix, or a more elegant nonlinear expression of features that are desirable in a calibrated parameter field, constitutes a much more flexible expression of expert knowledge, and the uncertainties associated therewith, than pre-inversion Kahanen-Loève transformation. Flexibility of Tikhonov regularization is further enhanced by the ability of the Tikhonov-regularized inversion process to adjust weights associated with Tikhonov constraints as part of the inversion process itself. On the other hand, singular value decomposition guarantees unequivocal numerical stability of the inversion process. The same applies to more sophisticated iterative solvers which are related to singular value decomposition but are much faster and can handle very large numbers of parameters; see the next section.

The best regularization method to use is a combination of the two. The inverse problem can be formulated using Tikhonov regularization; the equations that emerge from this formulation can then be solved using singular value decomposition, or a fast iterative solver which operates along similar orthogonalization lines.

8.4 Practical Implementation

8.4.1 Nonlinear Model Behaviour

Most models are, of course, nonlinear. Hence they can only be temporarily represented by a matrix during each iteration of the inverse problem solution process. During each of these iterations, the equations presented above are used to calculate *improvements* to the values of parameters, rather than the calibrated values of these parameters. These improvements are calculated from current residuals rather than directly from the calibration dataset \mathbf{h} . During each iteration, the Jacobian matrix replaces the \mathbf{Z} matrix. Each column of this matrix is a vector of length n (the number of observations comprising the calibration dataset) whose elements represent the partial derivatives of model outputs used in the calibration process with respect to a particular parameter; see section 7.4. The regularization function represented by \mathbf{Z}_w in equations presented in the preceding section may be linear or nonlinear. If it is nonlinear, then \mathbf{Z}_w is also updated during each iteration of the inversion process with a new matrix of partial derivatives calculated using updated parameter values.

Strategies through which nonlinear model behaviour can be best accommodated during each linearized iteration of the regularized inversion process are similar to those described for manually-regularized, parsimonious model parameterization in chapter 7. They include logarithmic transformation of some parameters, limiting the length of the parameter upgrade vector, employment of the Broyden Jacobian update procedure, and use of the Marquardt

lambda.

Where a model is highly nonlinear, the possibility of convergence to a local objective function minimum can never be completely removed. However the probability of this happening can often be greatly reduced through creative formulation of a multi-component measurement objective function in which some components are more immune to model nonlinearities than are other components (see chapters 15 and 16). If these objective function components are supplemented with others that encapsulate regularization constraints that guide the inversion process towards a sensible parameter set that achieves a good fit with the measurement dataset, then the purpose of the model calibration process has been well served. As will be recalled from earlier chapters, the purpose of model calibration is to obtain a parameter set whose status approaches that of minimum error variance. This parameter set can then be used to make predictions of minimum error variance. It can also provide a foundation for subsequent exploration of parameter and predictive uncertainty. As has been discussed, calibration should not be an end in itself.

The Marquardt lambda can be just as useful in nonlinear, highly parameterized inversion as it is in nonlinear parsimonious inversion that is the outcome of manual regularization. At early stages of the nonlinear inversion process, parameters can be guided directly down the objective function gradient through use of a moderate to high value of the Marquardt lambda. As has been discussed, PEST implements a trial-and-error procedure for selection of an appropriate Marquardt lambda value during each iteration of the inversion process; linearized equations for parameter upgrades are solved using different values of the Marquardt lambda. Upgraded parameters are then tested for their efficacy in lowering the overall objective function by running the model using these parameters. The Marquardt lambda which results in greatest objective function improvement is retained. It then forms the starting value for a similar Marquardt lambda testing process that is implemented during the following iteration, once a new Jacobian matrix has been filled. While this procedure increases the number of model runs per iteration beyond that which is required for filling of the Jacobian matrix, it can be implemented efficiently in a parallel computing environment where it is just as cost-effective to attempt many parameter upgrades as it is to attempt just one. Ideally the value of the Marquardt lambda should fall with iteration count. If it does not fall, this may indicate that formulation of Tikhonov constraints needs revision as the Marquardt lambda may be adopting a de facto regularization role.

With the Jacobian matrix used in place of the \mathbf{Z} matrix, with the inverse problem formulated in terms of parameter upgrades, and with use made of the Marquardt lambda, equation (8.3.2b) becomes

$$\underline{\mathbf{k}} - \underline{\mathbf{k}}_0 = [\mathbf{J}^t \mathbf{Q} \mathbf{J} + \mu^2 \mathbf{J}^t_w \mathbf{Q}_r \mathbf{J}_w + \lambda \mathbf{I}]^{-1} [\mathbf{J}^t \mathbf{Q} \mathbf{r} + \mu^2 \mathbf{J}^t_w \mathbf{Q}_r \mathbf{r}_w] \quad (8.4.1)$$

where the vector \mathbf{r} contains model-to-measurement residuals and \mathbf{r}_w contains residuals pertaining to “observed” regularization constraints. Meanwhile equation (8.3.3) becomes

$$\underline{\mathbf{k}} - \underline{\mathbf{k}}_0 = [\mathbf{J}^t \mathbf{Q} \mathbf{J} + \mu^2 \mathbf{J}^t_w \mathbf{Q}_r \mathbf{J}_w + \lambda \mathbf{I}]^{-1} \mathbf{J}^t \mathbf{Q} \mathbf{r} \quad (8.4.2)$$

While PEST employs the above equations to calculate parameter upgrades, other implementations of Tikhonov regularization do not use the Marquardt lambda in this same way. Instead, similar numerical advantages to use of the Marquardt lambda in solving a nonlinear inverse problem can be attained through varying the strength of enforcement of Tikhonov constraints. The role of λ is thus combined with that of μ .

8.4.2 Solution of Equations

Equations (8.4.1) and (8.4.2) can be solved for upgraded parameters using standard numerical libraries. In practice, they should be solved using singular value decomposition or a related method. This is because it is not a forgone conclusion that the matrix that must be inverted in these equations can actually undergo inversion without numerical difficulty, for it may approach singularity in spite of the use of Tikhonov regularization. Near singular behaviour can occur if μ^2 and λ are small. This can happen when PEST lowers μ^2 in order to pursue a better fit with a measurement dataset, and when PEST's Marquardt lambda testing procedure suggests the use of low Marquardt lambda values.

When using singular value decomposition to invert matrices that appear in equations (8.4.1) and (8.4.2), singular value truncation should not be employed as a regularization device in the manner described in section (8.2.5). Minimization of parameter and predictive error variance is now the task of Tikhonov regularization. Instead, the role of singular value truncation is to prevent amplification of numerical noise rather than amplification of measurement noise. Experience suggests that singular value truncation is often appropriate at a singular value that is about 5×10^{-7} times the maximum singular value. This threshold may need to be increased where calculation of the Jacobian matrix is accompanied by considerable measurement noise. This can occur where Jacobian matrix elements are filled using finite parameter differences and a model's solver experiences convergence difficulties.

Other options exist for solution of equations such as (8.4.1) and (8.4.2) as matrices become large because of the number of parameters that are featured in the inverse problem. Iterative solvers that perform similarly to singular value decomposition are often used to solve these equations. Termination of the solution process after an appropriate number of iterations offers the same numerical benefits as singular value truncation. Numerical solvers such as LSQR (Paige and Saunders, 1982a; 1982b) and PROPACK (Larsen, 1998) are easily accessed through appropriate libraries. Solvers of this type are particularly appealing in contexts where a model can calculate partial derivatives itself using adjoint methods. On any call to the model, a solver may require only that $\mathbf{J}^t \mathbf{x}$ or $\mathbf{J} \mathbf{y}$ be evaluated, where \mathbf{x} and \mathbf{y} are vectors that the solver provides and \mathbf{J} is the Jacobian matrix. Each such call requires that only one model run or one adjoint model run be undertaken. Parameters employed by the model can be progressively updated as these calls are made. Furthermore, during any iteration of the nonlinear inversion process, such model calls may be limited to a number that is much smaller than the number of parameters requiring estimation.

Use of randomized SVD (Halko et al., 2011) can also increase matrix equation solution speeds, as can parallelization of standard SVD algorithms that are implemented in packages such as LAPACK.

Even with use of these advanced numerical solvers, construction of the $\mathbf{Z}^t \mathbf{Q} \mathbf{Z}$ matrix, and inversion of matrices that appear in equations (8.4.1) and (8.4.2), can still be time consuming where parameter numbers are very large. Fortunately, however, where the number of observations is less than the number of parameters (as is often the case), another option exists for formulation, and then solution, of these equations. As we shall see later in this book, this alternative formulation is extensively used where the inverse problem is solved using ensemble methods.

Consider the matrix identity embodied in equation (2.15.6) which is repeated here.

$$(\mathbf{B}^t \mathbf{D}^{-1} \mathbf{B} + \mathbf{A}^{-1})^{-1} \mathbf{B}^t \mathbf{D}^{-1} = \mathbf{A} \mathbf{B}^t (\mathbf{B} \mathbf{A} \mathbf{B}^t + \mathbf{D})^{-1} \quad (8.4.3)$$

Now let us assume that (as mentioned above) the role of the Marquardt lambda is subsumed into that of μ^2 . Furthermore, define

$$\mathbf{H} = [\mathbf{J}_w^t \mathbf{Q}_w \mathbf{J}_w]^{-1} \quad (8.4.4)$$

Where Tikhonov constraints comprise respect for preferred parameter values, which are also prior mean parameter values, this becomes (see equation 8.3.12)

$$\mathbf{H} = \mathbf{C}^{-1}(\mathbf{k}) \quad (8.4.5)$$

With removal of the λ term, the above definition of \mathbf{H} , and use of the transformation described by (8.4.3), equation (8.4.2) becomes

$$\underline{\mathbf{k}} - \underline{\mathbf{k}}_0 = \mathbf{C}(\mathbf{k}) \mathbf{J}^t [\mathbf{J} \mathbf{C}(\mathbf{k}) \mathbf{J}^t + \mu^2 \mathbf{Q}^{-1}]^{-1} \mathbf{r} \quad (8.4.6)$$

If (as is often attempted)

$$\mathbf{Q} = \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \quad (8.4.7)$$

equation (8.4.6) becomes

$$\underline{\mathbf{k}} - \underline{\mathbf{k}}_0 = \mathbf{C}(\mathbf{k}) \mathbf{J}^t [\mathbf{J} \mathbf{C}(\mathbf{k}) \mathbf{J}^t + \mu^2 \mathbf{C}(\boldsymbol{\varepsilon})]^{-1} \mathbf{r} \quad (8.4.8)$$

The matrix that requires inversion in equation (8.4.8) is of size $n \times n$, where n is the number of observations that are featured in the inversion process. On the other hand, equation (8.4.2) requires inversion of an $m \times m$ matrix, where m is the number of parameters that are featured in the inverse problem. Where the number of parameters exceeds the number of observations, equation (8.4.8) is numerically easier to use.

8.4.3 Multiple Regularization Weight Factors

As was mentioned above, while Tikhonov regularization provides mechanisms for sophisticated and flexible expressions of expert knowledge, numerical stability of the Tikhonov-regularized inversion process cannot always be guaranteed. To understand why this is so, let us suppose that we are estimating parameters for a spatial model whose domain is two-dimensional and covers a large area. Suppose also that measurements comprising the calibration dataset are concentrated in some parts of the model domain but are virtually absent from other parts. In those parts of the model domain where measurement density is high, the calibration dataset may support unique estimation of system properties. However in other parts of the model domain the local information content of the calibration dataset may be too low to warrant alteration of parameter values from their defaults. It follows that in parts of the model domain where information density is high there is no need for the regularization weight factor to be high, as the local information content of the calibration dataset may not need expert knowledge enhancement. At the same time, the weight factor may need to be relatively high in those other parts of the model domain where information density is low or non-existent in order to ensure that parameters which occupy these areas do not deviate from their preferred values. Unfortunately, a single weight factor may not be able to accomplish both of these tasks. The assignment of a low value to a single weight factor in order to achieve the level of fit with the calibration dataset that is requested through the user-specified target measurement objective function may undermine numerical stability in those parts of the model domain where the calibration dataset is under-represented and solution of the inverse problem is nonunique.

Similar problems can occur where model parameters represent different system properties. Information density within the calibration dataset as it pertains to one of these system property types may be high, while system properties of other types may be relatively uninformed by the calibration dataset. In this context, differential weighting of expert knowledge as it pertains to

different parameter types is warranted in order to ensure stability of an inverse problem that is relatively well-posed with respect to some parameters and relatively ill-posed with respect to others. Such differential relative weighting of Tikhonov constraints can be provided by the modeller in his/her filling of the \mathbf{Q}_r matrix of equations (8.3.2) and (8.3.3). However it is unlikely that the relativities that are thus provided will complement relativities in measurement dataset information content in such a way as to guarantee stability of the inversion process while at the same time allowing maximum transfer of information from the measurement dataset to those parameters which are estimable.

Conceptually, the best way to handle this problem is to allow the inversion process itself to assign differential regularization weight factors to different subsets of Tikhonov regularization. Equation (8.3.3) thus becomes

$$\underline{\mathbf{k}} = [\mathbf{Z}^t \mathbf{Q} \mathbf{Z} + \mu_1^2 \mathbf{Z}_{w1}^t \mathbf{Q}_{r1} \mathbf{Z}_{w1} + \mu_2^2 \mathbf{Z}_{w2}^t \mathbf{Q}_{r2} \mathbf{Z}_{w2} + \mu_3^2 \mathbf{Z}_{w3}^t \mathbf{Q}_{r3} \mathbf{Z}_{w3} + \dots]^{-1} \mathbf{Z}^t \mathbf{Q} \mathbf{h} \quad (8.4.9)$$

The problem of estimating a suitable value for one regularization weight multiplier μ^2 then becomes that of estimating suitable values for multiple weight multipliers μ_i^2 . There are a number of ways to go about this. A method employed by PEST that has proven effective in many calibration contexts is to estimate the ratios of the μ_i^2 first and then apply a global weight multiplier to all Tikhonov regularization in the manner described in previous sections. As usual, this global weight multiplier is chosen in order to achieve a user-specified target measurement objective function. One easy option for choosing relative regularization weight multipliers is to assign them values such that composite observation sensitivities pertaining to all rows of each \mathbf{Z}_{wi} sum to the same value after each of these matrices is multiplied by its respective μ_i^2 ; see equation (7.3.4) for definition of composite observation sensitivity. Thus expressions of expert knowledge embodied in the different \mathbf{Z}_{wi} all make an impact on the total regularization objective function, regardless of the number of rows in each of the \mathbf{Z}_{wi} matrices.

Another option is to assign higher μ_i^2 values to groupings of expert knowledge that pertain to parameters that are ill-informed by the calibration dataset \mathbf{h} . One measure of the information content of the calibration dataset with respect to a particular parameter is the composite sensitivity of that parameter as calculated from the \mathbf{Z} matrix; see equation (7.3.1). Meanwhile the parameters that feature in the expert knowledge to which a particular regularization weight factor μ_i^2 pertains can be gleaned from the \mathbf{Z}_{wi} matrix with which it is associated. Those \mathbf{Z}_{wi} which feature parameters whose average or maximum composite sensitivity is low receive a higher μ_i^2 than those which feature parameters that are well informed by the calibration dataset.

An even more sophisticated attempt to weight regularization that pertains to inestimable parameters more highly than that which pertains to estimable parameters is to examine the projection of each row of each \mathbf{Z}_{wi} matrix onto the null space of the \mathbf{Z} matrix. Those which have a higher null space projection receive a greater relative regularization weight.

For all of these options, new values for μ_i^2 are calculated during every iteration of the inversion process. See PEST documentation for details.

8.5 Learning from History-Matching

When undertaking highly parameterized inversion, the attainment of a good fit between model outputs and field measurements may not be too difficult. However just because a high level of fit with a measurement dataset *can* be attained, this does not necessarily mean that it *should* be attained. The adverse consequences of over-fitting are dramatically illustrated by figure 8.4 and by equations that support it.

If using only singular value decomposition as a regularization device, over-fitting can be prevented through suitable singular value truncation. If using Tikhonov regularization, the target measurement objective function becomes the mechanism through which a modeller can prevent over-fitting. Presumably, the level of fit which is sought through either of these mechanisms should be no greater than that which is expected on the basis of measurement noise.

However, as has been mentioned many times in this text (and as will be discussed from a theoretical perspective in chapter 15), the level of model-to-measurement fit that *should* be sought through history-matching is not usually determined by measurement noise. It normally reflects, more than anything else, the adequacy, or otherwise, of the model as a simulator of real-world behaviour. It follows that the level of fit that should be sought in any particular decision-support modelling context is not known in advance. Instead it must be “discovered” through the inverse problem solution process itself.

As is discussed in chapter 7, where parameter uniqueness is achieved through manual regularization, the inverse problem is solved by minimizing an objective function whose value rises with model-to-measurement misfit. If parameter parsimonization achieved through manual regularization is sufficient to achieve inverse problem uniqueness, the objective function possesses a single, discrete minimum value. The value of this minimum reflects the amount of measurement and structural noise that is associated with a calibration dataset; as is discussed in section 7, some of this structural noise may have originated from manual regularization itself. Using equation (7.2.23) a reference variance σ_r^2 can be calculated from the minimized value of the objective function. This allows a modeller to quantify the actual amount of noise that is associated with a calibration dataset. Meanwhile, overfitting is deemed to occur if parameter values that are achieved through this process are unrealistic. If this is the case, then manual regularization must be repeated; fewer adjustable parameters must be defined, or the same number of parameters must be defined differently.

Where regularization is mathematical rather than manual, the objective function is not minimized. Instead, the desired level of model-to-measurement fit must be specified by the user. While selection of this desired value may be assisted by the availability of an L-curve, more often than not its specification is ad hoc (and a matter of some subjectivity) as construction of an L-curve is too numerically intensive.

In order to make the ensuing discussion a little easier follow, let us assume that Tikhonov regularization of the “preferred value” type is implemented. Solution to the inverse problem is achieved using the following equation

$$\underline{\mathbf{k}} = (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mu^2 \mathbf{C}^{-1}(\mathbf{k}))^{-1} \mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{h} \quad (8.5.1)$$

However, conclusions that are drawn below apply to other implementations of Tikhonov regularization as well.

In equation (8.5.1) parameters are represented as departures from their prior mean values. Prior means of parameters defined in this way are therefore zero. Note that equation (8.5.1) is the same as equation (8.3.11); however the weight matrix, i.e. $\mathbf{C}^{-1}(\boldsymbol{\varepsilon})$, is expressed explicitly rather than being incorporated into the \mathbf{Z} matrix. Obviously, weighting is in accordance with measurement noise.

Equation (8.5.1) is the solution of an objective function minimization problem. The objective function that is minimized is

$$\Phi = (\mathbf{h} - \mathbf{Z}\mathbf{k})^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) (\mathbf{h} - \mathbf{Z}\mathbf{k}) + \mu^2 \mathbf{k}^t \mathbf{C}^{-1}(\mathbf{k}) \mathbf{k} \quad (8.5.2)$$

If $C(\mathbf{k})$ and $C(\boldsymbol{\varepsilon})$ are perfectly known (which assumes that the real world and measurement/structural noise are both simple enough to be statistically characterized using covariance matrices), then μ^2 can be set to 1.0. Under these circumstances, Φ possesses a unique minimum. Parameters $\underline{\mathbf{k}}$ obtained through minimization of Φ are minimum error variance estimates of the true \mathbf{k} . If the model is linear, they are also posterior mean parameters that would be derived using Bayes equation. (This is why equation 8.5.1 is sometimes characterized as implementing “Bayesian Tikhonov” regularization.)

In practice, neither $C(\boldsymbol{\varepsilon})$ nor $C(\mathbf{k})$ is known. However we normally consider these matrices to be correct in a relative sense. Hence equations (8.5.1) and (8.5.2) are still applicable, but a solution must be found not just for $\underline{\mathbf{k}}$, but for μ^2 as well. μ^2 can be regarded as a “balancing parameter”. As such, it is used to trade off respect for a measurement dataset against respect for prior mean parameter values. Inclusion of $C(\boldsymbol{\varepsilon})$ and $C(\mathbf{k})$ in these equations attempts to ensure that departures of model outputs from field measurements respect the temporal/spatial characteristics of measurement noise, while departures of estimated parameter values from their prior means respect the temporal/spatial correlation structure of real-world hydraulic properties.

It goes without saying that if we are unsure of μ^2 , then we are probably unsure of all of the elements of $C(\mathbf{k})$ and $C(\boldsymbol{\varepsilon})$. However we adopt these matrices anyway for lack of a better alternative.

So what is the best way to determine an optimal value for μ^2 under these circumstances? If μ^2 is too low, then the matrix appearing in (8.5.1) may approach singularity. This onset of matrix noninvertibility therefore sets a lower bound on μ^2 . Perhaps we can solve equation (8.5.1) many times using many different values of μ^2 that are above this lower limit in order to plot the L curve of figure 8.5. This will work, but is numerically expensive.

A numerically cheaper option is to set different values for the target measurement objective function and see what we can learn from PEST. As is discussed in section 8.3.3, PEST calculates μ^2 from this target. It is often useful to commence a model calibration process by setting the target measurement objective function to an unachievably low value. PEST then obtains as good a fit as the model allows. Behind the scenes, PEST does not allow μ^2 to fall to zero as it achieves the best level of model-to-measurement fit that it can. Instead, as it implements the regularized inversion process, it employs a series of temporary target measurement objective functions that change from iteration to iteration. These are all somewhat lower than the current measurement objective function (the user decides how much lower). Hence, in obtaining the best level of model-to-measurement fit that the model allows, emergent parameter fields are not chaotic, for they obtain the benefits of regularization. If regularization is implemented using equations (8.5.1) and (8.5.2), estimated values of $\underline{\mathbf{k}}$ therefore depart from prior mean values of \mathbf{k} only to the smallest extent required to achieve this level of fit. While these departures are probably too great to be tolerable, a modeller is nevertheless made aware of the best fit that the model can attain with available data, and of the “parameter unreality” cost that is required to achieve this fit. If this fit is not good enough, the modeller is awakened to the fact that his/her numerical model is in need of improvement. Meanwhile, the nature of these improvements may be suggested by the nature of model-to-measurement misfit, and by the nature of parameter unreality that emerged during solution of the inverse problem.

Alternatively, a modeller may find that he/she is satisfied with his/her model, but not with the parameter field that has emerged through pursuit of the lowest possible objective function. Figure 8.4 suggests that the target measurement objective function may not need to be raised too much higher than the best attainable objective function in order to attain an acceptable fit

with a measurement dataset with a much more pleasing parameter field. A new target measurement objective function is easily tested by repeating the inversion process. Generally, successful inversion outcomes can be attained with only a 5% to 10% increase in the target measurement objective function above the best that the model can achieve. Meanwhile, because the modeller has initiated the model calibration process by pursuing the best possible fit with the calibration dataset that the model allows, he/she is in a position to justify the level of fit that he/she eventually decides to accept by pointing out to possible critics of the calibrated model that a better fit is indeed attainable, but that the cost of attaining this fit is parameter field unacceptability.

By foregoing model-to-measurement fit in order to gain parameter acceptability, a modeller recognizes that some or all elements of $C(\epsilon)$ are larger than expected. This has repercussions for post-calibration uncertainty/error analysis. A higher-than-expected level of noise associated with a calibration dataset reduces access of the history-matching process to information that is resident in that dataset, whether this noise is “real” or structural. So something has been learned from the regularized inversion process.

There may be occasions however, where regularized inversion exposes the fact that current assumptions pertaining to $C(k)$ should be questioned. If PEST is determined to introduce anomalous hydraulic properties to certain parts of a model domain despite the fact that regularization seeks minimum departures from prior parameter means, there can only be two reasons for this. The first is that aberrant parameters are hosting calibration-acquired information that is inappropriate for them; that is, these parameters are adopting roles that compensate for model imperfections. The second reason is that the emergent heterogeneity is real. While the shapes that this heterogeneity adopts in the calibrated parameter field may not be the same as those of causative geological features (because regularization was not tuned to exposure of these features), the emergence of anomalous parameter values may provide unequivocal evidence of the existence of these features.

Acceptance of the second of these reasons should precipitate a review of the current $C(k)$. There are two reasons for this.

Firstly, if history-matching has just exposed the inadequacy of the current $C(k)$ at one or a number of locations where measurements of system behaviour have had the opportunity to contradict it, then it may be inadequate at other locations where field measurements are sparse. If the model is required to make management-pertinent predictions at the latter locations, then the integrity of uncertainty assessment of those predictions requires a revised $C(k)$.

Secondly, if one aspect of $C(k)$ (for example the magnitude of one or more of its diagonal elements) has been exposed as inadequate, then it is probable that other aspects of $C(k)$ (for example correlations expressed by off-diagonal elements) are also inadequate. In most real-world instances of model calibration, history-match emergent patterns of anomalous parameter values do not reflect the true shapes of causative geological features because the current $C(k)$ matrix dictates that heterogeneity adopt other patterns. In many geological circumstances, anomalous features are discrete and elongate. Where model calibration exposes the presence of anomalous hydraulic properties, a modeller must then decide whether to attempt to represent the feature using a parameterizable geobody, or by introducing spatially varying anisotropy to his/her current $C(k)$ matrix. Regularized inversion that is conducted using either of these alternatives will probably result in significant alterations to patterns of emergent heterogeneity. If these new patterns are more geologically realistic, the potential for bias in at least some model predictions will have been reduced. At the same time, the integrity of post-calibration uncertainty analysis will have been raised. Once again, something has been learned from the

regularized inversion process.

8.6 Hierarchical Inversion

Use of equations such as (8.5.1) and (8.5.2) that include $C(\mathbf{k})$ and $C(\boldsymbol{\varepsilon})$ assumes, of course, that a modeller can provide these matrices, despite the problems that are associated with their formulation in real-world modelling contexts. Selection of an appropriate target measurement objective function prior to implementation of regularized inversion implies an appropriate value for μ^2 . This, in turn, implies factors that must be applied to $C(\boldsymbol{\varepsilon})$ and $C(\mathbf{k})$ that are learned through the inversion process. However, relativity of elements within each of the $C(\boldsymbol{\varepsilon})$ and $C(\mathbf{k})$ matrices is assumed to be correct.

The previous section discusses how outcomes of an inversion process may induce a modeller to question temporal and spatial correlations that are implied in off-diagonal elements of $C(\boldsymbol{\varepsilon})$ and $C(\mathbf{k})$. In practice, $C(\boldsymbol{\varepsilon})$ is normally assumed to be diagonal, even though model-to-measurement misfit is generally dominated by structural rather than measurement noise. As is discussed in chapter 15, the presence of structural noise can often be more intuitively accommodated by formulation of a multi-component objective function than by introduction of non-diagonal elements to $C(\boldsymbol{\varepsilon})$. The same does not apply to $C(\mathbf{k})$.

There are few, if any, groundwater modelling circumstances where spatial correlations that are required for construction of a $C(\mathbf{k})$ matrix are either known with any degree of certainty, or are even applicable. Media through which groundwater flows are generally too heterogeneous for such a simplistic description, even with abandonment of the assumption of geostatistical stationarity. It follows that hyperparameters that govern construction of $C(\mathbf{k})$ (see chapter 4) should be recognized as uncertain, and hence accommodated in prior and posterior parameter and predictive uncertainty analysis. What is relevant in the present chapter, however, are repercussions of their uncertainties for regularized inversion.

So let us suppose that a calibration dataset hosts information on correlation lengths and directions of hydraulic properties in different parts of a model domain. If this is the case, then the model calibration process should be formulated in such a way as to harvest this information. Heuristic harvesting of this information, in which manual adjustments are made to $C(\mathbf{k})$ based on calibration outcomes, are discussed in the previous section. We now discuss how this process can be automated. Automation can be achieved through adoption of a so-called “noncentred” model parameterization scheme; see section 4.2.

Suppose, for the moment, that we believe that spatial variability of subsurface hydraulic properties (and hence model parameters that represent these properties) can be described by a covariance matrix $C(\mathbf{k})$ which we acknowledge to be uncertain. First we imbue the inversion process with a set of parameters \mathbf{z} from which model hydraulic properties can be calculated using equation (4.2.2) which is repeated here as (8.6.1).

$$\mathbf{k} = \underline{\mathbf{k}} + C^{\frac{1}{2}}(\mathbf{k})\mathbf{z} \quad (8.6.1)$$

In this equation, $\underline{\mathbf{k}}$ is the mean value of \mathbf{k} . As is explained in chapter 4, use of equation (8.6.1) does not require an assumption of geostatistical stationarity. $\underline{\mathbf{k}}$ can exhibit spatial variability. Spatial correlation lengths encapsulated in $C(\mathbf{k})$ can also exhibit spatial variability. What is further explained in chapter 4 is that the prior mean of \mathbf{z} is $\mathbf{0}$, and that

$$C(\mathbf{z}) = \mathbf{I} \quad (8.6.2)$$

regardless of $C(\mathbf{k})$. That is, the elements of \mathbf{z} are independent standard normal variates (often referred to as “iids”).

Equation (8.6.2) provides us with a covariance matrix to use in equations such as (8.5.1) through which regularized inversion can be implemented to obtain minimum error variance estimates $\underline{\mathbf{z}}$ of \mathbf{z} . Meanwhile, geostatistical hyperparameters that govern the spatial distribution of $\underline{\mathbf{k}}$ and correlation lengths implied by $C(\mathbf{k})$ can also be featured in the inversion process. These hyperparameters have their own means and covariance matrices. These should be included in formulation of the inverse problem through which model calibration is achieved so that minimum error variance estimates of these hyperparameters can also be obtained.

Let us use the letter \mathbf{g} to denote geostatistical hyperparameters from which the elements of $\underline{\mathbf{k}}$ and $C(\mathbf{k})$ are calculated. We assume that elements of \mathbf{g} are defined as departures from prior mean values so that the prior mean of \mathbf{g} is $\mathbf{0}$. Meanwhile we denote the prior covariance matrix of \mathbf{g} as $C(\mathbf{g})$. Then the linearized, hierarchical inverse problem can be formulated as

$$\mathbf{h} = \mathbf{M}\mathbf{z} + \mathbf{N}\mathbf{g} + \boldsymbol{\varepsilon} \quad (8.6.3)$$

where, from (8.6.1),

$$\mathbf{M} = \mathbf{Z}C^{1/2}(\mathbf{k}) \quad (8.6.4)$$

and \mathbf{N} embodies sensitivities of model outputs to hyperparameters \mathbf{g} . This equation can be written as

$$\mathbf{h} = [\mathbf{M} \quad \mathbf{N}] \begin{bmatrix} \mathbf{z} \\ \mathbf{g} \end{bmatrix} + \boldsymbol{\varepsilon} \quad (8.6.5)$$

Once in the form of (8.6.5), and recalling equation (8.6.2), the solution to the hierarchical inverse problem is easily formulated as (see 8.5.1)

$$\begin{bmatrix} \mathbf{z} \\ \mathbf{g} \end{bmatrix} = \left([\mathbf{M}^t \quad \mathbf{N}^t] C^{-1}(\boldsymbol{\varepsilon}) \begin{bmatrix} \mathbf{M} \\ \mathbf{N} \end{bmatrix} + \mu^2 \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & C^{-1}(\mathbf{g}) \end{bmatrix} \right)^{-1} [\mathbf{M}^t \quad \mathbf{N}^t] C^{-1}(\boldsymbol{\varepsilon}) \mathbf{h} \quad (8.6.6)$$

The fact that \mathbf{M} is a function of \mathbf{g} because of the dependence of $C(\mathbf{k})$ on \mathbf{g} does not affect problem linearization that is embodied in the above equations. To be sure, sensitivities encapsulated in \mathbf{M} and \mathbf{N} change from iteration to iteration. But this is standard for nonlinear parameter estimation; the elements of \mathbf{M} change anyway because of model nonlinearity. What is most important is that neither $C(\mathbf{z})$ nor $C(\mathbf{g})$ change as the values of \mathbf{z} and \mathbf{g} change; $C(\mathbf{z})$ remains the identity matrix \mathbf{I} . This makes hierarchical regularized inversion possible.

9. Dimensional Reduction Methods

9.1 General

Previous chapters of this book have pointed out the benefits of using many parameters when undertaking environmental model history-matching. Highly parameterized inversion is essential for attainment of a solution of minimized error variance to the inverse problem that model calibration attempts to solve. Deployment of many parameters is also essential for the integrity of parameter and predictive uncertainty analysis.

This chapter begins in the same way as the previous chapter. It starts with an equation that presents a linearized representation of the action of a model \mathbf{Z} on a set of parameters \mathbf{k} . We assume that \mathbf{k} has as many elements as are required to represent prediction-salient aspects of an environmental system, and to harvest information that is salient to those predictions from measurements of system behaviour.

$$\mathbf{h} = \mathbf{Zk} + \boldsymbol{\varepsilon} \quad (9.1.1)$$

To maintain simplicity of equations that follow, we continue to define parameters \mathbf{k} as departures from their prior means.

The previous chapter discusses ways in which an inverse problem can be formulated, and then solved, in order to yield minimum error variance estimates $\underline{\mathbf{k}}$ of \mathbf{k} . Regularization that is necessary to achieve these estimates is pursued using singular value decomposition and Tikhonov methods.

Where parameter numbers are large, matrices appearing in equations that are derived in the previous chapter become large. Manipulation of these matrices becomes numerically costly. Inversion of these matrices becomes even more costly, as does implementation of singular value decomposition. In some circumstances, numerical savings can be made by using alternative formulations of these equations; see section 8.4.2. However in the groundwater modelling context, by far the largest numerical burden is incurred through construction of the sensitivity (i.e. \mathbf{Z}) matrix. Generally elements of this matrix are filled using finite parameter differences. This requires that the model be run at least once for each parameter. Where model run times are large this can become exceedingly costly, even where model runs are parallelized. Because \mathbf{Z} must be reconstructed during every iteration of a nonlinear inversion process, the numerical cost of inversion can become prohibitive.

This chapter discusses two methods through which the numerical burden of estimating many parameters can be reduced. Both of these methods achieve numerical benefits by reducing the dimensionality of the space in which an inverse problem is solved. Matrices therefore become smaller. Most importantly, the numerical burden of filling the sensitivity matrix is reduced.

The theory behind one of the methods that is discussed in the present chapter has been presented already; see section 8.2.7. Here it is briefly reviewed. The other method makes use of ensembles of parameters. Later chapters of this book discuss how ensemble methods can be used to explore parameter and predictive uncertainty. These methods also benefit from working in a reduced dimensional subspace of parameter space.

9.2 The Principle

We start with equation (7.1.2). This is the same equation that was used to begin our discussion of manual regularization; however we alter the name of the \mathbf{L} matrix to \mathbf{M} for notational

convenience of equations that follow.

$$\mathbf{k} = \mathbf{M}\mathbf{p} \quad (9.2.1)$$

The matrix \mathbf{M} has as many rows as the number of elements that comprise the parameter vector \mathbf{k} ; as usual, we characterize this number using the integer m . \mathbf{M} has as many columns as the number of elements that comprise a smaller parameter set \mathbf{p} ; we symbolize the number of elements of \mathbf{p} using the integer r (for “reduced”). Presumably r is less than m ; hence \mathbf{M} has fewer columns than rows. We seek numerical efficiencies by working in \mathbf{p} -space rather than \mathbf{k} -space. This is done when regularization is manual. However the approach to parameter reduction that is pursued in the present chapter is different from that of chapter 7 in two ways. First, we do not seek to formulate a well-posed inverse problem, for regularization will still be deployed when working in a reduced dimensional parameter space. Second, estimated parameters \mathbf{p} are defined using less primitive devices than zones of piecewise constancy. Ideally, the parameter field that emerges from the inversion process will therefore “look good” from a hydrogeological point of view. That is, it will be more reflective of $C(\mathbf{k})$.

Equation (9.2.1) states that a vector \mathbf{k} with many elements can be built from a vector \mathbf{p} with few elements. To see how this happens, let us represent the matrix \mathbf{M} using vectors that comprise its columns, so that

$$\mathbf{M} = [\mathbf{m}_1 \quad \mathbf{m}_2 \quad \dots \quad \mathbf{m}_r] \quad (9.2.2)$$

Equation (9.2.1) states that

$$\mathbf{k} = p_1\mathbf{m}_1 + p_2\mathbf{m}_2 + \dots + p_r\mathbf{m}_r \quad (9.2.3)$$

That is, the vector \mathbf{k} is a linear combination of the columns of \mathbf{M} . Presumably, these columns are linearly independent so that they collectively span an r -dimensional space. Nevertheless, unless r is equal to m , not every \mathbf{k} can be expressed in terms of a complementary \mathbf{p} . Hence use of \mathbf{p} in place of \mathbf{k} cannot capture all of the details that can possibly exist in \mathbf{k} -space. Ideally, however, use of \mathbf{p} instead of \mathbf{k} should allow representation of enough \mathbf{k} detail for a model to be capable of replicating field measurements of system behaviour. The elements of \mathbf{p} must therefore be sufficient in number to harvest the information content of a calibration dataset. At the same time, predictions of management interest that are made using this reduced-dimensional calibrated parameter set should suffer as little bias as possible despite \mathbf{p} ’s inability to represent all of the nuances of \mathbf{k} .

In PEST parlance, the elements of \mathbf{p} are referred to as “super parameters”.

Notice that the above conditions on \mathbf{p} do not include the ability to properly characterize the full range of uncertainty of decision-critical predictions. For the moment it is assumed that, while model calibration takes place in a reduced dimensional subspace, it does not follow that exploration of predictive uncertainty will be confined to this same subspace. This is in contrast to ensemble-based implementations of Bayes equation (described later in this book) wherein history-matching and uncertainty analysis do, indeed, take place in the same subset of parameter space.

9.3 SVD-Assist

9.3.1 How SVD-Assist Works

For the SVD-assist methodology that is supported by both PEST and PESTPP-GLM, the \mathbf{M} vector in the above equations is the \mathbf{V}_1 vector that emerges from singular value decomposition of the \mathbf{Z} matrix. The number of columns of this matrix determines the subspace of parameter

space in which parameter estimation takes place. The dimensions of this subspace should exceed that of the inverse problem solution space. That is, it should extend beyond the optimal singular value truncation point that is discussed in section 8.2.5 and depicted in figure 8.4. (In theory, the number of \mathbf{V}_1 matrix columns need only extend to the optimal singular value truncation point. However, this point is not known in advance of the inversion process; so it is best to use more columns than this, thereby estimating more super-parameters than are actually needed to fit a calibration dataset.)

Once the original Jacobian matrix has been calculated in \mathbf{k} -space, the cost of SVD-assisted inversion is relatively cheap. Filling of the \mathbf{M} matrix using finite parameter differences requires that only r model runs be carried out. This is the number of columns of \mathbf{V}_1 , and hence the number of super parameters that require estimation.

It is apparent that SVD-assist is only partly a dimensional reduction method. Because filling of the \mathbf{M} matrix relies on singular value decomposition of \mathbf{Z} , a \mathbf{Z} matrix must first be filled. This is generally undertaken before commencement of the inversion process using finite difference perturbations about parameter prior mean values. There can be no guarantee, however, that separation of parameter space into solution and null spaces based on singular value decomposition of a \mathbf{Z} matrix calculated in this way remains valid as parameters change; this only happens if a model is linear. Nevertheless, as long as the vectors that comprise the columns of \mathbf{V}_1 continue to span the evolving solution space, this does not compromise the ability of the inversion process to find a set of super parameters for which model outputs fit field measurements well. However bias may be introduced to some model predictions as some vectors that comprise this \mathbf{V}_1 matrix may possess non-zero projections onto the evolving (but unknown) null space of parameter space.

In general, protection against both model-to-measurement misfit and predictive bias engendered by model nonlinearity can be gained by estimating as many super parameters as computing resources allow. Further protection can be achieved by recomputing the full \mathbf{Z} matrix after every few iterations of the SVD-assisted inversion process.

9.3.2 Regularization for SVD-Assisted Inversion

It was stated in section 8.2.4 that singular value decomposition of a Jacobian matrix should ideally be preceded by Kahunen-Loëve transformation of model parameters, for this reduces the propensity for calibration-induced parameter/predictive bias. However this is not done in PEST's implementation of SVD-assisted inversion. Instead, bias reduction is pursued through application of Tikhonov regularization.

PEST's implementation of the SVD-assist methodology assumes that a modeller has created a PEST input dataset in \mathbf{k} -space. It is also assumed that formulation of the \mathbf{k} -space inverse problem includes Tikhonov regularization that is applied in this same space. The SVDAPREP program used by PEST to build a reformulated inverse problem in \mathbf{p} -space retains all \mathbf{k} -space regularization constraints.

To see how this works, let us consider that formulation of the inverse problem in \mathbf{k} -space is such that its solution would proceed according to equation (8.5.1) if undertaken in that space. Solution of the inverse problem using SVD-assisted inversion employs a modified equation. First, we note that, from (9.1.1) and (9.2.1) with \mathbf{V}_1 substituted for \mathbf{M}

$$\mathbf{h} = \mathbf{ZV}_1\mathbf{p} + \boldsymbol{\varepsilon} = \mathbf{Xp} + \boldsymbol{\varepsilon} \quad (9.3.1)$$

By direct substitution, equation (8.5.1) is then modified as follows

$$\underline{\mathbf{k}} = \mathbf{V}_1 \mathbf{p} = \mathbf{V}_1 (\mathbf{X}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{X} + \mu^2 \mathbf{p}^t \mathbf{V}_1 \mathbf{C}^{-1}(\mathbf{k}) \mathbf{V}_1 \mathbf{p})^{-1} \mathbf{X}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{h} \quad (9.3.2)$$

9.4 Ensemble Space Inversion

9.4.1 How ENSI Works

Ensemble space inversion (ENSI) employs a regularized, ensemble adjustment algorithm that bears some similarity to dimensional reduction schemes that are discussed by Iglesias et al (2013) and Chada et al (2018). However, it possesses some novel options that are not described by these authors. These facilitate ENSI's solution of highly nonlinear inverse problems that may challenge other inversion methods.

The subspace of parameter space in which ENSI operates is defined by random realizations of parameters drawn from the prior parameter probability distribution. Because of this, ENSI bears a strong resemblance to Bayesian ensemble methods that are discussed later in this book. However ENSI offers greater flexibility in definition of the ensemble-defined subspace in which solution of an inverse problem is sought than most other ensemble methods. Also, it enables some model parameters (so-called “nonrealization parameters”) to retain their own individual identities. They are not therefore tethered to other parameters in immutable ratios defined by an initial set of realizations. Experience demonstrates that strategic selection of nonrealization parameters can enhance optimization efficiency and overcome obstacles to optimal parameter adjustment posed by inverse problem nonlinearity. ENSI is particularly well suited to hierarchical inversion; see Kitlasten et al (2025) for an example. Geostatistical hyperparameters can be separated from other parameters in their own subspace, or can be granted nonrealization status.

ENSI is now explained in greater detail. As usual, model parameters are designated by the vector \mathbf{k} .

To begin the ENSI process, \mathbf{k} is subdivided into $N+1$ different groups so that

$$\mathbf{k} = \begin{bmatrix} \mathbf{k}_0 \\ \mathbf{k}_1 \\ \mathbf{k}_2 \\ \vdots \\ \mathbf{k}_N \end{bmatrix} \quad (9.4.1)$$

Definition of these groups must be such that prior parameter correlations do not cross parameter group divides. The prior covariance matrix of \mathbf{k} can therefore be written as

$$\mathbf{C}(\mathbf{k}) = \begin{bmatrix} \mathbf{C}_0(\mathbf{k}_0) & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{C}_1(\mathbf{k}_1) & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{C}_2(\mathbf{k}_2) & \dots & \mathbf{0} \\ \vdots & \dots & \dots & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{C}_N(\mathbf{k}_N) \end{bmatrix} \quad (9.4.2)$$

Based on this $\mathbf{C}(\mathbf{k})$ matrix, realizations of parameters can now be generated. It is important to note, however, that

- realizations are generated separately for each group of parameters using their own covariance matrices;
- these are centred on initial parameter values provided in a PEST control file;
- different numbers of realizations can be generated for each parameter group (referred to as “realization groups”);

- no realizations are generated for parameter group 0.

Let r_i denote the number of realizations pertaining to realization group i . These are placed into the columns of an $m_i \times r_i$ matrix \mathbf{M}_i , where m_i is the number of parameters that belong to this realization group. Then, for each group, we define r_i factors encapsulated in a vector \mathbf{f}_i . For group i , equation (9.2.1) then becomes

$$\mathbf{k}_i = \mathbf{M}_i \mathbf{f}_i \quad (9.4.3)$$

That is, the factors \mathbf{f}_i are applied to the columns of \mathbf{M}_i to form \mathbf{k}_i .

Elements of the \mathbf{f}_i vectors are our new parameters, for the task of the inversion process is to estimate values for these factors, as well as values for individual members of the \mathbf{k}_0 vector. Because of the way in which parameters \mathbf{k} are defined (i.e. as perturbations from their prior mean values), prior mean values of all of the elements of all \mathbf{f}_i are 0. (Actually, internally to ENSI, parameters are always defined as perturbations about their empirical means; hence prior mean values of all \mathbf{f}_i are indeed $\mathbf{0}$.)

Estimation of these new \mathbf{f}_i parameters using regularized inversion requires knowledge of their prior covariance matrices. Before deriving $C(\mathbf{f}_i)$ (the prior covariance matrix of \mathbf{f}_i), we first note that, because the columns of \mathbf{M}_i are realizations of \mathbf{k}_i parameters, an empirical covariance matrix of \mathbf{k}_i can be calculated from \mathbf{M}_i as

$$C(\mathbf{k}_i) = \frac{\mathbf{M}_i \mathbf{M}_i^t}{r_i - 1} \quad (9.4.4)$$

(Recall that, for ease of discussion, parameter means are assumed to be zero.)

However, from (9.4.3) and (3.9.2)

$$C(\mathbf{k}_i) = \mathbf{M}_i C(\mathbf{f}_i) \mathbf{M}_i^t \quad (9.4.5)$$

Comparing (9.4.5) with (9.4.4) it is apparent that

$$C(\mathbf{f}_i) = \frac{\mathbf{I}}{r_i - 1} \quad (9.4.6)$$

where \mathbf{I} is the identity matrix. Each element \mathbf{f}_{ij} of \mathbf{f}_i therefore has a standard deviation that is given by

$$\sigma_{f_{ij}} = \frac{1}{\sqrt{r_i - 1}} \quad (9.4.7)$$

As stated above, no realizations are generated for \mathbf{k}_0 parameters, for these are estimated individually; they therefore retain their \mathbf{k} -space identity. Hence, to calculate \mathbf{k} from \mathbf{k}_0 and a set of \mathbf{f}_i vectors (these collectively defining the subspace in which inversion takes place), the following equation is employed.

$$\mathbf{k} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{M}_1 & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{M}_2 & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{M}_N \end{bmatrix} \begin{bmatrix} \mathbf{k}_0 \\ \mathbf{f}_1 \\ \mathbf{f}_2 \\ \vdots \\ \mathbf{f}_N \end{bmatrix} \quad (9.4.8)$$

We repeat that the number of elements in each \mathbf{f}_i can be different; this is the modeller's choice. Equation (9.1.1) which is used to define the inverse problem that must be solved now becomes

$$\mathbf{h} = \mathbf{Z} \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{M}_1 & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{M}_2 & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{M}_N \end{bmatrix} \begin{bmatrix} \mathbf{k}_0 \\ \mathbf{f}_1 \\ \mathbf{f}_2 \\ \vdots \\ \mathbf{f}_N \end{bmatrix} + \boldsymbol{\varepsilon} \quad (9.4.9)$$

Solution of the inverse problem for minimum error variance estimates $\mathbf{k}_0, \mathbf{f}_1, \mathbf{f}_2 \dots \mathbf{f}_N$ of $\mathbf{k}_0, \mathbf{f}_1, \mathbf{f}_2 \dots \mathbf{f}_N$ can now employ Bayesian Tikhonov regularized inversion in the usual way. Once the inverse problem has been solved, \mathbf{k} can be calculated from these quantities using (9.4.8).

9.4.2 Regularization for ENSI

As is discussed in section (9.3.2), PEST's implementation of SVD-assisted inversion is such that regularization is applied to \mathbf{k} parameters, despite the fact that the inverse problem is solved in \mathbf{p} -space. The same does not apply to PEST's implementation of ENSI. In this case, \mathbf{k}_0 and \mathbf{f}_i parameters are subjected to preferred-value regularization. The prior covariance matrix of \mathbf{k}_0 parameters is supplied by the modeller in the usual fashion, for these are native model parameters. Prior covariance matrices for \mathbf{f}_i parameters are given by equation (9.4.6).

9.4.3 Strengths and Weaknesses

9.4.3.1 Model Run Efficiency

The number of model runs that is required for ENSI-based inversion depends on the number of “super parameters” that are estimated. As for normal inversion, during each iteration of the ENSI process, at least one model run must be conducted for each parameter in order to calculate sensitivities of all model outputs with respect to this parameter. The number of ENSI parameters is equal to the number of elements in \mathbf{k}_0 plus the sum of the number of elements in each of the \mathbf{f}_i vectors. The number of elements in an \mathbf{f}_i vector is equivalent to the number of realizations that occupy the columns of the corresponding \mathbf{M}_i matrix.

When designing an ENSI inversion process, it is necessary to first subdivide parameters into “realization groups”, and then to define the number of realizations that is associated with each group; nonrealization parameters are also specified. By doing all of this, the total number of parameters that require estimation is identified. In deciding on the acceptability of this number, the same rules apply as for any other inversion method. The number of parameters must exceed the dimensionality of the inverse problem solution space. If this is not the case, there is an insufficiency of receptacles for information that is resident in the calibration dataset. This has two repercussions. The first is that it will not be possible to attain a good fit between pertinent model outputs and the calibration dataset. The second is calibration-induced distortion of parameter values as they adopt compensatory roles for parameter insufficiency. This can bias model predictions.

Of course, the number of dimensions of the calibration solution space is not generally known in advance of an inversion process. However once a super-parameter sensitivity matrix has been filled, it can be estimated using principles that are described in chapter 8 of this book and implemented in PEST utility software; see chapter 10. Alternatively, failure to achieve a good fit with a calibration dataset, and/or estimation of parameter values that lack credibility, can signify that the inversion process is attempting to estimate too few parameters.

ENSI has an advantage over SVD-assist in that it does not require computation of a Jacobian matrix prior to defining a subspace of parameter space in which to solve an inverse problem. However this initial investment endows SVD-assist with a significant benefit. The \mathbf{V}_1 vectors

that are used by the SVD-assisted inversion process are designed to lie completely within the solution space of an inverse problem. For reasons that are stated in the previous chapter, this facilitates achievement of minimum error variance status for estimated parameters. Where an inverse problem is nonlinear this does not quite apply, for the boundaries between solution and null spaces may vary as parameter values vary. This may cause some \mathbf{V}_1 vectors to possess non-zero projections onto the evolving calibration null space. As an outcome of calibration-induced “null-space entrainment” of estimated parameter values, some model predictions may incur bias. This is because, by definition, a calibration dataset has nothing to say about parameter combinations that lie within the calibration null space. Hence any artificial expression of the null space in model predictions has as much chance of being wrong as it does of being right. As long as an inverse problem is not too nonlinear, however, the magnitude of null space entrainment is unlikely to be large, especially if a superfluity of super-parameters is employed.

In contrast to SVD-assist, there can be no guarantee that the subspace of parameter space in which ENSI operates does not have a substantial projection onto the calibration null space. Hence some amount of calibration-induced parameter and predictive bias is inevitable. (The same applies to ensemble methods in general.) ENSI bias can only be mitigated by increasing the number of realizations that define the subspace in which inversion takes place. The greater is the number of realizations on which the ENSI inversion process is based, the greater is the chance of null space cancellation as realizations are combined to create a parameter field that calibrates a model.

On the other hand, model run efficiency is also an important consideration, so trade-offs must be made.

In general, a good rule of thumb is to assign greater numbers of realizations to groupings of parameters that are likely to be most informed by a calibration dataset. This ensures that there are enough receptacles for information that the group must receive as a model is calibrated. Conversely, in order to conserve model runs, fewer realizations may be assigned to groups that are unlikely to be well informed by a calibration dataset.

As described above, once an ENSI inverse problem has been constructed, it is solved using standard regularization techniques that are discussed in preceding chapters of this book and that are implemented in PEST and PEST_ HP. A program named ENSIPREP that is supplied with the PEST_ HP suite constructs a PEST input dataset for ENSI inversion based on a PEST input dataset for standard inversion. Other PEST utilities translate ENSI super-parameters (i.e. **f**-parameters that are featured in the above equations) to **k**-parameters. When undertaking ENSI inversion, it is suggested that PEST’s Broyden Jacobian update functionality be activated; see section 7.4.2. This strengthens ENSI performance in difficult inversion settings such as those that arise from model nonlinearity and from joint estimation of native model parameters and geostatistical hyperparameters.

9.4.3.2 Strategic Definition of Parameter Groups

The ability to estimate some groups of parameters by working in subspaces that are specific to those groups while estimating other parameters individually increases inversion efficiency at the same time as it reduces the potential for predictive bias in some inversion settings.

Experience demonstrates that when estimating parameters that have spatial connotations (for example pilot point parameters), it is often a good idea to define these parameters as perturbations about their mean values so that they become addend or multiplier parameters.

Meanwhile, the mean hydraulic property value (which can be spatially variable) can also be denoted as a parameter so that it too is estimated through the inversion process. This happens naturally where pilot points are deployed in noncentred parameterization schemes that enable hierarchical inversion; see section 5.4. Where a mean is uniform-valued, it should be declared as a nonrealization parameter. Alternatively, if a mean is spatially variable, parameters which govern its spatial variation can be assigned to a different realization group from parameters which define perturbations about this mean. The same applies to other geostatistical hyperparameters. Separation of parameters that express the locations of hydraulic property heterogeneity from those that express the background value of a particular hydraulic property and/or its correlation structure reduces the chances of one parameter type adopting roles that should more properly be adopted by another parameter type. This, in turn, reduces the chances of parameter and predictive bias. Experience has repeatedly demonstrated the good performance of ENSI in undertaking hierarchical inversion.

9.4.3.3 Weaknesses of ENSI

The principal disadvantage of ENSI is the same as that which afflicts all dimensional reduction methods (including Bayesian ensemble methods that are discussed later in this book). This is the possibility of parameter and predictive bias incurred by working with a reduced number of parameters in a subspace of parameter space that probably does not include the correct parameter field, and may not even include the projection of the correct parameter field onto the calibration solution space. Attainment of a good fit with a calibration dataset indicates that the projection of the \mathbf{k} field that is obtained using ENSI onto the “true” solution space of the inverse problem is correct. However it says nothing about the projection of this field onto the “true” null space of the inverse problem. Ideally, this projection should be zero. As stated above, defence against bias of this kind can be achieved by increasing the number of realizations on which the ENSI inversion process is based. However, there is generally an upper limit to which this can be done because of the high numerical cost that this incurs in many inversion settings (this being the reason for the decision to employ ENSI in the first place).

The subspace in which ENSI works is defined by random realizations of groups of parameters, as well as by individual nonrealization parameters. Any random parameter generation process relies on a random number seed. An obvious question is whether use of a different random number seed would yield a different ENSI-calibrated parameter field. The answer to this question is “yes”. However the magnitude of this difference depends on the size of the calibration dataset and on the number of realizations that ENSI employs. Differences between estimated parameter fields where random seeds are different is a measure of potential bias incurred through dimensional reduction. The potential for predictive bias is higher in areas that are far removed from points at which calibration data were gathered than where data density is high. However predictive uncertainties are likely to be high in these same areas. Ideally, the potential for predictive bias is spanned by predictive uncertainty intervals. As stated many times in this book, calibration should always be followed by predictive uncertainty analysis.

9.4.4 Methodology Variations

As is explained in PEST documentation, preparation for ENSI is relatively easy. A user must first prepare a PEST input dataset based on \mathbf{k} parameters. This must be accompanied by specification of prior parameter uncertainty, including covariance matrices that pertain to different parameter groups. A modeller must then allocate different subsets of parameters to different realization groups while also identifying nonrealization parameters. A program named ENSIPREP then builds a complete PEST input dataset for ENSI inversion.

ENSIPREP can accommodate a number of different ENSI implementation variations. All of these variations retain the principle ENSI design objective of dimensional reduction. One such variation is referred to as “PCA inversion” in PEST documentation.

When implementing PCA inversion, instead of generating realizations of parameters based on covariance matrices that pertain to subsets of parameters, these covariance matrices are subjected to singular value decomposition. For example, suppose that

$$C(\mathbf{k}_i) = \mathbf{E}\mathbf{F}\mathbf{E}^t \quad (9.4.10)$$

where $C(\mathbf{k}_i)$ is defined in equation (9.4.2). \mathbf{M}_i matrices that appear in equations (9.4.8) and (9.4.9) are then constructed as

$$\mathbf{M}_i = \mathbf{E}_1 \mathbf{F}_1^{1/2} \quad (9.4.11)$$

The “1” subscript on the right of (9.4.11) indicates truncation at the column limit of r_i ascribed to each \mathbf{M}_i matrix. Because this variation of ENSI does not require random number generation, the propensity for calibration-induced predictive bias is reduced. Under these circumstances

$$C(\mathbf{f}_i) = \mathbf{I} \quad (9.4.12)$$

$C(\mathbf{f}_i)$ is used for ENSI super-parameter regularization. This strategy of ENSI parameter definition has much in common with Kahanen-Loève parameter transformation. However ENSI continues to allow estimation of individual parameters (i.e. “nonrealisation parameters”) along with estimation of factors that are applied to covariance matrix eigencomponents.

In still another variation of PCA inversion, ENSIPREP can be supplied with user-prepared realizations of parameters belonging to different realization groups. It does not, therefore, generate these realizations itself based on user-supplied covariance matrices. A collection of such realizations comprises the columns of an \mathbf{M}_i matrix. ENSIPREP calculates an empirical covariance matrix $C(\mathbf{k}_i)$ for this realization group as

$$C(\mathbf{k}_i) = \frac{(\mathbf{M}_i - \underline{\mathbf{M}}_i)(\mathbf{M}_i - \underline{\mathbf{M}}_i)^t}{r-1} \quad (9.4.13)$$

Where $\underline{\mathbf{M}}_i$ is a matrix of realization means; see section 3.5.1.

It then subjects this empirical covariance matrix to singular value decomposition in order to obtain its eigencomponents. PCA inversion is then undertaken in the manner described above.

In this second ENSI PCA option, the concept of “ \mathbf{k} parameter” starts to lose its meaning, for ENSIPREP does not generate realizations of \mathbf{k} itself; nor does it have any knowledge of how random \mathbf{k} -based hydraulic property fields are generated. Instead stochastic field generation is undertaken by another program. There is nothing to prevent this other program from generating these realizations using sophisticated geostatistical methodologies that do not require a covariance matrix. Parameter fields can therefore be fully or partly categorical. The ENSI PCA inversion process is not degraded by this. ENSI simply estimates factors by which eigencomponents of the empirical covariance matrix of equation (9.4.13) must be multiplied for model outputs to achieve good fits with field measurements of system behaviour. If required, individual nonrealization parameters can be estimated along with factors applied to these complex hydraulic property field eigencomponents.

If desired, all of the above dimensional reduction methodologies can be combined in a single ENSI inversion process.

9.4.5 Post-ENSI Uncertainty Analysis

We close this section on ENSI by looking ahead.

The next sections of this book address the subject of uncertainty analysis. This can be done following model calibration. Or it can be done instead of model calibration (as is often the case when ensemble methods are employed). The reader may be wondering how parameter/predictive uncertainty analysis is best undertaken following ENSI calibration of a model. This subject is discussed later in this book; see section 12.4.7.

10. Linear Uncertainty and Error Analysis

10.1 General

The present chapter discusses insights into inversion and uncertainty analysis that can be gained by treating a model as linear.

Some of the formulas that are derived herein are rather esoteric, and would not be used in everyday modelling practice. However they do provide insights into issues such as that of how information is stored in a field measurement dataset and how this information flows from observations to model parameters. Given the necessity for decision support modelling to harvest information from data, and to deliver that information to decision-makers in ways that they can understand, these insights may prove valuable in sharpening a modeller's intuition. They also illustrate the depth of understanding that singular value decomposition can yield when turned to the analysis of inverse problem-solving. This level of understanding is important when it is remembered that, even where ensemble-based Bayesian analysis is pursued, practical modelling, and model parameterization, is always confined to a subspace of reality. Singular value decomposition sets the standard for optimality of subspace definition by which other methods can be judged.

Formulas that are derived in the present chapter support linear analysis methodologies that can be readily applied either before or after calibration of a model. These analyses can provide approximate, but rapid, estimates of pre- and post-calibration parameter and predictive uncertainty. They can also be used to appraise the ability of different data types to reduce the uncertainties of predictions of interest, even if these data have not yet been gathered. Some of these formulas also form the basis for nonlinear uncertainty analysis methods that are discussed in subsequent chapters. Many of them are encapsulated in linear analysis utilities that are provided with the PEST suite.

All of the analyses that are discussed herein require a Jacobian matrix; such a matrix is generally computed as part of the parameter estimation process, or can easily be computed after the calibration process has been completed. Some analyses require predictive sensitivities as well.

Linear analysis is based on the premise that sensitivities embodied in a Jacobian matrix do not change with parameter values. Most models, however, are nonlinear. Nevertheless, rarely is a model so nonlinear as to invalidate insights that can be gained through linear analysis whether the Jacobian matrix is calculated using prior mean parameter values, or parameter values obtained through model calibration. Often, however, linear analysis is conducted following model calibration. In this case it is best to use sensitivities that are calculated using calibrated parameter values.

Some of the concepts that are discussed in the present chapter have already been discussed in previous chapters. However we seek justification for some repetition in the interests of keeping similar topics in one place. We also note that an additional linear analysis option is presented in chapter 15 of this book where the effects of model defects are explored.

As usual, the inverse problem that forms the basis for present discussions is embodied in the following equation

$$\mathbf{h} = \mathbf{Zk} + \boldsymbol{\varepsilon} \quad (10.1.1)$$

The prior covariance matrix of \mathbf{k} is denoted as $\mathbf{C}(\mathbf{k})$ and the covariance matrix of measurement

noise is denoted as $C(\boldsymbol{\varepsilon})$.

Where a model prediction s is considered, the sensitivity of this prediction to model parameters is denoted as \mathbf{y} . Hence

$$s = \mathbf{y}^t \mathbf{k} \quad (10.1.2)$$

10.2 Post-Calibration Subspace-Based Analyses

10.2.1 Parameter Identifiability

Doherty and Hunt (2009) define the identifiability of a parameter as the square of the cosine between a vector pointing in the direction of the parameter and the projection of that vector onto the calibration solution space. This projection is zero where a parameter lies entirely within the null space; it is one when a parameter lies entirely within the solution space. If the identifiability of a parameter is zero, then its uncertainty is not reduced through the history-matching process. However if its identifiability is unity, this does not mean that its estimation takes place without any potential for error. What it does indicate is that all of its potential for error is inherited from noise in the measurement dataset.

A parameter's identifiability is readily calculated from the outcomes of singular value decomposition undertaken on the Jacobian matrix that replaces \mathbf{Z} in equation (10.1.1). In fact, singular value decomposition should be undertaken on the weighted Jacobian matrix $\mathbf{Q}^{1/2}\mathbf{Z}$, where \mathbf{Q} is chosen such that

$$\mathbf{Q} = C^{-1}(\boldsymbol{\varepsilon}) \quad (10.2.1)$$

For convenience, $C(\boldsymbol{\varepsilon})$ is generally assumed to be diagonal despite the significant contribution to model-to-measurement misfit that is made by structural noise.

Ideally, a parameter set \mathbf{k} should have been Kahunen-Loëve transformed prior to calculation of identifiability in order that the prior parameter covariance matrix is equal to the identity matrix. The reasons for this are discussed in section 8.2.4. However this may not always be convenient or useful, as Kahunen-Loëve transformed parameters are difficult for a modeller to recognize because they pertain to the eigencomponents of the $C(\mathbf{k})$ prior parameter covariance matrix. Hence identifiabilities are normally calculated for \mathbf{k} parameters rather than for transformed parameters.

The repercussions of failure to acknowledge the non-diagonal nature of $C(\mathbf{k})$ in calculation of identifiabilities of individual elements of \mathbf{k} can be illustrated as follows. Suppose that two elements of \mathbf{k} , say k_i and k_j , exhibit prior correlation. Suppose also that the calibration dataset \mathbf{h} is directly informative of k_i but not of k_j . The posterior uncertainty of k_j is nevertheless smaller than its prior uncertainty because of the conditioning effect of k_i (whose uncertainty is directly reduced through the calibration process) on k_j . However this is not reflected in a reduced identifiability for k_j if identifiability is calculated using $\mathbf{Q}^{1/2}\mathbf{Z}$ based on the \mathbf{Z} matrix of equation (10.1.1). For the same reason, if the spatial density of a parameterization device such as pilot points is increased over a part of the model domain, then the identifiability of individual pilot point parameters will decrease. It follows that identifiability, applied to individual parameters, does not elucidate information flow per unit area because of failure to acknowledge spatial parameter correlation, which is likely to be high where parameters are close together. This does not diminish its usefulness (particularly when applied to lumped or zone-based parameters), as long as its shortcomings are understood. Note also that other, Bayesian-based, statistics that will be presented later in this chapter readily accommodate a non-diagonal $C(\mathbf{k})$ matrix and can, in fact, elucidate flow of information per unit area.

Suppose that the matrix $\mathbf{Q}^{1/2}\mathbf{Z}$ is subjected to singular value decomposition (SVD). As described in early chapters, this operation is expressed by the equation

$$\mathbf{Q}^{1/2}\mathbf{Z} = \mathbf{USV}^t = \mathbf{U}_1 \mathbf{S}_1 \mathbf{V}_1^t + \mathbf{U}_2 \mathbf{S}_2 \mathbf{V}_2^t \quad (10.2.2)$$

In equation (10.2.2) the orthogonal matrix \mathbf{V} is partitioned into submatrices whose columns span two orthogonal-complementary parameter subspaces. These are the inverse problem solution and null spaces.

The matrix operator that projects an arbitrary parameter vector \mathbf{k} onto the calibration solution space is $\mathbf{V}_1 \mathbf{V}_1^t$. As is discussed in section 8.2, partitioning of the \mathbf{V} matrix into \mathbf{V}_1 and \mathbf{V}_2 submatrices is not generally such that the singular values associated with \mathbf{V}_2 are all zero. Instead, singular value truncation takes place prior to singular values becoming zero in order to prevent amplification of measurement noise when estimating values for parameters. Selection of the optimum truncation point is often inexact; this can affect identifiabilities that are calculated for some parameters. This adds to the somewhat arbitrary nature of this statistic.

Let \mathbf{x} be an arbitrary vector, and let \mathbf{y} be its projection onto the calibration solution space spanned by the columns of \mathbf{V}_1 . Then

$$\mathbf{y} = \mathbf{V}_1 \mathbf{V}_1^t \mathbf{x} \quad (10.2.3)$$

The cosine of the angle between \mathbf{x} and \mathbf{y} is their scalar product divided by the magnitude of each. Let us refer to this angle as θ . Then

$$\cos \theta = \frac{\mathbf{x}^t \mathbf{V}_1 \mathbf{V}_1^t \mathbf{x}}{\sqrt{\mathbf{x}^t \mathbf{V}_1 \mathbf{V}_1^t \mathbf{x} \sqrt{\mathbf{x}^t \mathbf{x}}}} \quad (10.2.4)$$

So that

$$\cos^2 \theta = \frac{\mathbf{x}^t \mathbf{V}_1 \mathbf{V}_1^t \mathbf{x}}{\mathbf{x}^t \mathbf{x}} \quad (10.2.5)$$

If \mathbf{x} is a unit vector that points in the direction of one of the elements of \mathbf{k} , let us say k_i , then all of its elements are zero except for its i 'th element, which is one. For this vector

$$\cos^2 \theta = (\mathbf{V}_1 \mathbf{V}_1^t)_{ii} \quad (10.2.6)$$

From equation (8.2.10) it follows that the identifiability of parameter i is the i 'th diagonal element of the resolution matrix that emerges from use of singular value decomposition as a regularization device. Recall that the “ideal” resolution matrix is the identity matrix. However such an ideal resolution matrix can only emerge from a parameter estimation process where there is no null space. Where there is no null space, all parameters possess an identifiability of unity.

The columns of matrix \mathbf{V}_1 are comprised of unit orthogonal vectors which we denoted as \mathbf{v}_i . Thus

$$\mathbf{V}_1 = [\mathbf{v}_1 \ \mathbf{v}_2 \ \dots \ \mathbf{v}_i \ \dots \ \mathbf{v}_w] \quad (10.2.7)$$

where w is the number of singular values before truncation. If this is substituted into (10.2.6) and expanded we obtain

$$\cos^2 \theta = \sum_{j=1}^w (\mathbf{v}_j \mathbf{v}_j^t)_{ii} \quad (10.2.8)$$

$(\mathbf{v}_j \mathbf{v}_j^t)_{ii}$ is the square of the cosine of the angle between a unit vector in the direction of parameter k_i and its projection onto the unit \mathbf{v}_j vector. If this angle is denoted as θ_j , then

$$\cos^2 \theta = \sum_{j=1}^w \cos^2 \theta_j \quad (10.2.9)$$

The above equations can form the basis for some interesting and informative post-calibration plots. Where parameters have a spatial setting, identifiability maps can be drawn. Simple bar charts of identifiability are also readily plotted. These bar charts can be stacked so that contributions to total parameter identifiability by each of the \mathbf{v}_j vectors are visible in each bar. If these contributions are colour-coded according to j , with warmer colours (i.e. shades of red) assigned to j 's for which corresponding singular values are higher, the effect is not only stunning but informative. If the identifiability bars pertaining to two different parameters are the same height, but one bar has warmer colours than the other, this signifies that a value can be estimated for the former parameter with less potential for error than for the latter parameter as its estimate is less likely to be contaminated by measurement noise.

As stated, a parameter's identifiability is 1.0 or less. Perhaps the complement to this should be named “non-identifiability”. It is easy to show that the number obtained by subtracting a parameter's identifiability from 1.0 is the squared cosine of the angle between a vector pointing in the direction of that parameter and its projection onto the null space of $\mathbf{Q}^{1/2}\mathbf{Z}$. The latter is spanned by unit vectors comprising the columns of \mathbf{V}_2 . If a parameter has a high “non-identifiability” then the calibration dataset carries little or no information pertaining to that parameter. Note however that, as discussed above, if the parameter exhibits prior correlation with an identifiable parameter, then its uncertainty is nevertheless reduced through history-matching as an outcome of that correlation.

While parameter identifiability and non-identifiability statistics are informative, they must be used with caution. Indeed, as stated above, there is some degree of arbitrariness associated with these statistics, for their values depend on where singular value truncation takes place. Where singular value truncation aims to minimize parameter and predictive error variance, the truncation point may be based on a curve such as that depicted in figure 8.4. In practice the minimum of this curve may be broad, perhaps spanning tens of singular values. The identifiabilities of parameters with high projections onto \mathbf{v}_i vectors which correspond to this interval will be very sensitive to the truncation point.

10.2.2 Other Subspace Insights

10.2.2.1 Transformed Equations

In addition to parameter identifiability, a number of other post-calibration statistics can readily be computed based on subspace concepts. These can reveal much about the status of a recently-completed inversion process. Some of them may be considered for inclusion in a report on the calibration process because of insights that they provide into the information content of a calibration dataset, and into how that information is accessed by parameters.

Ideally, calculation of these statistics should take place in a Kahunen-Loëve transformed parameter environment, regardless of whether Kahunen-Loëve transformation was undertaken prior to the preceding calibration process. Where the inversion process employed Tikhonov regularization, then pre-calibration Kahunen-Loëve parameter transformation will probably not have been undertaken, as Tikhonov constraints will have been the primary mechanism for expressing expert knowledge. However as it can be difficult to include Tikhonov constraints, and weight them appropriately, when evaluating subspace-based post-calibration statistics, Kahunen-Loëve transformation constitutes the most practical means of incorporating prior knowledge into their calculation.

As usual we take (10.1.1) as our starting point and assume that a weight matrix has been chosen according to (10.2.1). We suppose that $C(\mathbf{k})$ can be decomposed through singular value decomposition as

$$C(\mathbf{k}) = \mathbf{E}\mathbf{F}\mathbf{E}^t \quad (10.2.10)$$

A Kahunen-Loève transformed parameter set \mathbf{m} can be calculated as

$$\mathbf{m} = \mathbf{F}^{-1/2}\mathbf{E}^t\mathbf{k} \quad (10.2.11)$$

so that

$$C(\mathbf{m}) = \mathbf{I} \quad (10.2.12)$$

Back-transformation from \mathbf{m} -space to \mathbf{k} -space is achieved through the equation

$$\mathbf{k} = \mathbf{E}\mathbf{F}^{1/2}\mathbf{m} \quad (10.2.13)$$

With appropriate weight transformation, equation (10.1.1) can now be written as

$$\mathbf{g} = \mathbf{Y}\mathbf{m} + \boldsymbol{\xi} \quad (10.2.14)$$

where

$$\mathbf{g} = \mathbf{Q}^{1/2}\mathbf{h} \quad (10.2.15)$$

$$\mathbf{Y} = \mathbf{Q}^{1/2}\mathbf{Z}\mathbf{E}\mathbf{F}^{1/2} \quad (10.2.16)$$

$$\boldsymbol{\xi} = \mathbf{Q}^{1/2}\boldsymbol{\epsilon} \quad (10.2.17)$$

Obviously, if \mathbf{Q} is chosen according to (10.2.1)

then

$$C(\boldsymbol{\xi}) = \mathbf{I} \quad (10.2.18)$$

Let us now undertake singular value decomposition on \mathbf{Y} , so that

$$\mathbf{Y} = \mathbf{U}\mathbf{S}\mathbf{V}^t \quad (10.2.19)$$

10.2.2.2 Model Matrix Products

From (10.2.19) it immediately follows that

$$\mathbf{Y}\mathbf{Y}^t = \mathbf{U}\mathbf{S}^2\mathbf{U}^t \quad (10.2.20)$$

and that

$$\mathbf{Y}^t\mathbf{Y} = \mathbf{V}\mathbf{S}^2\mathbf{V}^t \quad (10.2.21)$$

If measurement noise is ignored, then from (10.2.14)

$$C(\mathbf{g}) = \mathbf{Y}C(\mathbf{m})\mathbf{Y}^t = \mathbf{Y}\mathbf{Y}^t \quad (10.2.22)$$

Hence $\mathbf{Y}\mathbf{Y}^t$ is the covariance matrix that denotes the variability of model outputs \mathbf{g} as this arises from the natural variability of parameters \mathbf{m} . This is the “evidence” term of Bayes equation; see equation 6.2.4. This can be compared with the variability that \mathbf{g} possesses by virtue of its noise content; with transformations as above, this is described by the covariance matrix \mathbf{I} . The

variability of individual model outputs arising from natural parameter variability is therefore described by the diagonal elements of $\mathbf{Y}\mathbf{Y}^t$. This variability is a measure of their information content with respect to parameters. High parameter-induced variability compared to a noise variance of unity indicates high accessible information content with respect to the parameters that induce that variability. The diagonal elements of $\mathbf{Y}\mathbf{Y}^t$ are closely related to the composite observation sensitivities defined in equation (7.3.4).

$\mathbf{Y}^t\mathbf{Y}$, on the other hand, can be viewed in a number of ways. As is apparent from equation (7.3.1), each diagonal element of this matrix is closely related to the composite sensitivity of the corresponding Kahunen-Loëve-transformed parameter. Because Kahunen-Loëve transformation effectively scales parameters by their innate variabilities, the diagonal elements of $\mathbf{Y}^t\mathbf{Y}$ are somewhat reminiscent of the composite scaled sensitivity statistic of Hill and Tiedeman (2007); however scaling is a little more theoretically based. Recall from equation (7.2.18) that if an inverse problem is well-posed $(\mathbf{Y}^t\mathbf{Y})^{-1}$ is the post-calibration parameter covariance matrix. However if the inverse problem is not well-posed, then this matrix cannot be formed because $\mathbf{Y}^t\mathbf{Y}$ cannot be inverted. Because the diagonal elements of $(\mathbf{Y}^t\mathbf{Y})^{-1}$ denote individual post-calibration parameter uncertainties, the diagonal elements of $\mathbf{Y}^t\mathbf{Y}$ can be loosely viewed as denoting post-calibration “parameter certainties”. A low value of post-calibration “parameter certainty” denotes a low ability to be inferred through the calibration process. The post-calibration uncertainties of such parameters must therefore be constrained by expert knowledge, as information that is resident in the calibration dataset is insufficient to constrain them.

There is another, slightly more abstract way, to view the diagonal elements of $\mathbf{Y}^t\mathbf{Y}$. From (10.2.21) $\mathbf{Y}^t\mathbf{Y}$ can be written as

$$\mathbf{Y}^t\mathbf{Y} = \sum_{j=1}^m s_j^2 \mathbf{v}_j \mathbf{v}_j^t \quad (10.2.23)$$

where s_j^2 are singular values and m is the number of parameters featured in the inverse problem. From considerations presented in the preceding section, each diagonal element of $\mathbf{Y}^t\mathbf{Y}$ is therefore equal to the sum of weighted squared cosines between each transformed parameter m_i (see equation 10.2.11) and each unit vector \mathbf{v}_j emerging from the singular value decomposition process; weights used in the summation are corresponding squared singular values. If parameter m_i projects strongly onto \mathbf{v}_j unit vectors that span the null space, or onto \mathbf{v}_j unit vectors for which corresponding singular values are very low, then the i 'th diagonal element of $\mathbf{Y}^t\mathbf{Y}$ is small.

10.2.2.3 Singular Values

The \mathbf{S} matrix has only diagonal elements, these being singular values. Singular values are listed from highest to lowest down the diagonal; all are positive. Where an inverse problem is ill-posed their values diminish to zero or near zero. An inspection of these singular values can be most informative. If they decay rapidly to zero, this informs a modeller that the information content of a calibration dataset is directed to only a small number of receptacles. These receptacles comprise the few \mathbf{v}_i vectors for which corresponding singular values are high. Alternatively, if singular value decay is slow, then the information content of the calibration dataset is informative of many combinations of parameters (i.e. many \mathbf{v}_i vectors). Where singular values decay to about 10^{-4} of the highest singular value, this marks the point where, from numerical considerations alone, the workable solution space ends. This is especially the case in parameter estimation settings where the Jacobian matrix is calculated using finite parameter differences, for this marks the point where numerical noise incurred through finite-

difference derivatives calculation may be too grossly amplified to allow parameter combinations defined by corresponding \mathbf{v}_i vectors to be estimated with integrity.

In practice, as has already been discussed, singular value truncation should take place before this point is reached because of the presence of measurement noise in the calibration dataset. A method for determining the optimal truncation point was suggested in section 8.2.5.3. In the present context this consists of determining the singular value at which $s_i^2(\mathbf{v}_i^t \mathbf{v}_i)^2$ is greater than $(\mathbf{v}_i^t \mathbf{v}_i)^2$, that is where s_i^2 is greater than 1 (as $\mathbf{v}_i^t \mathbf{v}_i$ is always equal to 1). This, of course, occurs where s_i falls below 1. This is the point where the calibration process has the potential to incur more error in estimation of the \mathbf{v}_i parameter combination than the uncertainty that this combination of parameters already possesses based on expert knowledge alone. This defines the optimal dimensionality of the solution space (and hence the optimal dimensionality of the orthogonal complementary null space). This is highly informative, and is worth reporting following calibration of a model. The number of singular values that are greater than 1 is equal to the number of accessible pieces of information that a calibration dataset hosts.

10.2.2.4 Transformed Parameter Status

Once a singular value truncation point has been established, and parameter space has thereby been subdivided into mutually orthogonal solution and null subspaces, the identifiabilities and non-identifiabilities of \mathbf{m} -space parameters are easily computed as the diagonal elements of $\mathbf{V}_1 \mathbf{V}_1^t$ and $\mathbf{V}_2 \mathbf{V}_2^t$ respectively.

The diagonal elements of $\mathbf{V}_1 \mathbf{S}_1^2 \mathbf{V}_1^t$ and $\mathbf{V}_1 \mathbf{S}_1^{-2} \mathbf{V}_1^t$ are also worthy of consideration. If the calibration solution space comprises the whole of parameter space, the diagonal elements of $\mathbf{V}_1 \mathbf{S}_1^2 \mathbf{V}_1^t$ are equal to the diagonal elements of $\mathbf{Y}^t \mathbf{Y}$. If not, they are equal to or less than these. If the solution space comprises the entirety of parameter space, $\mathbf{V}_1 \mathbf{S}_1^{-2} \mathbf{V}_1^t$ is equal to $(\mathbf{Y}^t \mathbf{Y})^{-1}$; this is the post-calibration covariance matrix of Kahanen-Loève transformed parameter error. If not, each diagonal element of $\mathbf{V}_1 \mathbf{S}_1^{-2} \mathbf{V}_1^t$ is equal to the solution space component of post-calibration Kahanen-Loève transformed parameter error. In both cases the i 'th diagonal element pertains to the i 'th element of \mathbf{m} . For each such parameter, the sum of this diagonal element and the respective diagonal element of $\mathbf{V}_2 \mathbf{V}_2^t$ is the total post-calibration error variance. That is

$$\mathbf{C}(\underline{\mathbf{m}} - \mathbf{m}) = \mathbf{V}_2 \mathbf{V}_2^t + \mathbf{V}_1 \mathbf{S}_1^{-2} \mathbf{V}_1^t \quad (10.2.24)$$

From (10.2.13) and (3.9.2) the post-calibration covariance matrix of native model parameter error is

$$\mathbf{C}(\underline{\mathbf{k}} - \mathbf{k}) = \mathbf{E} \mathbf{F}^{1/2} \mathbf{V}_2 \mathbf{V}_2^t \mathbf{F}^{1/2} \mathbf{E}^t + \mathbf{E} \mathbf{F}^{1/2} \mathbf{V}_1 \mathbf{S}_1^{-2} \mathbf{V}_1^t \mathbf{F}^{1/2} \mathbf{E}^t \quad (10.2.25)$$

(Note that this formulation differs from equation (8.2.32) because, in the present subsection, singular value decomposition is undertaken on \mathbf{Y} of equation (10.2.16) rather than on \mathbf{Z} . Back-transformation to \mathbf{k} -space is therefore necessary.)

10.2.2.5 Relative parameter error variance reduction

The i 'th diagonal element of $\mathbf{C}(\underline{\mathbf{k}} - \mathbf{k})$ is the post-calibration error variance of parameter i . This can be denoted as σ_{ei}^2 . Let σ_i^2 denote the prior error variance of this same parameter; this is the same as the variance of its prior uncertainty. We denote the relative error variance reduction of parameter i as rev_i . This can be written as

$$rev_i = 1 - \frac{\sigma_{ei}^2}{\sigma_i^2} \quad (10.2.26)$$

Ideally, the relative error variance reduction of a parameter should, like the identifiability of a parameter, range between zero and one. However there is no guarantee that rev_i will always be non-negative. Negativity of rev_i will occur if singular value truncation takes place at a point where the parameter's post-calibration error variance, as depicted in a plot such as that shown in figure 8.4, is greater than its pre-calibration error variance. The potential for this to occur will, of course, be exacerbated for those parameters for which the minimum of the error variance curve is only shallow.

10.2.2.6 Flow of Information

The solution to the inverse problem of estimating Kahunen-Loëve transformed parameters \mathbf{m} is

$$\underline{\mathbf{m}} = \mathbf{V}_1 \mathbf{S}^{-1} \mathbf{U}^t \mathbf{g} \quad (10.2.27)$$

which can also be written as

$$\mathbf{S}^{-1} \mathbf{U}^t \mathbf{g} = \mathbf{V}^t \underline{\mathbf{m}} \quad (10.2.28)$$

As is discussed in section 8.2.6.3, the orthogonal unit vectors \mathbf{u}_i comprising the columns of \mathbf{U} constitute the “signals” or “characteristics” of a measurement dataset that contain the information which it hosts. For the i 'th such signal, the information itself is the i 'th singular value s_i . This information flows uniquely and entirely to the i 'th parameter receptacle; the latter is the vector \mathbf{v}_i comprising the i 'th column of \mathbf{V} .

Identification of “information signals” and “information receptacles” can sometimes be useful. Where observations and/or parameters have spatial and/or temporal connotations, then it can sometimes be instructive to plot the individual elements that comprise the elements of \mathbf{u}_i and \mathbf{v}_i in space and/or time. For example, if the calibration dataset is comprised of a streamflow time series, decomposition of the time series in this manner may allow a modeller to see the features of this time series which are most information-rich, at the same time as he/she sees the parameter combinations that these various streamflow characteristics inform. However before making such a plot, Kahunen-Loëve parameters must be back-transformed to native parameters, and weight-transformed measurements must be back-transformed to native measurements. From (10.2.11) and (10.2.15), (10.2.28) becomes

$$\mathbf{S}^{-1} \mathbf{U}^t \mathbf{Q}^{1/2} \mathbf{h} = \mathbf{V}^t \mathbf{F}^{-1/2} \mathbf{E}^t \underline{\mathbf{k}} \quad (10.2.29)$$

The vectors comprising the columns of $\mathbf{Q}^{1/2} \mathbf{U}_1$ and $\mathbf{E} \mathbf{F}^{-1/2} \mathbf{V}_1$, even though they are no longer orthogonal unit vectors, define “information signals” in native observation space on the one hand, and the receptacles for that information in native parameter space on the other hand. The elements of these vectors thus have space and/or time connotations and can be plotted as such.

10.2.2.7 Influence Statistics

Observation influence statistics are discussed in section 7.3.3. Though formulated for well-posed inverse problems, they can still be used in highly parameterized contexts if an ill-posed problem is re-formulated as a well-posed problem. This can be done by supplementing the calibration dataset with appropriately weighted prior information. Alternatively it may be accomplished by restricting the search for the solution to the ill-posed inverse problem to the subspace of parameter space that is spanned by the vectors comprising the columns of the \mathbf{V}_1 matrix obtained through singular value decomposition of \mathbf{Y} . However the information gained from such an analysis is unlikely to be useful.

To the extent that a prediction of interest is sensitive to parameters, or combinations of

parameters, that occupy the null space, its uncertainty is not reduced through history-matching. If such a prediction is of management interest, there may be a strong incentive to reduce its uncertainty. The most effective means by which to achieve this reduction is to reduce the dimensionality of the null space in a manner that is specific to this prediction. In particular, parameters and parameter combinations to which the prediction is sensitive should be shifted from the null space to the solution space. It follows that influence statistics must accommodate the existence of the null space if they are to identify observations, or potential observations, which can accomplish this most successfully. Analyses which serve this purpose are presented shortly.

10.3 Post-Calibration Uncertainty Analysis

10.3.1 Formulation of Equations

Equations derived in the previous section are based on optimal calibration of a linear model achieved using truncated singular value decomposition applied to a Kahunen-Loëve-transformed parameter set. In these equations, the calibration process is conceptual only. Nevertheless evaluation of these equations using a Jacobian matrix derived from calibration of a real-world model can provide important insights into what the actual (rather than conceptual) model calibration process was able to achieve.

The same theme is continued in the present section. However, rather than adopting a data assimilation approach that is based on calibration followed by error analysis, the matrix equations that are discussed in the present section have Bayesian roots. Hence they describe conditioning of model parameters, and of predictions that are sensitive to these parameters, by field measurements of system behaviour. The same measurement dataset is employed as in the previous section. However the focus is on posterior uncertainty rather than on post-calibration propensity for parameter error. As we shall see, these two quantities are related. However evaluated uncertainty is generally smaller than error potential when evaluation of the latter is based on subspace principles. Hence uncertainty provides a more “statistically efficient” characterization of the benefits of history-matching (provided that all of the assumptions on which its calculation is based are correct).

As usual, equation (10.1.1) is our starting point. To this we add a rather obvious equation.

$$\mathbf{k} = \mathbf{Ik} \quad (10.3.1)$$

This equation says nothing other than that parameters are related to themselves through the identity matrix. If this is combined with (10.1.1) in a single matrix equation we obtain

$$\begin{bmatrix} \mathbf{k} \\ \mathbf{h} \end{bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{Z} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{k} \\ \boldsymbol{\varepsilon} \end{bmatrix} \quad (10.3.2)$$

Applying equation (3.9.2) to propagate variance, the joint covariance matrix of parameters and observations of system state can be written as

$$\mathbf{C} \left(\begin{bmatrix} \mathbf{k} \\ \mathbf{h} \end{bmatrix} \right) = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{Z} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{C}(\mathbf{k}) & \mathbf{0} \\ \mathbf{0} & \mathbf{C}(\boldsymbol{\varepsilon}) \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{Z}^t \\ \mathbf{0} & \mathbf{I} \end{bmatrix}$$

That is

$$\mathbf{C} \left(\begin{bmatrix} \mathbf{k} \\ \mathbf{h} \end{bmatrix} \right) = \begin{bmatrix} \mathbf{C}(\mathbf{k}) & \mathbf{C}(\mathbf{k})\mathbf{Z}^t \\ \mathbf{Z}\mathbf{C}(\mathbf{k}) & \mathbf{Z}\mathbf{C}(\mathbf{k})\mathbf{Z}^t + \mathbf{C}(\boldsymbol{\varepsilon}) \end{bmatrix} \quad (10.3.3)$$

Using equation (3.8.5) we can now formulate the conditional covariance matrix of \mathbf{k} , the

condition being that we have gained perfect knowledge of \mathbf{h} through having acquired a measurement dataset. We denote this matrix as $C'(\mathbf{k})$. It must be remembered, however, that use of equation (3.8.5) rests on the assumption that the probability distributions with which the $C(\mathbf{k})$ and $C(\boldsymbol{\varepsilon})$ covariance matrices are associated are multiGaussian. We obtain

$$C'(\mathbf{k}) = C(\mathbf{k}) - C(\mathbf{k})\mathbf{Z}^t[\mathbf{Z}C(\mathbf{k})\mathbf{Z}^t + C(\boldsymbol{\varepsilon})]^{-1}\mathbf{Z}C(\mathbf{k}) \quad (10.3.4)$$

The first term on the right of equation (10.3.4) is the prior parameter covariance matrix. The second term on the right shows how history-matching to the calibration dataset \mathbf{h} alters this matrix. One or more of the diagonal elements of $C(\mathbf{k})$ is reduced through this subtraction. The extent to which it is reduced for a particular parameter depends on the information content of the calibration dataset with respect to that parameter.

$C'(\mathbf{k})$ is the covariance matrix associated with the posterior probability distribution of model parameters. Equation (10.3.4) can also be derived directly from Bayes equation under the assumptions of linear model behaviour and multiGaussianity of prior parameter uncertainty and measurement noise.

Calculation of $C'(\mathbf{k})$ using equation (10.3.4) requires that a $n \times n$ matrix be inverted, where n is the number of observations comprising the history-matching dataset. An equivalent formula can be derived in which a matrix of size $m \times m$ must be inverted, where m is the number of parameters in the parameter vector \mathbf{k} . From the Sherman-Morrison-Woodbury formula (2.15.14)

$$C'(\mathbf{k}) = [\mathbf{Z}^t C^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + C^{-1}(\mathbf{k})]^{-1} \quad (10.3.5)$$

It is important (and somewhat amazing) to note that equation (10.3.5) is equivalent to equation (7.2.39b). Thus where $C(\mathbf{k})$ and $C(\boldsymbol{\varepsilon})$ are known and pertain to multiGaussian distributions, and where the observation dataset is supplemented by preferred-value prior information that is weighted according to prior parameter uncertainty (while members of the history-matching dataset are weighted according to measurement noise), the post-calibration distribution of parameter error following Gauss-Newton parameter estimation is the same as the posterior parameter uncertainty distribution obtained through Bayes equation.

It is also easy to show that under the same circumstances parameter values estimated using the Gauss-Newton method are equal to their posterior expected values when conditioned by members of the history-matching dataset. Recalling that prior expected values of \mathbf{k} and \mathbf{h} are $\mathbf{0}$, equation (3.8.4) becomes

$$\underline{\mathbf{k}} = C(\mathbf{k})\mathbf{Z}^t[\mathbf{Z}C(\mathbf{k})\mathbf{Z}^t + C(\boldsymbol{\varepsilon})]^{-1}\mathbf{h} \quad (10.3.6)$$

Notice the similarity of this equation to the second equation on the right of (8.3.14). With appropriate weight matrix and Kahanen-Loève transformation they are the same.

With application of the identity expressed by (2.15.16) equation (10.3.6) becomes

$$\underline{\mathbf{k}} = [\mathbf{Z}^t C^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + C^{-1}(\mathbf{k})]^{-1} \mathbf{Z}^t C^{-1}(\boldsymbol{\varepsilon}) \mathbf{h} \quad (10.3.7)$$

This equation is identical to (7.2.38).

10.3.2 Relative Parameter Uncertainty Variance Reduction

The diagonal elements of $C'(\mathbf{k})$ express the posterior variances of parameter uncertainty. A similar statistic to the relative parameter error variance reduction can be formulated for parameter uncertainty. This denotes the ability (or otherwise) of a history-matching dataset to reduce the uncertainty of a particular parameter.

Let the i 'th diagonal element of $C'(\mathbf{k})$ be denoted by σ_{ui}^2 . Let σ_i^2 denote its prior uncertainty variance, this being the corresponding diagonal element of $C(\mathbf{k})$. We denote the relative uncertainty variance reduction of parameter i as ruv_i . This can be written as

$$ruv_i = 1 - \frac{\sigma_{ui}^2}{\sigma_i^2} \quad (10.3.8)$$

Like parameter identifiability, ruv_i ranges between zero and one. Unlike relative parameter error variance reduction, it cannot be negative. It is a more robust statistic than either identifiability or relative parameter error variance reduction as it does not depend on selection of a singular value at which the subdivision of parameter space into solution and null subspaces is deemed to occur. Furthermore, where parameters have spatial connotation (for example pilot points) it is worth plotting on a map. This allows modelling stakeholders to understand at a glance those parts of a study area where history-matching has been most and least successful in illuminating properties of the subsurface.

10.4 Model Predictions

10.4.1 Formulation of Equations

The focus of this chapter so far has been on the parameter estimation process (or the parameter conditioning process, depending on whether one adopts a subspace or Bayesian perspective). It has been demonstrated that, using linear analysis, it is relatively easy to derive a suite of statistics which reflect the capacity or otherwise of a measurement dataset to inform model parameters.

Of course, decision-support models are rarely built simply for the purpose of estimating system properties. Most models are built to make predictions of future environmental behaviour under similar or altered stress regimes to those which prevailed at the time of model history-matching. So it is to the making of model predictions that attention is now turned.

From (10.1.2) and (10.3.4) the posterior uncertainty of a prediction s can be evaluated as

$$\sigma_s^2 = \mathbf{y}^t \mathbf{C}(\mathbf{k}) \mathbf{y} - \mathbf{y}^t \mathbf{C}(\mathbf{k}) \mathbf{Z}^t [\mathbf{Z} \mathbf{C}(\mathbf{k}) \mathbf{Z}^t + \mathbf{C}(\boldsymbol{\varepsilon})]^{-1} \mathbf{Z} \mathbf{C}(\mathbf{k}) \mathbf{y}^t \quad (10.4.1)$$

From (10.3.5) this can also be written as

$$\sigma_s^2 = \mathbf{y}^t [\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k})]^{-1} \mathbf{y} \quad (10.4.2)$$

From (10.2.25) the post-calibration error variance of this same prediction can be written as

$$\sigma_{s-s}^2 = \mathbf{y}^t \mathbf{E} \mathbf{F}^{1/2} \mathbf{V}_2 \mathbf{V}_2^t \mathbf{F}^{1/2} \mathbf{E}^t \mathbf{y} + \mathbf{y}^t \mathbf{E} \mathbf{F}^{1/2} \mathbf{V}_1 \mathbf{S}^{-2} \mathbf{V}_1^t \mathbf{F}^{1/2} \mathbf{E}^t \mathbf{y} \quad (10.4.3)$$

It must not be forgotten that in equation (10.4.3) the \mathbf{U} , \mathbf{S} and \mathbf{V} matrices are calculated through singular value decomposition undertaken on the \mathbf{Y} matrix of equation (10.2.16). This matrix accounts for weight transformation of observations and Kahunen-Loëve transformation of parameters. Obviously, equation (10.4.3) is more difficult to evaluate than either of (10.4.2) and 10.4.1. Also, as will be demonstrated shortly, posterior error variance is a less efficient statistic than variance of posterior uncertainty. However it is an appropriate statistic to use where model calibration is less than optimal.

10.4.2 Effect of Transformation

The importance of weight-transformation of measurements and Kahunen-Loëve transformation of parameters in minimizing the error variances of estimated parameters and predictions made by a model that is calibrated using truncated singular value decomposition

has been discussed extensively in this text. It is because of pre-calibration Kahunen-Loëve transformation of parameters that the \mathbf{E} and \mathbf{F} matrices that emerge from singular value decomposition of the $\mathbf{C}(\mathbf{k})$ prior covariance matrix appear in equation (10.4.3). Expert knowledge that is encapsulated in $\mathbf{C}(\mathbf{k})$ is thereby granted entry into the inversion process.

It was demonstrated in section 8.2.4.2 that orthogonal projection in the space of Kahunen-Loëve transformed parameters (i.e. \mathbf{m} -space) does not constitute orthogonal projection in native parameter space (i.e. \mathbf{k} -space). If solution to the inverse problem is therefore undertaken in \mathbf{k} -space using singular value decomposition, the estimated $\underline{\mathbf{k}}$ vector, once transformed to \mathbf{m} -space, may have a non-zero projection onto the null space of the model as it operates in that space. The minimum error variance status of the solution to the inverse problem is thereby lost. The potential for parameter and predictive error is thereby increased; however it remains quantifiable.

Similar considerations do not apply to parameter and predictive uncertainty that is explored using Bayesian methods. Of course, as will be discussed in the next chapter, implementation of Bayesian methods entails many approximations that have their own cost in terms of analysis-incurred bias and potential for predictive error that is not included in quantified uncertainty. However, we demonstrate here that, if undertaken in a linear world using no incorrect assumptions, there is no need to subject parameters to Kahunen-Loëve transformation based on $\mathbf{C}(\mathbf{k})$ prior to Bayesian analysis, as the $\mathbf{C}(\mathbf{k})$ matrix is already built into Bayes equation. However, there is no loss in the integrity of Bayesian analysis if this is done.

From (10.3.7) applied in \mathbf{m} -space, together with (10.2.12) and (10.2.18), the posterior mean of Kahunen-Loëve transformed parameters can be calculated as

$$\underline{\mathbf{m}} = [\mathbf{Y}^t \mathbf{Y} + \mathbf{I}]^{-1} \mathbf{Y}^t \mathbf{g} \quad (10.4.4)$$

With pertinent substitutions for $\underline{\mathbf{m}}$, \mathbf{Y} and \mathbf{g} , this becomes

$$\underline{\mathbf{k}} = \mathbf{E} \mathbf{F}^{1/2} [\mathbf{F}^{1/2} \mathbf{E}^t \mathbf{Z}^t \mathbf{Q} \mathbf{Z} \mathbf{E} \mathbf{F}^{1/2} + \mathbf{I}]^{-1} \mathbf{F}^{1/2} \mathbf{E}^t \mathbf{Z}^t \mathbf{Q}^{1/2} \mathbf{Q}^{1/2} \mathbf{h} \quad (10.4.5)$$

From (10.2.10)

$$\mathbf{I} = \mathbf{F}^{1/2} \mathbf{E}^t \mathbf{C}^{-1}(\mathbf{k}) \mathbf{E} \mathbf{F}^{1/2} \quad (10.4.6)$$

If this and (10.2.1) are substituted into (10.4.5) we obtain

$$\underline{\mathbf{k}} = \mathbf{E} \mathbf{F}^{1/2} \{ \mathbf{F}^{1/2} \mathbf{E}^t [\mathbf{Z}^t \mathbf{Q} \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k})] \mathbf{E} \mathbf{F}^{1/2} \}^{-1} \mathbf{F}^{1/2} \mathbf{E}^t \mathbf{Z}^t \mathbf{C}^{-1}(\mathbf{e}) \mathbf{h} \quad (10.4.7)$$

Applying (2.6.4) and recalling that \mathbf{E} is an orthonormal matrix, this becomes

$$\underline{\mathbf{k}} = \mathbf{E} \mathbf{F}^{1/2} \mathbf{F}^{-1/2} \mathbf{E}^t [\mathbf{Z}^t \mathbf{Q} \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k})]^{-1} \mathbf{E} \mathbf{F}^{-1/2} \mathbf{F}^{1/2} \mathbf{E}^t \mathbf{Z}^t \mathbf{C}^{-1}(\mathbf{e}) \mathbf{h} \quad (10.4.8)$$

from which (10.3.7) follows.

A similar procedure can be followed to show that the computed uncertainty of a prediction is also immune to Kahunen-Loëve and weight transformation. We start with (10.4.2). If measurement dataset conditioning is done in \mathbf{m} -space we obtain

$$\sigma^2_s = \mathbf{w}^t [\mathbf{Y}^t \mathbf{Y} + \mathbf{I}]^{-1} \mathbf{w} \quad (10.4.9)$$

where \mathbf{w} expresses the sensitivity of a prediction s to parameters \mathbf{m} . From (10.1.2) and (10.2.13)

$$\mathbf{w} = \mathbf{F}^{1/2} \mathbf{E}^t \mathbf{y} \quad (10.4.10)$$

Through expansion of \mathbf{w} and \mathbf{Y} , equation (10.4.9) becomes

$$\sigma^2_s = \mathbf{y}^t \mathbf{E} \mathbf{F}^{1/2} [\mathbf{F}^{1/2} \mathbf{E}^t \mathbf{Z}^t \mathbf{Q}^{1/2} \mathbf{Q}^{1/2} \mathbf{Z} \mathbf{E} \mathbf{F}^{1/2} + \mathbf{I}]^{-1} \mathbf{F}^{1/2} \mathbf{E}^t \mathbf{y} \quad (10.4.11)$$

From (10.4.6) and (10.2.1) this becomes

$$\sigma^2_s = \mathbf{y}^t \mathbf{E} \mathbf{F}^{1/2} \{ \mathbf{F}^{1/2} \mathbf{E}^t [\mathbf{Z}^t \mathbf{C}^{-1}(\mathbf{\epsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k})] \mathbf{E} \mathbf{F}^{1/2} \}^{-1} \mathbf{F}^{1/2} \mathbf{E}^t \mathbf{y} \quad (10.4.12)$$

From (2.6.4) we then obtain

$$\sigma^2_s = \mathbf{y}^t \mathbf{E} \mathbf{F}^{1/2} \mathbf{F}^{-1/2} \mathbf{E}^t [\mathbf{Z}^t \mathbf{C}^{-1}(\mathbf{\epsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k})]^{-1} \mathbf{E} \mathbf{F}^{-1/2} \mathbf{F}^{1/2} \mathbf{E}^t \mathbf{y} \quad (10.4.13)$$

from which (10.4.2) follows.

10.4.3 Error and Uncertainty

Equation (10.4.3) can be written as

$$\sigma^2_{s-s} = \mathbf{w}^t \mathbf{V}_2 \mathbf{V}_2^t \mathbf{w} + \mathbf{w}^t \mathbf{V}_1 \mathbf{S}^{-2} \mathbf{V}_1^t \mathbf{w} \quad (10.4.14)$$

Equation (10.4.9) is the equation for the variance of predictive uncertainty under equivalent conditions to those described above. It was shown above that parameter and predictive uncertainty are the same, regardless of whether a measurement dataset conditions native or transformed parameters. Hence parameter and predictive uncertainty are functions solely of expert knowledge and the information content of a history-matching dataset. Unlike parameter and predictive error variance, they do not depend on how parameter uniqueness is sought; nor do they depend on regularization details such as the assumed dimensionality of the solution space or the target measurement objective function. (We repeat that Bayesian methods are burdened with their own assumptions, some of which will be discussed later in this book.)

Parameter and predictive error variance are minimized when regularization is undertaken using singular value decomposition applied to estimation of Kahanen-Loëve transformed parameters. At the same time, the model matrix to which singular value decomposition is applied must be weight-matrix-transformed, with the weight matrix being proportional to the inverse of the covariance matrix of measurement noise. Equation (10.4.14) depicts the error variance of a prediction calculated under such conditions.

Equation (10.4.14) can be written as

$$\sigma^2_{s-s} = \sum_{i=1}^w s_i^{-2} (\mathbf{w}^t \mathbf{v}_i)^2 + \sum_{i=w+1}^m (\mathbf{w}^t \mathbf{v}_i)^2 \quad (10.4.15)$$

where m is the number of parameters, w is the number of pre-truncation singular values and \mathbf{v}_i is the i 'th column of the \mathbf{V} matrix. Note that the first and second terms on the right of the equality sign have been interchanged between equations (10.4.14) and (10.4.15) in order to preserve the order of summation.

Now let us examine equation (10.4.9). From (10.2.21) and the fact that

$$\mathbf{V} \mathbf{V}^t = \mathbf{I} \quad (10.4.16)$$

(10.4.9) can be written as

$$\sigma^2_s = \mathbf{w}^t [\mathbf{V} \mathbf{S}^2 \mathbf{V}^t + \mathbf{V} \mathbf{V}^t]^{-1} \mathbf{w} \quad (10.4.17)$$

Hence

$$\sigma^2_s = \mathbf{w}^t [\mathbf{V} (\mathbf{S}^2 + \mathbf{I}) \mathbf{V}^t]^{-1} \mathbf{w} \quad (10.4.18)$$

$(\mathbf{S}^2 + \mathbf{I})$ is a diagonal matrix. The i 'th term along its diagonal is $s_i^2 + 1$. Its inverse is thus also a diagonal matrix. The i 'th term along the diagonal of its inverse is $1/(s_i^2 + 1)$. Equation (10.4.18) therefore becomes

$$\sigma^2_s = \mathbf{w}^t [\mathbf{V} (\mathbf{S}^2 + \mathbf{I})^{-1} \mathbf{V}^t] \mathbf{w} \quad (10.4.19)$$

This can be written in the same form as (10.4.15) as

$$\sigma'_s^2 = \sum_{i=1}^m \frac{(\mathbf{w}^t \mathbf{v}_i)^2}{(s_i^2 + 1)} \quad (10.4.20)$$

The scalar multiplier of each term of (10.4.20) is always less than or equal to 1 (it is equal to 1 where the singular value is equal to 0 for a particular i), and always less than s_i^{-2} . Hence every term in (10.4.20) is less than or equal to the corresponding term in (10.4.15). Posterior predictive error variance is therefore greater than posterior predictive uncertainty variance. Even the minimized value of predictive error variance (which occurs where singular value truncation is optimal) is therefore greater than the uncertainty of the prediction. Posterior uncertainty is a function purely of prior uncertainty and of the capacity of information contained in a measurement dataset to reduce it. It is not a function of how this information is harvested.

In section 10.2.2.3 it was shown that where Kahunen-Loëve and weight transformation take place, singular value truncation should exclude singular values below 1; these should be set to zero in (10.4.20). Where singular values are large, the terms of equation (10.4.20) are almost equal to the respective terms of (10.4.15). Where singular values are small, the terms of equation (10.4.20) are again almost equal to the respective terms of (10.4.15); where singular values are zero, exact equality prevails. Significant differences between the corresponding terms of the two equations exist only where singular values that are close to unity. Hence the difference between predictive error variance and predictive uncertainty variance is an outcome of the fact that calibration based on singular value decomposition creates a sharp cutoff between acceptance and rejection of information from a calibration dataset. In contrast, for Bayesian analysis, the transition between acceptance and rejection of information is smoother as access to information becomes more and more difficult because of measurement noise.

To see why predictive error variance needs to be higher than predictive uncertainty when calibration is undertaken using singular value decomposition, we continue the above analysis. Consider the estimate of the conditional mean for \mathbf{m} obtained using equation (10.4.4). After subjecting \mathbf{Y} to singular value decomposition this equation becomes

$$\underline{\mathbf{m}} = [\mathbf{V}\mathbf{S}^2\mathbf{V}^t + \mathbf{I}]^{-1}\mathbf{V}\mathbf{S}\mathbf{U}^t\mathbf{g} \quad (10.4.21)$$

Now let us substitute (10.2.14) for \mathbf{g} while expanding terms inside the inverted matrix.

$$\underline{\mathbf{m}} = [\mathbf{V}(\mathbf{S}^2 + \mathbf{I})\mathbf{V}^t]^{-1}\mathbf{V}\mathbf{S}\mathbf{U}^t\mathbf{U}\mathbf{S}\mathbf{V}^t\mathbf{m} + [\mathbf{V}(\mathbf{S}^2 + \mathbf{I})\mathbf{V}^t]^{-1}\mathbf{V}\mathbf{S}\mathbf{U}^t\xi \quad (10.4.22)$$

This becomes

$$\underline{\mathbf{m}} = \mathbf{V}(\mathbf{S}^2 + \mathbf{I})^{-1}\mathbf{S}^2\mathbf{V}^t\mathbf{m} + \mathbf{V}(\mathbf{S}^2 + \mathbf{I})^{-1}\mathbf{S}\mathbf{U}^t\xi \quad (10.4.23)$$

which can be written as

$$\underline{\mathbf{m}} = \sum_{i=1}^m \frac{s_i^2}{(1+s_i^2)} \mathbf{v}_i \mathbf{v}_i^t \mathbf{m} + \sum_{i=1}^m \frac{1}{(1+s_i^2)} \mathbf{v}_i \mathbf{u}_i^t \xi \quad (10.4.24)$$

As the terms that comprise the second summation are random, the second term is random. Where singular values are large, its effect on $\underline{\mathbf{m}}$ is small. As singular values diminish in value the effects of these terms grow but are limited because the factor that is applied to these terms grows to an asymptotic value of 1.

Terms of the first summation are large where singular values are large. This means that pertinent real-world components of \mathbf{m} are transferred directly to $\underline{\mathbf{m}}$ with little attenuation. However as singular values diminish, pertinent components of \mathbf{m} are attenuated as they make

their way into $\underline{\mathbf{m}}$. When singular values become zero, $\underline{\mathbf{m}}$ assimilates nothing of \mathbf{m} . However, until this point the transition from total assimilation to zero assimilation is gradual. Contrast this with the “all or nothing” approach of singular value decomposition whereby pertinent components of \mathbf{m} are either entirely assimilated into $\underline{\mathbf{m}}$ (but perturbed by measurement noise), or are ignored. In fact $\underline{\mathbf{m}}$ obtained using singular value decomposition and $\underline{\mathbf{m}}$ obtained using Bayes equation can only be the same if measurement noise is zero and singular value truncation takes place at the point where singular values become zero. In this case the “1” is omitted from the denominator of the first summation of (10.4.24) and the second term is absent. It follows that estimates of $\underline{\mathbf{m}}$ based on singular value decomposition suffer bias in the presence of measurement noise. It is therefore necessary for post-calibration parameter/predictive error variance to exceed posterior parameter/predictive uncertainty; this ensures that this potential bias incurred through model calibration is spanned by error bounds. The same applies to any form of regularized inversion. Fortunately, these effects are generally small compared with overall parameter and predictive uncertainty.

10.5 Data Worth Analysis

In a linear multiGaussian world, the posterior uncertainty of a model prediction can be calculated using either of equations (10.4.1) or (10.4.2). Use of these equations in the real world requires that a Jacobian matrix be calculated to replace the matrix \mathbf{Z} that appears in them. It also requires that a predictive sensitivity vector \mathbf{y} be calculated. A covariance matrix expressing expert knowledge, namely $\mathbf{C}(\mathbf{k})$, and a covariance matrix expressing measurement noise, namely $\mathbf{C}(\boldsymbol{\varepsilon})$, must also be supplied.

There is no reason, of course, why a “prediction” cannot be a parameter. All that is required of \mathbf{y} under these circumstances is that its only non-zero element be that which corresponds to the parameter in question; this element should be 1.

Neither equation (10.4.1) nor equation (10.4.2) requires the values of parameters. Similarly, neither of these equations requires the values of observations, nor of the prediction whose uncertainty is being evaluated. Hence, for a linear model, it is only sensitivities and covariances that affect posterior predictive uncertainty.

This has an important repercussion. It is not necessary for the history-matching process that is expressed through the matrix \mathbf{Z} to have actually been accomplished for the posterior uncertainty of a prediction to be evaluated. Nor, in fact, do the measurements comprising a history-matching dataset need to have even been made. All that is required for evaluation of parameter and predictive uncertainty is that the sensitivities comprising the matrix \mathbf{Z} and the vector \mathbf{y} be computed.

If the worth of data is judged by their ability to reduce the uncertainties of predictions of management interest, then data worth assessment based on linear analysis becomes a comparatively simple task. The worth of data with respect to a particular prediction can be assessed before these data are gathered. The comparative worth of different data acquisition strategies can be assessed before these strategies are actually implemented. The uncertainties of one or a number of predictions of management interest can be calculated using the above equations with and without the addition of extra data to an existing history-matching dataset. The extent to which uncertainties of predictions of interest are reduced through acquisition of these data can thereby be assessed before money and time are invested in their acquisition.

Similar strategies can be implemented in order to evaluate the worth of data that already exists. One option for evaluating the worth of an existing data element, or of all data of a certain type,

is to employ equation (10.4.1) or (10.4.2) to evaluate the uncertainty of a prediction of interest under the assumption that these data comprise the entirety of a history-matching dataset. The amount by which the uncertainty of that prediction is reduced below its prior uncertainty (the latter is calculated using the first term on the right of equation 10.4.1) is a measure of the worth of that data to the prediction.

An alternative way to assess the worth of existing data is to calculate the uncertainty of a prediction of interest with and without inclusion of that data in the existing history-matching dataset. The extent to which the uncertainty of the prediction rises with omission of certain data from the existing dataset is a measure of the uniqueness of their information with respect to that prediction. See Dausman et al (2010), Suescun et al (2025) and references cited therein for further discussion and practical applications.

Data whose worth is assessed need not be restricted to measurements of system state. Direct measurements of system properties can also be included in the analysis. A row of the \mathbf{Z} matrix for data of this type is comprised of zeros except for the element pertaining to the parameter whose value is measured. This element has a value of 1.

An advantage of linear data worth analysis is the ease with which it is implemented. A disadvantage is the linearity assumption on which it is based. Hence linear data worth analysis can only be approximate. Nevertheless, it can still be useful. In the example provided by Dausman et al (2010), the prediction of interest was the future position of a coastal saltwater wedge. Linear analysis demonstrated that the worth of new data in predicting the future location of the salt water wedge is highly dependent on the current location of this wedge; this, in turn, is highly dependent on system hydraulic properties. This however, does not invalidate the utility of linear data worth analysis. Repetition of the analysis for a few hydraulic property realizations demonstrated that unless head or concentration measurements are made close to the wedge, their worth is diminished. At the same time it was able to readily demonstrate the greater worth of concentration measurements than of head measurements.

In a later chapter of this book the use of data space inversion in data worth assessment is explored. In certain circumstances, the numerical cost of DSI-based data worth analysis is lower than that of linear data worth analysis; see section 14.3.3. DSI also provides opportunities for nonlinear data worth analysis.

10.6 Parameter Contributions to Predictive Uncertainty

Let us define “the contribution that an individual parameter makes to the uncertainty of a prediction” as the fall in the uncertainty of the prediction that would be accrued through the acquisition of perfect knowledge of the value of that parameter. A similar definition can be made for the contribution that a group of parameters makes to the uncertainty of a prediction. These contributions are readily evaluated using equation (10.4.1) or (10.4.2) in one of two ways. One way is to add rows to the \mathbf{Z} matrix simulating direct measurements of system properties in the manner described in the previous section; zero measurement noise would be associated with such measurements. An alternative procedure is to re-calculate the uncertainty of the prediction with the number of elements of \mathbf{k} reduced to account for the reduced number of unknown parameters; $C(\mathbf{k})$ requires conditioning if there is prior correlation between parameters whose values are now notionally known and those whose values are still unknown.

If desired, the contributions that individual or grouped parameters make to the uncertainty of a prediction can be graphed or plotted in space. Inspection of such plots can be highly informative.

Parameter contribution analysis of this type may help to resolve arguments about such issues as correctness or proximity of boundary conditions, and relevance of inadequate knowledge of historical system stresses. “Parameters” that were not included in a model history-matching process can be included in parameter contribution analysis. These parameters may include boundary value properties and/or the values of historical stresses that were, for the sake of convenience, assumed to be known when model history-matching was carried out. Appropriate prior uncertainties can be attributed to these new parameters when implementing linear uncertainty analysis.

If the contributions that parameters such as these make to the uncertainties of predictions of management interest are small, then possible errors in boundary conditions or historical stresses that they represent will have been shown to contribute little to the potential for error that is incurred in making these predictions. Approximations that were made in designing these boundary conditions and implementing these stresses will therefore have been shown to be acceptable. In the meantime, these approximations may have enabled the area of the model domain to be reduced, or have enhanced the numerical stability of the model. This, in turn, may have allowed history-matching and uncertainty analysis to proceed where it would have otherwise been difficult or impossible. The losses incurred by making these approximations will have been shown to be smaller than the gains that they enabled.

Parameter contributions to predictive uncertainty can be readily computed under both prior and posterior conditions. In the former case only the first term of equation (10.4.1) is employed in calculation of uncertainty.

In many cases it is found that the post-history-matching contribution of a parameter, or parameter group, to the uncertainty of a prediction is less than its pre-history-matching contribution. The size of the parameter contribution reduction may illustrate why history-matching was successful in reducing the uncertainty of the prediction. This is an indication that the prediction of interest is sensitive to the parameter or parameter group whose posterior contribution to predictive uncertainty is significantly less than its prior contribution.

In other cases the contribution that a parameter makes to the uncertainty of a prediction is actually seen to increase, rather than decrease, through the history-matching process. In explaining this seemingly paradoxical situation, the definition given above for “contribution to uncertainty” must be remembered. Sometimes a parameter makes little contribution to the uncertainty of a prediction prior to the history-matching process because the prediction is not sensitive to that parameter. Meanwhile, the information contained within a history-matching dataset may be shared between this parameter and another parameter. If this is the case, posterior correlation between the two parameters may be high for they share the same information. In a case such as this, acquisition of knowledge of the first parameter will reduce the posterior uncertainty of the second parameter. If the model prediction of interest is sensitive to the second parameter, then its uncertainty will be reduced through acquisition of further knowledge of the first parameter, notwithstanding the lack of direct sensitivity of the prediction to the first parameter.

11. Basics of Nonlinear Uncertainty Analysis

11.1 Introduction

11.1.1 General

Direct use of equations that are presented in the previous chapter of this book for exploration of the uncertainties of decision-critical model predictions is generally infeasible for two reasons. The first is that relationships between model outputs and field measurements are usually nonlinear, rather than linear. The second is that all of these equations have at their core a matrix \mathbf{Z} which is likely to have many columns because of the large number of parameters that are employed by a complex, spatial model such as a groundwater model. As we shall see, nonlinear uncertainty analysis, like inversion, is an iterative process. The filling, and then re-filling, of a sensitivity matrix to replace the \mathbf{Z} matrix during multiple iterations is expensive. It follows that, in order to be numerically tractable, nonlinear uncertainty analysis must work in a subspace of parameter space. A challenge that is faced by nonlinear uncertainty analysis is that of ensuring that this subspace is large enough for outcomes of the analysis to have integrity.

11.1.2 Bayesian Analysis

Where a model is complex and nonlinear, it is not possible to derive an analytical expression for the posterior probability distribution, or even for the posterior covariance matrix, of its parameters. The same applies to model predictions that are sensitive to these parameters. Instead, knowledge is gained of these posterior probability distributions by sampling them.

As discussed in section 6.3 of this book, despite their Bayesian credentials, methods such as rejection sampling and Markov chain Monte Carlo (MCMC) are infeasible where parameter numbers are large. So application of Bayes equation must be approximate. This should not be a cause for disappointment. After all, most other aspects of environmental modelling are approximate. This includes definition of a prior parameter probability distribution, and characterization of measurement noise. So if the two terms on the right side of Bayes equation are compromised, it does not matter too much if its implementation is also somewhat compromised, as long as the analysis is not misleading. Nevertheless, misleading outcomes of methodologies that are discussed herein are possible. So a modeller should be aware of their limitations.

For context, we repeat Bayes equation as it pertains to parameters of a model; this is equation (6.2.4).

$$f(\mathbf{k}|\mathbf{h}) = \frac{f(\mathbf{h}|\mathbf{k})f(\mathbf{k})}{f(\mathbf{h})} \quad (11.1.1a)$$

or more simply

$$f(\mathbf{k}|\mathbf{h}) \propto f(\mathbf{h}|\mathbf{k})f(\mathbf{k}) \quad (11.1.1b)$$

As discussed in chapter 6, $f(\mathbf{k})$ denotes the prior probability distribution of model parameters while $f(\mathbf{h}|\mathbf{k})$ is the likelihood function. $f(\mathbf{k}|\mathbf{h})$ is the posterior probability distribution of model parameters. That is, it is the probability distribution of parameters, conditional on measurements embodied in the vector \mathbf{h} . In plain terms, equation (11.1.1) states that for a vector \mathbf{k}_i to be considered as a sample of the posterior parameter probability distribution, it must satisfy two constraints. The first is that it be considered as a member of the prior parameter probability distribution. The second is that if the model is run using these parameters, model-

to-measurement misfit must be commensurate with the level of noise that is associated with these measurements. $f(\mathbf{h}|\mathbf{k})$ is the probability distribution of measurement noise.

This plain language description of Bayes equation can be considered as a “mission statement” as far as sampling of the posterior parameter probability is concerned. It even suggests a way of achieving this mission.

Conceptually, it is not difficult to sample the prior parameter probability distribution; see chapter 4. There are many ways to do this, ranging from complex geological process emulators to use of a stationary covariance matrix. The next task that is required for implementation of Bayes equation is to modify these sampled parameter fields so that pertinent model outputs are compatible with field measurements. This is a history-matching problem. Ways in which history-matching problems can be solved in highly parameterized contexts have been discussed extensively in previous chapters of this book. A pre-condition for their solution, however, is that changes in model outputs be continuous with respect to changes in parameters. This restricts the types of parameter fields that can be adjusted, and the means through which their adjustment can be implemented. It may require that model hydraulic properties and stresses be represented in ways that are somewhat abstract when compared with their real-world counterparts. Nevertheless, this still leaves us with many options for their realistic representation, including use of nonstationary correlation functions that are discussed in chapter 5.

A multiGaussian probability distribution is completely characterized by its mean and covariance matrix. Let us suppose that the parameters of a particular model possess a multiGaussian probability distribution. Let us further suppose that the mean of this distribution is $\mathbf{0}$ because parameters are defined as perturbations from this prior mean. Let $\mathbf{C}(\mathbf{k})$ represent the prior covariance matrix of model parameters. Let us also suppose that measurement noise is characterized by a multiGaussian distribution with zero mean; we characterize its covariance matrix as $\mathbf{C}(\boldsymbol{\varepsilon})$. And let us assume a linear model. We repeat the familiar equation that characterizes the action of this model on its parameters using the matrix \mathbf{Z} .

$$\mathbf{h} = \mathbf{Z}\mathbf{k} + \boldsymbol{\varepsilon} \quad (11.1.2)$$

We have already seen the outcomes of Bayesian analysis applied to this situation. The posterior mean parameter vector $\underline{\mathbf{k}}$ can be calculated using either of two equivalent formulas which are linked by the Sherman-Morrison-Woodbury identity.

$$\underline{\mathbf{k}} = (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k}))^{-1} \mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{h} \quad (11.1.3a)$$

$$\underline{\mathbf{k}} = \mathbf{C}(\mathbf{k}) \mathbf{Z}^t [\mathbf{Z} \mathbf{C}(\mathbf{k}) \mathbf{Z}^t + \mathbf{C}(\boldsymbol{\varepsilon})]^{-1} \mathbf{h} \quad (11.1.3b)$$

These expressions for $\underline{\mathbf{k}}$ are obtained by minimizing the following objective function.

$$\Phi = (\mathbf{h} - \mathbf{Z}\mathbf{k})^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) (\mathbf{h} - \mathbf{Z}\mathbf{k}) + \mathbf{k}^t \mathbf{C}^{-1}(\mathbf{k}) \mathbf{k} \quad (11.1.4)$$

Meanwhile the posterior parameter covariance matrix can be calculated using either of the following equivalent formulas. These are also linked to each other through the Sherman-Morrison-Woodbury identities.

$$\mathbf{C}'(\mathbf{k}) = \mathbf{C}(\mathbf{k}) - \mathbf{C}(\mathbf{k}) \mathbf{Z}^t [\mathbf{Z} \mathbf{C}(\mathbf{k}) \mathbf{Z}^t + \mathbf{C}(\boldsymbol{\varepsilon})]^{-1} \mathbf{Z} \mathbf{C}(\mathbf{k}) \quad (11.1.5a)$$

$$\mathbf{C}'(\mathbf{k}) = [\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k})]^{-1} \quad (11.1.5b)$$

Equipped with these familiar formulas, we now commence our journey into approximate implementation of Bayes equation in highly parameterized, nonlinear model settings.

11.2 Randomized Maximum Likelihood

11.2.1 Methodology

We begin our journey into the realm of nonlinear uncertainty analysis by discussing the randomized maximum likelihood (RML) methodology. While this method finds common usage in disciplines such as weather and climate modelling, it is not often used in groundwater model uncertainty analysis because of its high numerical burden. We explore it here, however, because of the insights that it provides and because, in many ways, it is a gateway to ensemble methods.

The RML method is conceptually simple. Steps in its implementation are as follows.

1. Generate a random sample of the prior parameter probability distribution. We refer to this sample as \mathbf{k}_i .
2. Generate a random sample of measurement noise; add this to \mathbf{h} to form \mathbf{h}_i .
3. Using preferred-value Tikhonov regularization, calculate a parameter field that is as close as possible to \mathbf{k}_i while allowing model outputs to fit field data to within measurement noise. It is now a sample of the posterior parameter probability distribution. The sample of the prior parameter probability distribution has therefore been altered to the minimum extent required for it to respect field measurements of system behaviour.
4. Repeat steps 1 to 3 to obtain more samples of the posterior parameter probability distribution.

Step 3 is achieved by minimizing the following objective function. This is slightly modified from (11.1.4) in order to accommodate minimal departure of adjusted \mathbf{k} from \mathbf{k}_i .

$$\Phi = (\mathbf{h}_i - \mathbf{Z}\mathbf{k})^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) (\mathbf{h}_i - \mathbf{Z}\mathbf{k}) + (\mathbf{k} - \mathbf{k}_i)^t \mathbf{C}^{-1}(\mathbf{k}) (\mathbf{k} - \mathbf{k}_i) \quad (11.1.6)$$

Equation (11.1.3a) must be modified if it is to solve this minimization problem.

$$\mathbf{k}'_i = (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k}))^{-1} (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{h}_i + \mathbf{C}^{-1}(\mathbf{k}) \mathbf{k}_i) \quad (11.1.7)$$

The extra term in the second bracket on the right of (11.1.7) arises from the necessity for the solution parameter field \mathbf{k}'_i to fit the sample of the prior parameter field \mathbf{k}_i to within constraints set by the prior covariance matrix $\mathbf{C}(\mathbf{k})$ at the same time as model outputs respect noise-enhanced field measurements \mathbf{h}_i to within constraints set by the measurement noise covariance matrix $\mathbf{C}(\boldsymbol{\varepsilon})$.

Because of model nonlinearity, solution for each \mathbf{k}'_i is an iterative process that involves use of the Marquardt lambda in ways that have been discussed. Because this process is iterative, and because the \mathbf{Z} sensitivity matrix must be rebuilt during each iteration of this process, the numerical cost of obtaining a sample \mathbf{k}'_i of the posterior parameter probability distribution is high. The cost of obtaining many samples is prohibitive.

11.2.2 Respect for Bayes Equation

We now show that where a model is linear, \mathbf{k}'_i obtained in this way samples the posterior parameter probability distribution. See Blatter et al (2022) for more details. To do this, it is necessary to demonstrate that \mathbf{k}'_i has the correct expected value (in the statistical sense), and that it has the correct posterior covariance matrix. The first part is easy. We employ the symbol $E[]$ to denote expected value. From (11.1.7)

$$E[\mathbf{k}'_i] = (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k}))^{-1} (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) E[\mathbf{h}_i] + \mathbf{C}^{-1}(\mathbf{k}) E[\mathbf{k}_i]) \quad (11.1.8)$$

As usual, we assume that the prior mean of \mathbf{k} is $\mathbf{0}$. Hence $E[\mathbf{k}_i]$ is $\mathbf{0}$. Because the mean value of measurement noise is $\mathbf{0}$, $E[\mathbf{h}_i]$ is \mathbf{h} . Equation (11.1.8) therefore becomes (11.1.3). This is indeed $E[\mathbf{k}'_i]$.

We now examine the covariance matrix of \mathbf{k}'_i . This is the covariance matrix of samples obtained through repeated calculation of \mathbf{k}'_i using (11.1.7). First, expand (11.1.7) as

$$\mathbf{k}'_i = (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k}))^{-1} \mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{h}_i + (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k}))^{-1} \mathbf{C}^{-1}(\mathbf{k}) \mathbf{k}_i \quad (11.1.9)$$

Now invoke propagation of covariance (3.9.2) to obtain the covariance matrix of \mathbf{k}'_i over many i . At the same time we note that the covariance matrix of \mathbf{h}_i over many i is $\mathbf{C}(\boldsymbol{\varepsilon})$ while the covariance matrix of \mathbf{k}_i over many i is $\mathbf{C}(\mathbf{k})$.

$$\begin{aligned} \mathbf{C}(\mathbf{k}'_i) &= (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k}))^{-1} \mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{C}(\boldsymbol{\varepsilon}) \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k}))^{-1} \\ &\quad + (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k}))^{-1} \mathbf{C}^{-1}(\mathbf{k}) \mathbf{C}(\mathbf{k}) \mathbf{C}^{-1}(\mathbf{k}) (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k}))^{-1} \end{aligned} \quad (11.1.10)$$

That is

$$\mathbf{C}(\mathbf{k}'_i) = (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k}))^{-1} (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}(\mathbf{k})) (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k}))^{-1} \quad (11.1.11)$$

So that

$$\mathbf{C}(\mathbf{k}'_i) = (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k}))^{-1} \quad (11.1.12)$$

which is in agreement with (11.1.5b).

It follows that RML is a Bayesian method as it samples the posterior parameter probability distribution. However, as Evensen et al (2022) point out, this is only the case if a model is linear. Where a model is nonlinear, it still yields parameter fields that respect the prior parameter probability distribution and that respect field measurements. But these cannot be construed as providing a complete, or bias-free, set of samples of the posterior parameter probability distribution. However, as Bladde et al (2022) show, the set of parameter samples that it provides is generally good enough. This is fortunate, as use of methods such as MCMC whose Bayesian credentials are excellent, are impractical where parameter numbers are high.

Bladde et al (2022) extend the RML method to accommodate selection and updating of a regularization weight factor. They note that addition of a stochastic accept/reject step (similar to MCMC) improves the methodology's Bayesian credentials in the face of model nonlinearity.

11.3 Approximate and Numerically Cheap RML

A methodology is now briefly discussed that resembles RML but is numerically cheaper to implement. This is the so-called “null-space Monte Carlo” method that is supported by the PEST suite. This method is not as widely used today as it was in the past as it is not as efficient as ensemble methods. However it offers an alternative to ensemble methods if the latter are challenged by problem nonlinearity, with the attendant possibility of parameter/predictive bias or failure to achieve an adequate level of model-to-measurement fit.

Use of null space Monte Carlo rests on the premise that a model has just been calibrated, and that a Jacobian matrix is available. Ideally, model calibration has yielded a parameter field that is at, or close to, the posterior mean. A linear approximation to the posterior parameter covariance matrix can be calculated using the prior parameter covariance matrix and the Jacobian matrix; see equations (11.1.5). Armed with a posterior mean, an approximate posterior covariance matrix, and an assumption of multiGaussianity, samples can be drawn from a linear approximation to the posterior parameter probability distribution.

If the model is run using these samples, it is generally found that the objective function is higher than that attained during model calibration, by more (sometimes much more) than measurement noise allows (even if the structural component of “measurement noise” is taken into account). However model-to-measurement fit is generally much better than would be attained using samples of the prior parameter probability distribution. Adjustment of these approximate samples to satisfy history-matching constraints is not, therefore, nearly as burdensome as for standard RML.

Further numerical advantages can be gained if adjustment of all of these samples of the linearized posterior parameter probability distribution is based on the same Jacobian matrix for the first iteration of their adjustment. This can be the same Jacobian matrix as that on which calculation of the linearized posterior covariance matrix was based. If model outputs calculated using these adjusted samples do not all yield satisfactory model-to-measurement fits, then a second iteration of parameter adjustment may be required; however this requires calculation of new, sample-specific Jacobian matrices. Alternatively, some parameter fields may be rejected if one iteration of adjustment does not attain the desired level of model-to-measurement fit. Another option is to employ SVD-assist (see section 9.3) for parameter adjustment. This reduces the cost of a possible second iteration to a number of model runs that is commensurate with the dimensionality of the calibration solution space.

12 Iterative Ensemble Smoother

12.1 General

The literature on ensemble methods is vast. For reasons that will become obvious shortly, they have established a firm foothold for implementation of posterior uncertainty analysis of parameters and predictions made by large models in many fields of science and engineering, including groundwater, petroleum reservoir, geothermal and meteorology. They can be used not just for estimating the values of model parameters, but also for updating model-domain-wide estimates of system state based on measurements of system state at a relatively small number of locations. The number of variables that they can estimate and update is enormous.

In this book we focus only on the ensemble smoother. This is generally used to estimate the values of model parameters rather than system states. Instead of performing sequential data assimilation in which parameters and states are progressively upgraded as real and simulated time progresses, a smoother performs numerical tasks that bear some similarity to that of model calibration. That is, a model is run over a historical time period. Pertinent model outputs are compared with field measurements of system behaviour. Parameters are then adjusted in order for model outputs to achieve a suitable level of fit with these measurements. However an important difference between ensemble smoothing and inverse problem solving that has been discussed so far in this book is that multiple sets of parameters are adjusted simultaneously. These parameter sets start out as samples of the prior parameter probability distribution. After iterative adjustment, they morph to become samples of the posterior parameter probability distribution.

Parameter adjustment is iterative for an iterative ensemble smoother (IES) for the same reason that it is iterative for regularized inversion. This is to accommodate model nonlinearity. Recall that a model is nonlinear if sensitivities of model outputs to parameters change as the values of parameters change. However, IES does not fill or store a Jacobian matrix. Instead, it calculates and stores matrices which can be conceptualized as the actions of a Jacobian matrix on certain other matrices. As parameters are updated, these matrices must also be updated. Meanwhile, numerical expediencies such as the Marquardt lambda are deployed to hasten solution convergence that would otherwise be challenged by model nonlinearity.

Conceptually, IES performs the same task as the randomized maximum likelihood (RML) method. So RML is our starting point for deriving the equations on which IES is based. Like RML, IES has Bayesian roots. And like RML, IES promises sampling of the posterior probability distribution of parameters employed by a linear model. Unfortunately, like RML, the theoretical basis of IES is eroded by model nonlinearity. Hence posterior parameter fields that it delivers must be viewed with caution. Nevertheless, it has a well-proven ability to compute suites of diverse parameter fields that all allow a model to fit field measurements of system behaviour remarkably well with unparalleled numerical efficiency. The parameter fields that it yields can therefore be used to assess post-history-matching model predictive uncertainty.

The PESTPP-IES ensemble smoother is described by White (2018) and by its documentation. It, together with other members of the PEST++ suite, can be downloaded from its GitHub site at <https://github.com/usgs/pestpp>

12.2 Mathematical Basics

12.2.1 Simultaneous Parameter Adjustment

To see how ensemble-based parameter adjustment works, let us start by generating N samples \mathbf{k}_i (where i ranges from 1 to N) of the prior probability distribution of a model's parameters. We assign each of these \mathbf{k}_i vectors to successive columns of a matrix \mathbf{K} . \mathbf{K} therefore possesses m rows and N columns, where m is the number of elements of \mathbf{k} (i.e. the number of parameters with which a model is endowed). The task of IES is to derive a matrix \mathbf{K}' whose columns comprise vectors \mathbf{k}'_i that are samples of the posterior parameter probability distribution.

We now sprinkle N samples of measurement noise on the measurement dataset \mathbf{h} to yield N noise-enhanced vectors \mathbf{h}_i . These vectors then become the columns of a matrix \mathbf{H} . \mathbf{H} is thus an $n \times N$ matrix, where n is the number of elements in the history-matching dataset.

For now, we consider the model to be linear so that its action on parameters can be represented by the matrix \mathbf{Z} . Let $\mathbf{C}(\mathbf{k})$ denote the prior covariance matrix of parameters, and let $\mathbf{C}(\boldsymbol{\varepsilon})$ denote the covariance matrix of measurement noise. Equation (11.1.7) that is employed by the randomized maximum likelihood (RML) method can now be used to estimate N sets of parameters \mathbf{k}'_i . These can be assigned to the columns of matrix \mathbf{K}' so that

$$\mathbf{K}' = (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k}))^{-1} (\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{H} + \mathbf{C}^{-1}(\mathbf{k}) \mathbf{K}) \quad (12.2.1)$$

The first major difference between RML and IES is that equation (12.2.1) is solved for all columns of \mathbf{K}' using the same \mathbf{Z} matrix. Recall that \mathbf{Z} is a sensitivity (i.e. Jacobian) matrix. Where a model is nonlinear, equation (12.2.1) must be solved iteratively. However, despite the fact that many sets of parameters are adjusted, the matrix \mathbf{Z} is filled only once per iteration, and then used to adjust all parameter sets, regardless of the number of parameter sets that occupy the columns of \mathbf{K} and therefore \mathbf{K}' . Simultaneous adjustment of all of these parameter fields using a single sensitivity matrix comprises a huge numerical saving.

However further numerical efficiencies await.

12.2.2 Empirical Covariance Calculation

Though implied in the equations on which they are based, ensemble methods neither calculate nor store Jacobian matrices; recall that the Jacobian matrix is the cornerstone of regularized inversion that is discussed in previous chapters of this book. Instead, ensemble methods work with covariance matrices. These are calculated empirically.

Let the vector \mathbf{o} designate the model-calculated counterpart to the measurement vector \mathbf{h} . Hence, similar to equation (11.1.2) but without measurement noise

$$\mathbf{o} = \mathbf{Z}\mathbf{k} \quad (12.2.2a)$$

Realizations of \mathbf{o} are readily calculated from realizations of \mathbf{k} . These can be assigned to the columns of a matrix \mathbf{O} . Therefore

$$\mathbf{O} = \mathbf{Z}\mathbf{K} \quad (12.2.2b)$$

We now define a new matrix $\underline{\mathbf{K}}$ that is easily calculated from \mathbf{K} . All of the columns of $\underline{\mathbf{K}}$ are the same; they are occupied by a vector $\underline{\mathbf{k}}$. Each of the elements of $\underline{\mathbf{k}}$ is the average of all elements in the same row of \mathbf{K} . Therefore the i 'th element of $\underline{\mathbf{k}}$ is calculated as

$$\underline{k}_i = \frac{\sum_{j=1}^N k_{ij}}{N} \quad (12.2.3a)$$

or, to put it another way,

$$\underline{\mathbf{K}} = \frac{\mathbf{K}\mathbf{1}}{N} \quad (12.2.3b)$$

where $\mathbf{1}$ is an $N \times N$ matrix in which every element is 1. Similarly

$$\underline{\mathbf{Q}} = \frac{\mathbf{0}\mathbf{1}}{N} \quad (12.2.4)$$

Empirical covariance matrices are now easily calculated. The covariance matrix that links model outputs to parameters \mathbf{C}_{ok}^e is calculated as

$$\mathbf{C}_{\text{ok}}^e = \frac{(\mathbf{o} - \underline{\mathbf{o}})(\mathbf{k} - \underline{\mathbf{k}})^t}{N-1} \quad (12.2.5)$$

Similarly, an empirical $\mathbf{C}(\mathbf{k})$, which we denote as \mathbf{C}_{kk}^e , is calculated as

$$\mathbf{C}_{\text{kk}}^e = \frac{(\mathbf{k} - \underline{\mathbf{k}})(\mathbf{k} - \underline{\mathbf{k}})^t}{N-1} \quad (12.2.6)$$

while

$$\mathbf{C}_{\text{oo}}^e = \frac{(\mathbf{o} - \underline{\mathbf{o}})(\mathbf{o} - \underline{\mathbf{o}})^t}{N-1} \quad (12.2.7)$$

To see what these matrices can tell us about parameter sensitivities, consider the vector

$$\begin{bmatrix} \mathbf{o} \\ \mathbf{k} \end{bmatrix} = \begin{bmatrix} \mathbf{Z} \\ \mathbf{I} \end{bmatrix} \mathbf{k} \quad (12.2.8)$$

The empirical covariance matrix of $\begin{bmatrix} \mathbf{o} \\ \mathbf{k} \end{bmatrix}$ can be partitioned into submatrices as follows

$$\mathbf{C}^e \left(\begin{bmatrix} \mathbf{o} \\ \mathbf{k} \end{bmatrix} \right) = \begin{bmatrix} \mathbf{C}_{\text{oo}}^e & \mathbf{C}_{\text{ok}}^e \\ \mathbf{C}_{\text{ko}}^e & \mathbf{C}_{\text{kk}}^e \end{bmatrix} \quad (12.2.9)$$

From (12.2.8) this same covariance matrix can be calculated from the covariance matrix of \mathbf{k} using the standard formula for propagation of covariance (3.9.3).

$$\mathbf{C} \left(\begin{bmatrix} \mathbf{o} \\ \mathbf{k} \end{bmatrix} \right) = \begin{bmatrix} \mathbf{Z} \\ \mathbf{I} \end{bmatrix} \mathbf{C}(\mathbf{k}) \begin{bmatrix} \mathbf{Z}^t & \mathbf{I}^t \end{bmatrix} \quad (12.2.10)$$

Multiplying out

$$\mathbf{C} \left(\begin{bmatrix} \mathbf{o} \\ \mathbf{k} \end{bmatrix} \right) = \begin{bmatrix} \mathbf{ZC}(\mathbf{k})\mathbf{Z}^t & \mathbf{ZC}(\mathbf{k}) \\ \mathbf{C}(\mathbf{k})\mathbf{Z}^t & \mathbf{C}(\mathbf{k}) \end{bmatrix} \quad (12.2.11)$$

Comparing (12.2.9) and (12.2.11) it is apparent that

$$\mathbf{C}_{\text{oo}}^e = \mathbf{ZC}(\mathbf{k})\mathbf{Z}^t \quad (12.2.12a)$$

$$\mathbf{C}_{\text{ok}}^e = \mathbf{ZC}(\mathbf{k}) \quad (12.2.12b)$$

$$\mathbf{C}_{\text{ko}}^e = \mathbf{C}(\mathbf{k})\mathbf{Z}^t \quad (12.2.12c)$$

$$\mathbf{C}_{\text{kk}}^e = \mathbf{C}(\mathbf{k}) \quad (12.2.12d)$$

The last of these relationships is obvious.

Note, in particular, the second and third of the above equations. These give us the ability to calculate a sensitivity matrix from an empirical covariance matrix. For example, from (12.2.12b) and (12.2.12d)

$$\mathbf{Z} = C_{ok}^e [C_{kk}^e]^{-1} \quad (12.2.13)$$

Conceptually, this could be used in equation (12.2.1) to solve for \mathbf{K}' . The number of model runs that is required to compute this approximate \mathbf{Z} is equal to the number of columns of \mathbf{K} , \mathbf{H} and \mathbf{O} . That is, it is equal to the number of realizations N that comprise the ensemble (generally a few hundred), regardless of the number of parameters m that populate the model. This sounds too good to be true.

To be sure, the matrix \mathbf{Z} that emerges from (12.2.13) is a poor approximation to the real Jacobian matrix. It is rank-deficient. This is because its rank can only be as high as the number of realizations N that comprise the ensemble; it cannot be as high as the number of parameters for which sensitivities are required to fill the matrix \mathbf{Z} . Furthermore, the “inverse” of C_{kk}^e that is featured in equation (12.2.13) can only be the Moore-Penrose, generalized inverse of C_{kk}^e ; it is computed using truncated singular value decomposition. It too, has a diminished rank of N , whereas the true $C(\mathbf{k})$ is an $m \times m$ matrix with a full rank of m . Nevertheless, the rank-deficient Jacobian matrix that can be computed using equation (12.2.13) (or rather, as we shall see, matrices that are related to this matrix) is generally good enough for columns of \mathbf{K}' that are calculated using this matrix to support a good fit between model outputs and field measurements.

Now, let us modify equation (12.2.1) by ignoring the last term in the right brackets. This is the term that compels parameter adherence to initial parameter values, which are random parameter values drawn from the prior parameter probability distribution. Parameters continue to be adjusted from these random initial values; however, because this term is ignored, adherence to these initial values is not balanced against model-to-measurement fit. (This term is often omitted in ensemble deployment; it simplifies the mathematics of ensembles, and can greatly reduce the number of model runs required for data assimilation. So we omit it now, but reinstate it later.) Equation (12.2.1) therefore becomes

$$\mathbf{K}' = (\mathbf{Z}^t C^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + C^{-1}(\mathbf{k}))^{-1} \mathbf{Z}^t C^{-1}(\boldsymbol{\varepsilon}) \mathbf{H} \quad (12.2.14)$$

Using the Sherman-Morrison-Woodbury formula (2.15.6), this can be re-written as

$$\mathbf{K}' = C(\mathbf{k}) \mathbf{Z}^t (Z C(\mathbf{k}) \mathbf{Z}^t + C(\boldsymbol{\varepsilon}))^{-1} \mathbf{H} \quad (12.2.15)$$

Now, making substitutions from (12.2.12), \mathbf{K}' can be computed from the outcomes of empirical covariance analysis emerging from N model runs based on parameter realizations encapsulated in \mathbf{K} .

$$\mathbf{K}' = C_{ko}^e (C_{oo}^e + C(\boldsymbol{\varepsilon}))^{-1} \mathbf{H} \quad (12.2.16)$$

The repercussions of equation (12.2.16) are astounding. In a linear world, this equation allows calculation of N parameter sets encapsulated in the columns of \mathbf{K}' that sample the posterior parameter probability distribution. The cost of obtaining these samples is equal to the number of parameter realizations, that is N . There is no need to calculate a Jacobian matrix. Instead, matrix products of a notional Jacobian matrix with other matrices are assembled simply by calculating empirical covariances between model outputs and other model outputs, and between model outputs and parameters. Furthermore, because the N model runs on which these covariances are based are independent of each other (each employing a different prior parameter realization), these runs are parallelizable.

12.2.3 Alternative Formulation

From (11.1.5b), equation (12.2.14) can be re-written as

$$\mathbf{K}' = \mathbf{C}'(\mathbf{k}) \mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{H} \quad (12.2.17)$$

From (11.1.5a) $\mathbf{C}'(\mathbf{k})$ can also be formulated as

$$\mathbf{C}'(\mathbf{k}) = \mathbf{C}(\mathbf{k}) - \mathbf{C}(\mathbf{k}) \mathbf{Z}^t [\mathbf{Z} \mathbf{C}(\mathbf{k}) \mathbf{Z}^t + \mathbf{C}(\boldsymbol{\varepsilon})]^{-1} \mathbf{Z} \mathbf{C}(\mathbf{k}) \quad (12.2.18)$$

Making substitutions from (12.2.12) this becomes

$$\mathbf{C}'(\mathbf{k}) = \mathbf{C}_{kk}^e - \mathbf{C}_{ko}^e (\mathbf{C}_{oo}^e + \mathbf{C}(\boldsymbol{\varepsilon}))^{-1} \mathbf{C}_{ok}^e \quad (12.2.19)$$

Meanwhile, from (12.2.13)

$$\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) = \mathbf{C}_{kk}^{e-1} \mathbf{C}_{ko}^e \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \quad (12.2.20)$$

So that (12.2.17) becomes

$$\mathbf{K}' = \left[\mathbf{C}_{kk}^e - \mathbf{C}_{ko}^e (\mathbf{C}_{oo}^e + \mathbf{C}(\boldsymbol{\varepsilon}))^{-1} \mathbf{C}_{ok}^e \right] \mathbf{C}_{kk}^{e-1} \mathbf{C}_{ko}^e \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{H} \quad (12.2.21)$$

This equation is generally less convenient to use than (12.2.16).

12.3 Numerical Implementation of IES

12.3.1 General

Equations that are presented above provide an extremely convenient implementation of Bayes equation where a model is linear, and can therefore be represented by a matrix. Their Bayesian credentials also rest on the assumption that the prior probability distributions of \mathbf{k} and $\boldsymbol{\varepsilon}$ are both multiGaussian. If this is not the case, it may be possible to transform \mathbf{k} and or $\boldsymbol{\varepsilon}$ to make this so. If this is not possible, then parameter uncertainties provided by the above equations must be construed as approximate only. However, this does not diminish their usefulness, especially when the small numerical burden of their implementation is considered.

Of greater concern are the impacts of model nonlinearity. The impact of model nonlinearity on the Bayesian integrity of ensemble methods is unavoidable, and must therefore be tolerated. At the same time, steps must be taken to ensure good numerical performance of the history-matching process under nonlinear conditions. Some of these steps are similar to those that must be taken to ensure good performance of regularized inversion under these conditions. They include the necessity for the ensemble adjustment process to be iterative. Practical numerical devices such as the Marquardt lambda must also be employed. Strategic use of singular value decomposition also becomes necessary (see below). Other practical measures that are discussed later in this chapter can also ensure good numerical behaviour of ensemble-based history-matching in awkward nonlinear contexts. As stated above, the literature on ensemble methods is vast; other useful strategies that are not mentioned herein can also be deployed to enhance history-matching alacrity in difficult inversion settings.

To illustrate some numerical practicalities of ensemble deployment in concert with complex nonlinear models, a few details of the iterative ensemble algorithm developed by Chen and Oliver (2013) are now presented. This algorithm is used by the PESTPP-IES ensemble smoother.

12.3.2 The Chen and Oliver Algorithm

We start with equation (12.2.1). By expanding terms, this equation becomes

$$\mathbf{K}' = [\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k})]^{-1} \mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{H} + [\mathbf{Z}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k})]^{-1} \mathbf{C}^{-1}(\mathbf{k}) \mathbf{K} \quad (12.3.1)$$

Application of the Sherman-Morrison-Woodbury identity to the first term yields

$$\mathbf{K}' = \mathbf{C}(\mathbf{k})\mathbf{Z}^t[\mathbf{Z}\mathbf{C}(\mathbf{k})\mathbf{Z}^t + \mathbf{C}(\boldsymbol{\varepsilon})]^{-1}\mathbf{H} + [\mathbf{Z}^t\mathbf{C}^{-1}(\boldsymbol{\varepsilon})\mathbf{Z} + \mathbf{C}^{-1}(\mathbf{k})]^{-1}\mathbf{C}^{-1}(\mathbf{k})\mathbf{K} \quad (12.3.2)$$

We now introduce the Marquardt lambda, while noting that the resulting equation must be solved iteratively, and that some trial and error may be required to determine the optimal Marquardt lambda to use during each iteration of the history-matching process. In the following equation, we also replace \mathbf{K} with \mathbf{K}_0 , indicating that the prior probability distribution is sampled to fill the columns of this matrix.

$$\mathbf{K}' = \mathbf{C}(\mathbf{k})\mathbf{Z}^t[\mathbf{Z}\mathbf{C}(\mathbf{k})\mathbf{Z}^t + (1+\lambda)\mathbf{C}(\boldsymbol{\varepsilon})]^{-1}\mathbf{H} + [\mathbf{Z}^t\mathbf{C}^{-1}(\boldsymbol{\varepsilon})\mathbf{Z} + (1+\lambda)\mathbf{C}^{-1}(\mathbf{k})]^{-1}\mathbf{C}^{-1}(\mathbf{k})\mathbf{K}_0 \quad (12.3.3)$$

Next we substitute empirical covariance matrices in the first term of this equation, as was done in the preceding section.

$$\mathbf{K}' = \mathbf{C}_{ko}^e[\mathbf{C}_{oo}^e + (1+\lambda)\mathbf{C}(\boldsymbol{\varepsilon})]^{-1}\mathbf{H} + [\mathbf{Z}^t\mathbf{C}^{-1}(\boldsymbol{\varepsilon})\mathbf{Z} + (1+\lambda)\mathbf{C}^{-1}(\mathbf{k})]^{-1}\mathbf{C}^{-1}(\mathbf{k})\mathbf{K}_0 \quad (12.3.4)$$

$\mathbf{C}(\mathbf{k})$ in the above equation is the prior covariance matrix. However, in order to achieve computational conveniences that will become apparent in a moment, $\mathbf{C}(\mathbf{k})$ inside the second bracket is now replaced by an empirical covariance matrix that is calculated from parameter samples that pertain to the current iteration. This covariance matrix is calculated using equation (12.2.6). Equation (12.3.4) then becomes

$$\mathbf{K}' = \mathbf{C}_{ko}^e[\mathbf{C}_{oo}^e + (1+\lambda)\mathbf{C}(\boldsymbol{\varepsilon})]^{-1}\mathbf{H} + [\mathbf{Z}^t\mathbf{C}^{-1}(\boldsymbol{\varepsilon})\mathbf{Z} + (1+\lambda)\mathbf{C}_{kk}^{e-1}]^{-1}\mathbf{C}^{-1}(\mathbf{k})\mathbf{K}_0 \quad (12.3.5)$$

The presence of \mathbf{Z} in the second term is inconvenient. Calculation and storage of this large sensitivity matrix must be avoided. So equation (12.2.13) is used to substitute replacement matrices that can be calculated empirically from the ensemble on which the current iteration is based.

$$\mathbf{K}' = \mathbf{C}_{ko}^e[\mathbf{C}_{oo}^e + (1+\lambda)\mathbf{C}(\boldsymbol{\varepsilon})]^{-1}\mathbf{H} + [\mathbf{C}_{kk}^{e-1}\mathbf{C}_{ko}^e\mathbf{C}_{ok}^e\mathbf{C}_{kk}^{e-1} + (1+\lambda)\mathbf{C}_{kk}^{e-1}]^{-1}\mathbf{C}^{-1}(\mathbf{k})\mathbf{K}_0 \quad (12.3.6)$$

Recall that \mathbf{C}_{kk}^e is rank-deficient; hence its numerical inverse can only be approximate.

For convenience, we now replace $\mathbf{C}(\boldsymbol{\varepsilon})$ by the identity matrix. This involves no loss of generality if \mathbf{H} is comprised of weighted observations, and the matrix \mathbf{O} that is computed from \mathbf{K} is comprised of weighted model outputs. Equation (12.3.6) then becomes.

$$\mathbf{K}' = \mathbf{C}_{ko}^e[\mathbf{C}_{oo}^e + (1+\lambda)\mathbf{I}]^{-1}\mathbf{C}^{-1/2}(\boldsymbol{\varepsilon})\mathbf{H} + [\mathbf{C}_{kk}^{e-1}\mathbf{C}_{ko}^e\mathbf{C}_{ok}^e\mathbf{C}_{kk}^{e-1} + (1+\lambda)\mathbf{C}_{kk}^{e-1}]^{-1}\mathbf{C}^{-1}(\mathbf{k})\mathbf{K}_0 \quad (12.3.7)$$

Note that $\mathbf{C}^{-1/2}(\boldsymbol{\varepsilon})$ is usually a diagonal matrix whose elements are weights that appear in a PEST control file.

We now focus on the inconvenient second term of the above equation. At the same time we make computation of empirical covariance matrices explicit. See equations (12.2.5) to (12.2.7). The second term of (12.3.7) then becomes, after cancellation of some ($N-1$) terms

$$\mathbf{K}'_2 = \left[\left\{ (\mathbf{K} - \underline{\mathbf{K}})(\mathbf{K} - \underline{\mathbf{K}})^t \right\}^{-1} (\mathbf{K} - \underline{\mathbf{K}})(\mathbf{O} - \underline{\mathbf{O}})^t (\mathbf{O} - \underline{\mathbf{O}})(\mathbf{K} - \underline{\mathbf{K}})^t \left\{ (\mathbf{K} - \underline{\mathbf{K}})(\mathbf{K} - \underline{\mathbf{K}})^t \right\}^{-1} + (1+\lambda)(N-1) \left\{ (\mathbf{K} - \underline{\mathbf{K}})(\mathbf{K} - \underline{\mathbf{K}})^t \right\}^{-1} \right]^{-1} \mathbf{C}^{-1}(\mathbf{k})\mathbf{K}_0 \quad (12.3.8)$$

Repetition of terms in the above equation renders its calculation a little less inconvenient. Now we introduce singular value decomposition to evaluation of inverses. This brings further convenience as it will allow cancellation of some matrices. First, subject the following two matrices to singular value decomposition

$$(\mathbf{K} - \underline{\mathbf{K}}) = \mathbf{U}_K \mathbf{S}_K \mathbf{V}_K^t \quad (12.3.9a)$$

$$(\mathbf{O} - \underline{\mathbf{O}}) = \mathbf{U}_O \mathbf{S}_O \mathbf{V}_O^t \quad (12.3.9b)$$

The following matrix products can then be expressed in terms of matrices that appear on the

right side of the above equations. Note that, in writing the following equations, products of transposed and untransposed orthogonal matrices have been omitted, as these are equal to the identity matrix.

$$[(\mathbf{K} - \underline{\mathbf{K}})(\mathbf{K} - \underline{\mathbf{K}})^t]^{-1} = \mathbf{U}_K \mathbf{S}_K^{-2} \mathbf{U}_K^t \quad (12.3.10a)$$

$$(\mathbf{K} - \underline{\mathbf{K}})(\mathbf{O} - \underline{\mathbf{O}})^t (\mathbf{O} - \underline{\mathbf{O}})(\mathbf{K} - \underline{\mathbf{K}})^t = \mathbf{U}_K \mathbf{S}_K \mathbf{V}_K^t \mathbf{V}_O \mathbf{S}_O^2 \mathbf{V}_O^t \mathbf{V}_K \mathbf{S}_K \mathbf{U}_K^T \quad (12.3.10b)$$

If the above matrix products are substituted into (12.3.8) and orthogonal matrix products cancelled out, we obtain.

$$\mathbf{K}'_2 = \mathbf{U}_K \mathbf{S}_K [\mathbf{V}_K^t \mathbf{V}_O \mathbf{S}_O^2 \mathbf{V}_O^t \mathbf{V}_K + (1 + \lambda)(N - 1)\mathbf{I}]^{-1} \mathbf{S}_K \mathbf{U}_K^t \mathbf{C}^{-1}(\mathbf{k}) \mathbf{K}_0 \quad (12.3.11)$$

So that finally

$$\mathbf{K}' = \mathbf{C}_{ko}^e [\mathbf{C}_{oo}^e + (1 + \lambda)\mathbf{I}]^{-1} \mathbf{C}^{-\frac{1}{2}}(\boldsymbol{\epsilon}) \mathbf{H} + \mathbf{U}_K \mathbf{S}_K [\mathbf{V}_K^t \mathbf{V}_O \mathbf{S}_O^2 \mathbf{V}_O^t \mathbf{V}_K + (1 + \lambda)(N - 1)\mathbf{I}]^{-1} \mathbf{S}_K \mathbf{U}_K^t \mathbf{C}^{-1}(\mathbf{k}) \mathbf{K}_0 \quad (12.3.12)$$

Some details, such as matrix scaling before undertaking singular value decomposition, have been omitted from the above derivation. See the original paper for details. However the above derivation serves to illustrate some of the steps that are required for robust, iterative implementation of an ensemble smoother. The second term of (12.3.12) is computationally more expensive than the first term. However use of singular value decomposition simplifies it and makes it easier to compute. Furthermore, this term can optionally be omitted. Its omission not only removes the numerical burden of computing the matrices from which it is formed. It can also reduce the number of iterations required for the IES process to achieve a good fit with a history-matching dataset, as the need for maximum respect of the original parameter ensemble embodied in \mathbf{K}_0 is neglected. However, every convenience comes with a cost; in some cases, diversity of the final ensemble suffers because of this.

Selection of an iteration-specific value of the Marquardt lambda is just as important for IES as it is for regularized inversion. During each iteration of the IES process, a number of model runs is devoted to a trial-and-error assessment of its optimal value. Normally a variety of different lambdas are tested on a small number of realizations by calculating upgrades for these realizations to fill the new \mathbf{K}' matrix. The most successful lambda is then used to upgrade the remaining realizations.

Termination criteria for the ensemble smoothing process can be formal or informal. If formal, they may be based on slow improvement of model-to-measurement fit, or on small alterations to individual parameter realizations. These same assessments can be made manually.

12.3.3 Multiple Data Assimilation

The MDA (for “multiple data assimilation”) scheme of Emerick and Reynolds (2013) provides another strategy for accommodation of model nonlinearity through iterative data assimilation. MDA is supported by PESTPP-IES. This scheme brings similar numerical benefits to use of the Marquardt lambda.

We start from equation (12.2.16). For iteration i of the ensemble-based inversion process, (12.2.16) is modified to

$$\mathbf{K}' = \mathbf{C}_{ko}^e (\mathbf{C}_{oo} + \alpha_i \mathbf{C}(\boldsymbol{\epsilon}))^{-1} \mathbf{H} \quad (12.3.13)$$

α_i must exceed 1.0; hence measurement noise is effectively inflated.

In implementing MDA, a modeller must choose the number of iterations of the IES process

before that process is initiated. Suppose that this number is R . Then, in order to ensure compliance with Bayes equation, values that are assigned to α_i during each iteration must be such that

$$\sum_{i=1}^R \frac{1}{\alpha_i} = 1 \quad (12.3.14)$$

Another important Bayes-compliant aspect of MDA implementation is that fresh realizations of measurement noise must be added to columns of \mathbf{H} during every iteration.

While the above steps ensure Bayesian compliance in a linear world, the same does not apply for a nonlinear model. Nevertheless, as a practical measure, history-matching performance of the ensemble smoothing process is often improved through the “gradual” assimilation of information contained in the field measurement dataset \mathbf{h} that is implemented through the MDA process.

12.3.4 Localization

Recall from section 2.4 that when a matrix multiplies a vector, the outcome of the matrix multiplication process is a vector that is a linear combination of the columns of the matrix. The factors by which the matrix columns are multiplied are the elements of the vector.

The outcome of equation (12.2.16) is a set of vectors that comprise the columns of matrix \mathbf{K}' . Each of these columns is a linear combination of the columns of the first matrix on the right side of this equation, namely \mathbf{C}^e_{ko} . These are calculated using the transpose of equation (12.2.5), that is

$$\mathbf{C}^e_{ko} = \frac{(\mathbf{K} - \underline{\mathbf{K}})(\mathbf{o} - \underline{\mathbf{o}})^t}{N-1} \quad (12.3.15)$$

The columns of \mathbf{C}^e_{ko} are linear combinations of perturbations of prior random parameter sets from the prior mean parameter set. It follows that posterior random parameter sets calculated using (12.2.16) are also linear combinations of prior parameter perturbations about their means. The subspace of parameter space in which these new parameter vectors lie does not therefore change. Confinement to this same parameter subspace continues from iteration to iteration of the ensemble-based inversion process.

Now consider equation (12.3.12). The same considerations as above apply to the first term of this equation. Meanwhile the leading matrix of the second term of (12.3.12) is \mathbf{U}_k which, from (12.3.9a), also spans the $\mathbf{K} - \underline{\mathbf{K}}$ subspace of parameter space. Once again, the subspace of parameter space spanned by initial parameter realizations is not breached.

Restriction of the search for an improved parameter set to a particular subset of parameter space diminishes the ability of ensemble-based inversion to attain a good fit with a measurement dataset if this subspace is not large enough. It can also cause problems if the true prior probability distribution of hydraulic properties is not the same as that assumed by a modeller. Under these circumstances it may still be possible to fit individual items of a measurement dataset; however this may require that unusually large factors be applied to some prior parameter sets \mathbf{k}_i in order to attain this fit. This, in turn, may precipitate unusually large or small parameter values in those parts of a model domain where field measurements are absent. Predictive bias may result.

This problem can be viewed in another way. Use of a limited number of parameter realizations (and hence restriction of the ensemble-based inversion process to a small dimensional subset of parameter space) may create the illusion of correlation between some model outputs and

some parameters where no such correlation exists. This phenomenon, known as “spurious correlation”, is further discussed below. This is an outcome of empirical calculation of a covariance matrix where the averaging process that is central to this calculation does not yield a value of zero where the true correlation is, in fact, zero because cancellation of random, uncorrelated numbers is incomplete.

There may be instances where a modeller knows that certain observation-to-parameter sensitivities/correlations are small or zero. For example, he/she may know that a parameter that pertains to a system stress at a certain time cannot influence model outputs at an earlier time. This correlation should be directly ascribed a value of zero by the modeller.

Predisposition in time may not be the only indicator of diminished or zero correlation. Measurement-pertinent model outputs may exhibit low or zero correlations with spatial parameters that are located far from them. In recognition of this, a modeller may override empirical, ensemble-based correlation calculation in order to assign zero or low values to these correlations him/herself. The concept of diminishing correlation with distance gives rise to the term “localization”. However in groundwater modelling, this strategy must be applied with caution. For example, conductance parameters at or close to an outflow model boundary can influence all upstream heads, regardless of separation.

User-prescribed diminution of empirical correlations is often effected by superimposition of a so-called “tapering matrix” on an ensemble-based covariance matrix. This superimposition process is known as the “Schur product”; elements of one matrix are directly multiplied by elements of another matrix. This dampens or removes spurious correlations in accordance with modeller expectations of where true correlations are likely to be small. See Chen and Oliver (2017), and references cited therein, for a full discussion of this topic.

Algorithms that can automatically detect and reduce spurious correlation have been proposed in the literature. One of these (see Luo et al, 2018) is implemented in the PESTPP-IES ensemble smoother. The Luo et al algorithm assesses the statistical significance of empirical correlations against background “correlation noise” by repeatedly applying a circular shifting operation to observation or parameter realizations and re-calculating correlations. If a specific parameter-to-observation correlation is invisible above correlation noise, then it is either reduced, or declared to be zero.

The conceptual benefits of reduction or elimination of spurious correlation are obvious. Furthermore, it can be shown that strategic alterations to empirical correlations attained by setting some of them to zero, and/or by Schur multiplication by a tapering matrix, can add dimensions to the subspace of parameter space in which history-matching takes place. Problems associated with working in a limited-dimensional parameter subspace are thereby ameliorated, to some extent at least.

As stated above, however, localization must be applied with caution in the groundwater modelling context. In many cases a “safer” way to increase the dimensionality of the space in which the ensemble-based inverse problem is solved is to simply increase the number of realizations on which the IES process is based.

12.3.5 Other Performance Enhancing Devices

This subsection describes two algorithms, both of which are available through the PESTPP-IES ensemble smoother, that can enhance the performance of ensemble-based inversion in difficult history-matching contexts. Other options are provided by this and other IES packages. No doubt, algorithms that increase the history-matching alacrity of ensemble methods will

continue to evolve over time in order to meet the challenges that are posed by nonlinear uncertainty analysis in general, and by working in reduced dimensional subspaces in particular.

12.3.5.1 Variance Reinflation

Experience demonstrates that, under some circumstances, periodic recentering of parameter ensembles, and simultaneous reinflation of realization variability, can breathe new life into a stalled ensemble-based inversion process. During an iteration in which recentering is undertaken, the current means of all parameter values over all realizations are calculated. These new means then replace initial parameter means. At the same time, initial variability about initial means is reinstated; however variability is now centred on the new means. With parameter realizations altered in this manner, the IES process then proceeds in the usual manner.

This reinflation and recentering step can be especially useful where it becomes apparent that modeller-provided initial mean hydraulic property values do not reflect real-world hydraulic properties. Under these circumstances, it may be difficult to calculate new realizations from old realizations that support a high level of model-to-measurement fit, for the factors applied to some realizations may require considerable “stretching” of those realizations about an inappropriate mean. Balancing of stretching in one direction at locations where data are dense by concomitant stretching in the other direction in places where data are sparse may induce local parameter bias.

Note that this problem can also be partially overcome by defining spatial parameters as multipliers about a mean (which may itself be spatially variable). Mean values should then be defined as parameters so that they, too, are adjustable.

Recentering and reinflation can also be useful where nonlinearity of an inverse problem adversely affects IES performance. It allows the IES process to “take a deep breath and start again”, but from a more advanced position in parameter space than that defined by initial realizations.

12.3.5.2 Local Updating

Zhang et al (2018) describe a methodology that enhances IES performance in solving an inverse problem that is characterized by multiple objective function minima, and/or by a continuous objective function minimum whose shape in parameter space exhibits considerable curvature. Under these circumstances, normal operation of the IES process in which all empirical correlations are used to calculate an upgrade for all parameter realizations may be inappropriate. In situations such as these, it may be better to encourage different groups of realizations to take different trajectories in parameter space. This grants separate parts of parameter space representation in posterior parameter realizations; this is necessary where the posterior parameter probability distribution is bimodal or multimodal. Equations through which realizations that comprise different groups of parameters are adjusted are the same as those presented above. However covariances that are represented in these equations are realization-group-specific.

Realizations are partitioned into groups on the basis of proximity in parameter space, and on the basis of similarity of objective functions computed using these realizations. Variables which control this partitioning are user-selectable.

As is discussed above, validity of use of equations such as (12.2.16) requires that the ensemble size be large enough to span a suitably-large subspace of parameter space. The size of this

subspace is reduced where parameters are subdivided into groups and treated separately for the purpose of parameter upgrade calculation. Hence, when invoking local updating in this manner, it is recommended that the initial parameter ensemble contain a large number of realizations, possibly larger than would normally be used.

12.4 Some Practicalities

12.4.1 How Many Realizations?

The answer to this question is obvious: the more the better. However, where model run times are large, and where access to computing hardware that can increase the level of model run parallelization is limited, practical considerations limit the number of realizations that can be employed.

IES gains numerical efficiencies through dimensional reduction. Unless localization is employed, it works in a subspace of parameter space that is spanned by parameter realizations that are drawn from the prior parameter probability distribution. Ideally, randomization ensures that each of these realizations has a different projection onto the solution space of parameter space that would emerge from singular value decomposition of the Kahunen-Loèeve-transformed \mathbf{Z} matrix of equation (11.1.2). The number of realizations on which an IES parameter adjustment process is based should exceed the dimensionality of this space, for this ensures that each \mathbf{k}' vector that emerges from the IES process has roughly the same projection $\underline{\mathbf{k}}$ onto this space as does the true parameter field \mathbf{k} . $\underline{\mathbf{k}}$ is the solution to the inverse problem that would emerge from regularized inversion using methods that are discussed in chapter 8. Correct parameter projection onto the calibration solution space ensures good model-to-measurement fit. Note, however, that some solution space scatter is necessary in order to reflect measurement-noise-induced parameter uncertainty.

Ideally, while vectors that emerge from the IES process should all possess roughly the same projection $\underline{\mathbf{k}}$ onto the idealized calibration solution space, they should still be diverse. This is schematized in Figure 12.1. This allows them to explore null space parameter uncertainty. Attainment of constrained diversity in this manner requires that the number of parameter realizations considerably exceed the dimensionality of the calibration solution space.

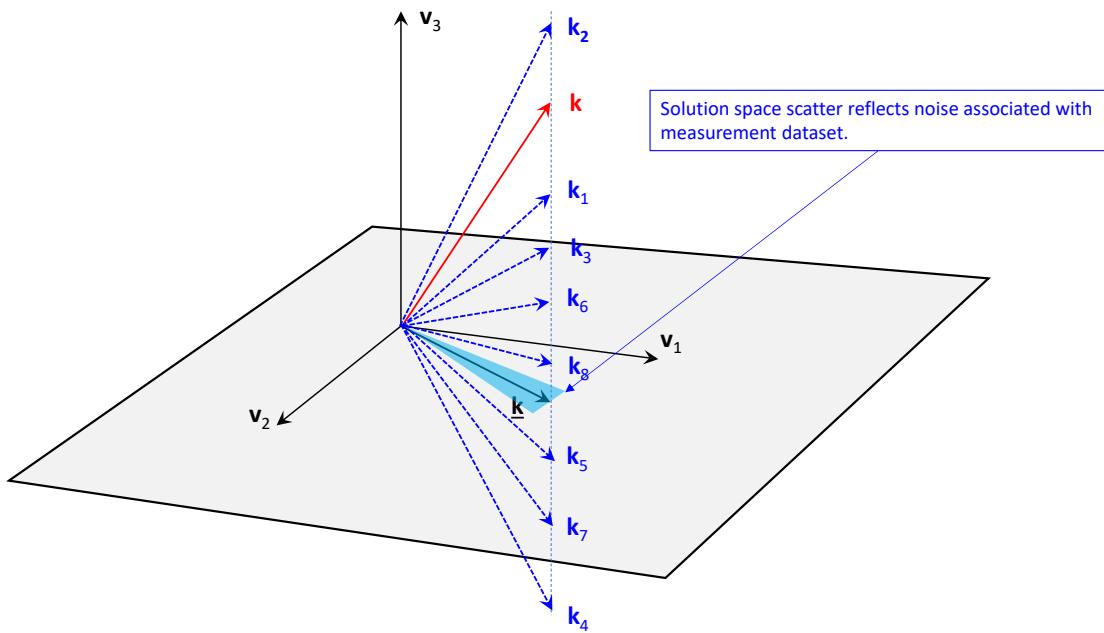


Figure 12.1. The true (but unknown) parameter set is shown in red. Its projection onto the calibration solution space is shown in black. Parameter fields that emerge from the IES process are shown in blue. Compare this figure with figure 8.2.

Of course, a modeller does not know the dimensionality of the calibration solution space when choosing the number of realizations to employ in ensemble-based history-matching. Furthermore, because the number m of elements that comprise a parameter set \mathbf{k} is often very large, there is no opportunity to formulate a sensitivity matrix \mathbf{Z} , and to then calculate solution space dimensionality using linear analysis methodologies that are described in chapter 10. So the number of realizations must be guessed. In general, the larger is a measurement dataset, the larger is the dimensionality of the solution space. Where measurements are dense in time and space, the dimensionality of the solution space may be large.

As has already been discussed, another problem that attends deployment of too few realizations is that of spurious correlation. Covariance matrices that are used in equations such as (12.2.16) are empirical. If an observation \mathbf{o}_i is entirely insensitive to a parameter \mathbf{k}_j , then \mathbf{o}_i and \mathbf{k}_j should exhibit no statistical correlation. However where empirical covariance calculation is based on a limited number of realizations in which all elements of \mathbf{k} are simultaneously varied, it is difficult to separate parameters that induce \mathbf{o}_i variability from those that do not. In contrast, where covariance calculation is based on many realizations, cancellation of spurious effects is able to occur. Nevertheless, the number of realizations that is required to ensure complete cancellation can be very large.

The outcomes of realization insufficiency can be viewed in yet another way. All realizations that are sampled from the prior parameter probability distribution necessarily possess a solution space component and a null space component. As stated above, attainment of a good fit between model outputs and field measurements requires collective spanning of the calibration solution space by these realizations. Meanwhile, quantification of parameter uncertainty requires spanning of enough dimensions of the calibration null space for this uncertainty to be expressed. However, if the number of realizations is sufficient to span the calibration solution space, but small enough for the average of all realizations to possess a nonzero projection onto the calibration null space (because of lack of complete cancellation of null space components),

this may engender parameter and predictive bias. This is because the posterior mean parameter field then gives expression to this null space component that (by definition) is unsupported by the measurement dataset. As is discussed in chapter 9, the possibility of parameter and predictive bias is a danger that is associated with all dimensional reduction methods.

12.4.2 Weights and Noise

Theoretically, observation weights in a PEST control file should be inversely proportional to the standard deviations of noise associated with respective measurements. Previous chapters of this book discuss how a proportionality constant for these weights may emerge through solution of a manually or numerically regularized inverse problem.

However, as has already been mentioned, most model-to-measurement misfit arises from simulator and parameterization inadequacies. As such, it cannot be considered to be a random supplement to field measurements. Rather, it has “structure”, and is therefore often referred to as “structural noise”. As will be discussed later in this book, certain measures can be adopted to minimize the ability of structural noise to induce bias in estimated parameters and in predictions that are sensitive to these parameters. These include the following.

- Formulation of “secondary observations” or “derived observations”, particularly spatial and temporal differences of primary observations; these are matched with identical functions of model outputs.
- Ensuring that different groups of observations (both primary and secondary) receive equal visibility in the objective function that is reduced through history-matching.

While weighting within each observation group may still reflect measurement credibility (and hence measurement noise), weighting between groups may become somewhat subjective, being more reflective of structural noise than measurement noise. Furthermore, while use of weights instead of the inverse of a measurement noise covariance matrix implies lack of measurement noise correlation, re-use of observations in groups that characterize spatial and temporal differences does (in fact) imply a noise correlation structure (as it should).

These considerations, and the necessity for observations \mathbf{h} to be supplemented by realizations of measurement noise in formulation of the \mathbf{H} matrix that features in equations such as (12.2.16), poses the question of how realizations of “measurement noise” should be generated. In particular, modellers should beware of default settings that may assume that user-supplied weights are inverted standard deviations of measurement noise. Obviously, weights also serve other purposes.

Implicit in IES history-matching is minimization of an objective function; see equation (11.1.6), and recall that IES mimics the action of RML. The objective function that IES minimizes is different for each \mathbf{h}_i that comprises a column of the matrix \mathbf{H} . Part of the reason for this is the fact that each \mathbf{h}_i is accompanied by a different realization of measurement noise. Nevertheless, these realizations should not be so different as to invalidate primary design specifications for any objective function that is minimized in order to solve an ill-posed inverse problem. These are to ensure maximum visibility of information that is hosted by a measurement dataset, and to defend the inverse problem solution process against the deleterious effects of structural noise. These specifications require deployment of strategies such as those listed above.

These considerations suggest that generation of measurement noise should not be based on weights that are provided in a PEST control file. Instead, a modeller should provide a $C(\mathbf{e})$

matrix specifically for the purpose of random measurement noise generation. (This is allowed by the PESTPP-IES ensemble smoother.) This $C(\epsilon)$ may reflect the true noise that is thought to contaminate field measurements. Or it may imply a level of model-to-measurement fit which, if exceeded, would expose estimated parameters to the deleterious effects of overfitting. These choices are subjective; furthermore, they will generally be based on a diagonal $C(\epsilon)$.

Alternatively, a modeller may decide not to enhance a measurement dataset \mathbf{h} with realizations of measurement noise at all. It is often discovered that model-to-measurement fit attained by IES is good, but not great, especially if the ensemble-based inversion process is allowed to proceed for only a few iterations. Acceptance of a good, but imperfect, fit with a measurement dataset may be construed as “noise enough”. Furthermore, where posterior parameter diversity is more an outcome of parameter nonuniqueness arising from data insufficiency than of noise associated with field measurements, this strategy is unlikely to compromise uncertainty analysis.

12.4.3 Ensemble Collapse

The term “ensemble collapse” is widely used, greatly feared, but often misunderstood by groundwater modellers.

Sometimes it is found that all parameter fields that emerge from ensemble-based inversion provide an equally good fit with a history-matching dataset, and that this fit is very good. This is a common occurrence where a measurement dataset is sparse. Some modellers mistakenly refer to this condition as “ensemble collapse”.

This is not ensemble collapse.

Ensemble collapse occurs where diversity between posterior realizations is artificially lost. Where a measurement dataset is sparse, and parameter nonuniqueness is therefore likely to prevail, lack of diversity between realizations that emerge from an IES process should be treated with considerable suspicion. However, it is important to realize that attainment of a high level of model-to-measurement fit by many parameter fields on the one hand, and lack of parameter diversity in achieving this fit on the other hand, are two different things. Where a null space has many dimensions, the addition of any null-space-projected parameter field to a parameter field that calibrates a model, leaves model-to-measurement fit unchanged (in a linear world).

Use of too few realizations in an IES process can precipitate ensemble collapse. If the number of realizations does not exceed the dimensionality of the calibration solution space, then ensemble collapse is inevitable, for under these conditions the dimensionally-reduced inverse problem is well-posed. Even with deployment of a greater number of realizations than this, ensemble diversity may be artificially reduced through failure to include the second term in the second bracket of equation (12.2.1) in the ensemble adjustment process (as is often done). However, as is discussed above, inclusion of this term may slow the ensemble-based inversion process while rendering attainment of a high level of model-to-measurement fit difficult.

12.4.4 Integrity of Uncertainty Intervals

In many practical applications of IES-based inversion, parameter fields that emerge from a history-matching process provide varying levels of fit with a measurement dataset. As stated in the previous subsection, this should not be construed as avoidance of ensemble collapse, and celebrated as such. In fact, it raises serious questions about how these parameter fields should

be used when the model to which they belong is used to make predictions of management interest.

As discussed above, IES parameter adjustment is based on an implicit, rank-deficient Jacobian matrix. Sometimes, use of this matrix promotes rapid attainment of good model-to-measurement fit (with the possibility of some parameter/predictive bias for reasons that have already been stated). On other occasions, attainment of a good fit with a history-matching dataset may require a considerable number of iterations; it may also require implementation of strategies such as parameter recentering and variance reinflation. These can erode the efficiency of the history-matching process. A modeller may then have to decide whether an IES process should be allowed to continue for more iterations in the hope that more parameter fields will yield better fits with a history-matching dataset. However, he/she may worry that those parameter fields that have already attained an expected level of fit may lose their diversity if ensemble-based inversion is allowed to continue.

Failure of an IES process to attain a high level of model-to-measurement fit may also be an outcome of sampling an incorrect prior parameter probability distribution. For example, spatial correlations that are expressed by a prior parameter probability distribution may be too large to account for locally high or low piezometric gradients that are induced by small scale heterogeneity. Or the prior parameter probability distribution may be incapable of expressing thinning of an aquitard that reduces local vertical head gradients. These are incidences of prior-data conflict – an issue of which a modeller should always be aware. As discussed later in this text, prior-data conflict can be formally examined using an ensemble of parameter fields; see section 14.5. Alternatively, if a modeller is worried about the suitability of his/her prior parameter probability distribution, he/she may decide to implement hierarchical inversion in order to allow the prior to adjust to field data. However experience suggests that this can slow the pace of inversion, and may not always be successful.

Where model-to-measurement fit is not equally good for all adjusted realizations, a modeller may be forced to make some unpleasant choices. He/she may decide to abandon noncompliant realizations. However this may artificially reduce parameter diversity, and therefore result in underestimation of predictive uncertainty. Alternatively, a modeller may decide that parameter and predictive diversity that result from model-to-measurement misfit can be considered as a viable surrogate for diversity that results from parameter nonuniqueness that the IES process may not be capable of expressing, given the number of realizations that it employs and the prior parameter probability distribution that these realizations sample. There may be no theoretical justification for this decision; however it may nevertheless be taken in order to prevent uncertainty underestimation.

Another issue that may affect the integrity of IES-calculated predictive uncertainty intervals is that of model nonlinearity. As is discussed above, an outcome of working in a limited dimensional subspace of parameter space is the possibility of history-match-induced parameter and predictive bias. This problem is exacerbated where relationships between model outputs and model parameters are nonlinear. The Bayesian credentials of ensemble methods are lost under these circumstances. Furthermore, there is no way to know the extent to which assessment of uncertainty intervals has suffered because of this.

12.4.5 Selection of the Prior

As discussed above, sampling from an inappropriate prior parameter probability distribution may impede attainment of a satisfactory level of model-to-measurement fit. However this is not always the case; use of an incorrect prior may result in a perfectly good fit with field

measurements. Therefore, lack of prior-data conflict should not be construed as confirmation of the validity of a particular prior parameter probability distribution.

An IES process cannot be initiated without selection of a prior parameter probability distribution, for it requires generation of parameter realizations on which to base model runs. Covariances that emerge from model runs are used to calculate parameter improvements. In some ways, therefore, parameter realizations can be considered as replacements for finite difference derivatives as the primary mechanism for parameter adjustment and model-to-measurement misfit reduction. Of course, the prior parameter probability distribution has far greater importance than this, for it has a large impact on the posterior parameter probability distribution. This can sometimes be forgotten as the success of ensemble-based history-matching is often based on how good a fit is attained with a measurement dataset, and perhaps a little less on whether posterior predictive uncertainties are overstated or understated.

Of particular importance in selection of a prior parameter probability distribution is representation of connectedness of high and/or low permeabilities. This is expressed by correlation lengths, and variability of the magnitude and direction of spatial correlation in different parts of a model domain. Opportunities afforded by non-stationarity hydraulic property field generation that are described in chapter 4 of this book are of particular relevance in this context. These can enhance IES utility, not just in attaining good fits with measurement datasets, but also in ensuring that parameter fields that emerge from ensemble-based history-matching are able to express the full range of possibilities that characterize predictions of management interest.

Where the prior is uncertain, IES can be asked to implement hierarchical history-matching. Hierarchical regularized inversion is discussed in sections 8.6 and 9.4. Hierarchical history-matching can also be implemented in conjunction with IES for reasons that are discussed in these sections. So-called noncentred parameterization of a model domain allows use of iid parameters to represent the locations of hydraulic property heterogeneity. Meanwhile, patterns that heterogeneity adopts are set by independent geostatistical hyperparameters. The former can be ascribed to pilot points, while the latter can be ascribed to conceptual points or other parameterization devices such as zones of piecewise constancy. History-match adjustment of hyperparameters does not affect the prior probability distribution of iid parameters, notwithstanding changes to the geostatistical characterization of model-based hydraulic property fields that hyperparameters inform.

When equipped with hierarchical parameterization, an IES inversion process can express uncertainty in prior uncertainties at the same time as it expresses the predictive consequences of spatial heterogeneity. Oliver (2022) points out that nonlinearities that are associated with hierarchical inversion may challenge IES history-matching alacrity. Strategies discussed above that can enhance IES performance in nonlinear contexts may therefore be warranted.

12.4.6 Giving IES a Head Start

Where a model is calibrated using regularized inversion prior to implementation of IES-based uncertainty analysis, the latter can receive a head start. Ideally, calibration has achieved a good fit with a history-matching dataset. It may have been undertaken traditionally, or using ENSI. In either case, theory suggests that regularized inversion has yielded a parameter field that is somewhere near the centre of the posterior parameter probability distribution. It therefore makes sense to centre random realizations of parameters on this parameter field when commencing an IES run. This achieves a similar purpose to recentering and variance reinflation that IES can undertake automatically if asked to do so; see section 12.3.5.1. However the hard

work of finding a suitable parameter field on which to centre realizations has already been done during the model calibration process.

The IES process can be given a further head start if, instead of using the prior covariance matrix for calculation of initial parameter realizations, a linear approximation to the posterior covariance matrix is employed for this purpose. Using formulas that are provided in chapter 10 (that are employed by utilities available through the PEST and PyEMU suites), a posterior parameter covariance matrix can be readily calculated using a Jacobian matrix that was filled during regularized inversion, or that is specially calculated for the occasion using optimized parameter values. Parameter fields that are sampled using this matrix should require only a few iterations of IES adjustment in order to provide just as good a fit with a measurement dataset as the parameter field that was achieved through model calibration itself.

12.4.7 IES Following ENSI

As is described in section 9.4, ensemble space inversion (ENSI) seeks model run efficiency when undertaking regularized inversion by working in a subspace of parameter space that is spanned by realizations of the prior parameter probability distribution. The same considerations apply to selection of a prior parameter covariance matrix for ENSI inversion as those that apply for IES-based inversion. Where necessary, it should represent the possibility of local or widespread connected permeability (or impermeability) whose omission may compromise model-to-measurement fit, introduce bias to some model predictions, and/or result in assessment of a reduced range of predictive possibilities.

Following ENSI inversion, IES can be used to explore parameter and predictive uncertainty. Efficiencies in IES performance can be gained if initial IES realizations are centred on the ENSI-calibrated parameter field. Even if these realizations are based on the prior parameter covariance matrix, their centring on a parameter field that fits a measurement dataset well should assist the IES process in deriving a suite of parameter fields that do the same.

Alternatively, because the ENSI process requires calculation of a Jacobian matrix, a modeller may decide to commence an ensuing IES process by sampling a linearized approximation to the posterior parameter probability distribution. As stated above, this can enhance the ability of an IES process to quickly attain a high level of model-to-measurement fit. A problem with this strategy, however, is that the Jacobian matrix that is yielded by an ENSI process pertains to “super parameters”. These may include some native model parameters. However most super parameters are actually factors by which realizations are multiplied before being combined to form the calibrated parameter field. The Jacobian matrix is therefore specific to the subspace of parameter space in which ENSI inversion was undertaken.

This does not have to be a problem, however. Provided ENSI inversion was based on a superfluity of realization factors, highly efficient implementation of IES can be achieved by continuing to work in the same subspace of parameter space as that in which ENSI inversion took place. In this context, “superfluity” requires the use of a greater number of realizations than the dimensionality of the calibration solution space.

As it prepares files for an ENSI inversion process, the ENSIPREP utility that is supplied with the PEST_HP suite writes a prior covariance matrix for ENSI parameters. This, together with the subspace-specific Jacobian matrix yielded by the ENSI process, can be used to calculate a linearized approximation to the posterior covariance matrix of ENSI parameters that were estimated during that process using equations that are presented in chapter 10. Random realizations of super parameters, and any native model parameters that were also estimated

during the ENSI process, can be generated using this covariance matrix. IES can then be used to adjust these random super parameters and native model parameters so that they all fit the measurement dataset while maintaining diversity. This process should be numerically fast as these samples of the linearised ENSI posterior covariance matrix should require little adjustment. Following their adjustment using IES, realizations of native model parameters can then be calculated from realizations of ENSI super parameters using utility programs provided with the PEST_HP suite. See the PEST_HP manual for further details.

12.4.8 Data Rich Environments

Where field measurements are plentiful, the dimensionality of the calibration solution space is likely to be high. At the same time, high spatial data densities may expose nuances of subsurface hydraulic property variability that are difficult to express using conventional or nonconventional prior parameter probability distributions. Taken together, these may compromise the ability of the IES process to achieve a high level of model-to-measurement fit.

In these circumstances, IES should be used with caution. Furthermore, because of the large number of realizations that it must use in these circumstances, efficiencies gained through IES deployment may not be great. At the same time, the possibility of predictive bias, and of overstatement of predictive uncertainty through underfitting, may be increased.

Where field measurements are plentiful in space and time, and where predictions of future system behaviour are required under conditions that are not too dissimilar from those that prevailed under historical conditions, predictions of management interest may be “data driven”. These types of predictions are discussed in greater detail in chapter 15. There it is shown that, under these circumstances, model history-matching should be considered as a kind of “machine learning” activity. The need for model parameters to provide realistic representation of the hydraulic properties after which they are named is secondary to the requirement for parameters to be sufficient in number and disposition to enable model outputs to fit field data well, for this is sufficient to guarantee predictive integrity.

In cases such as these, history-matching in support of decision-making may be most effectively conducted through regularized inversion. Under these circumstances, model predictive uncertainty is determined more by measurement noise than by parameter nonuniqueness.

Alternatively, other predictive strategies may be considered. Given the data-driven nature of key model predictions, and the importance of good statistical characterization of measurement noise for elucidation of predictive uncertainty, a modeller may consider the use of time series analyses methods based on forecasting theory. See, for example, Collenteur et al (2019), Petropoulos et al (2022) and references cited therein.

13. Other Parameter-Based Methods

13.1 General

In this section, two other methodologies that enable nonlinear predictive uncertainty analysis are examined. As for all uncertainty analysis methods that have been discussed so far, both of them require adjustability of model parameters. They can therefore be computationally expensive when deployed in conjunction with complex models. In the next chapter we show how a quantum gain in numerical efficiency can be attained by using a complex simulator to build statistical linkages between the measured past and the managed future; no parameter adjustment is required. Conditioning of the statistical model that emerges from this process is fast, regardless of the structural and parametric complexity of the simulator.

13.2 Constrained Predictive Maximization/Minimization

13.2.1 Concepts

Figure 13.1 schematizes two-dimensional parameter space. Note that while the methodology that is now described is readily deployed in higher dimensional parameter spaces than this, the concepts that underlie it are more easily depicted in two dimensions.

In figure 13.1, objective function contours are shown as full lines. The minimum of the objective function is marked by a cross. Problem well-posedness is assumed, as this minimum is unique. Perhaps manual regularization was employed in order to make it so. Perhaps uniqueness was achieved using prior information accompanied by a $C(k)$ matrix.

We also assume that the inverse problem has been provided with a weight matrix that is proportional to $C^{-1}(\varepsilon)$. Generally the weight matrix is diagonal. If any prior information is supplied, it is weighted according to $C^{-1}(k)$ while weighting of observations on the one hand and prior information on the other hand is properly balanced.

We denote the value of the minimized objective function in figure 13 as Φ_{\min} . The point in parameter space at which it is minimized is surrounded by closed contours. As stated above, their closure is an outcome of problem well-posedness. One of these contours is of particular interest to us. This corresponds to an objective function value of Φ_0 . We assume that if the objective function rises higher than Φ_0 , then confidence is lost in corresponding parameter values as the model is no longer calibrated. In other words we assume that, given the level of noise that accompanies the measurement dataset, an objective function value of Φ_0 is so unlikely that parameter values that give rise to this objective function, or to higher objective functions, are also statistically unlikely. The value of Φ_0 relative to Φ_{\min} is discussed shortly.

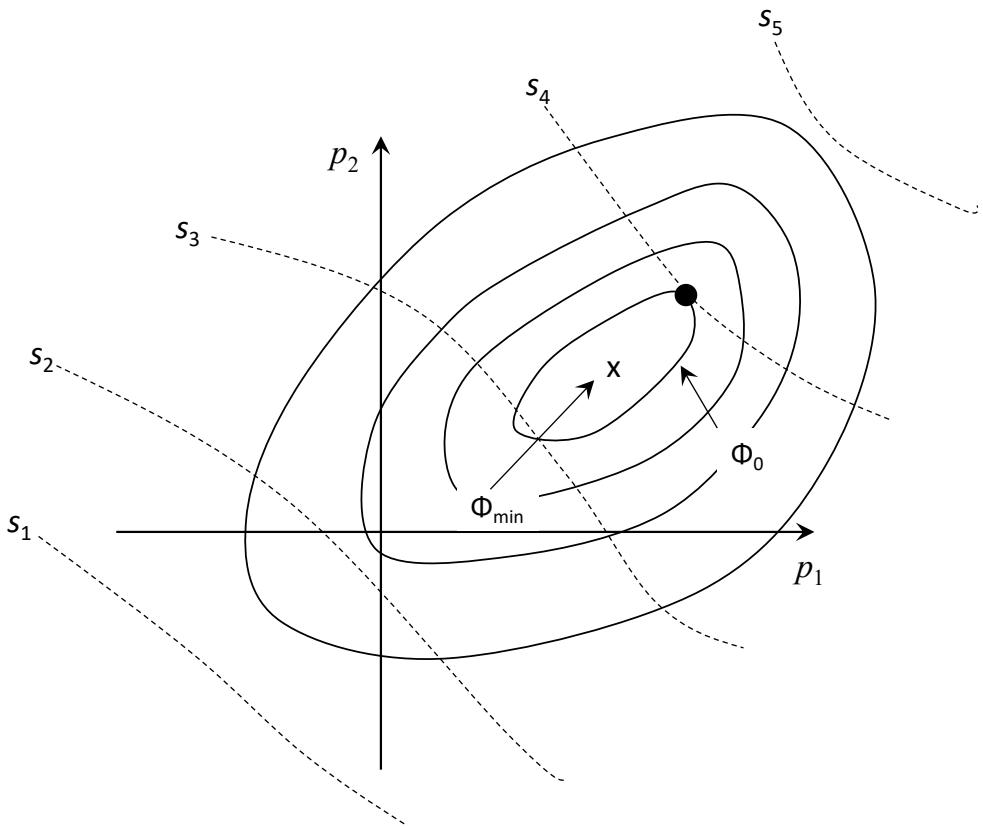


Figure 13.1. The dashed contours depict the dependence of a prediction on a set of parameters. The full contours depict the calibration objective function. The contour labelled Φ_0 is that at which model-to-measurement misfit is deemed to be too great for the model to be considered as calibrated. The dark spot marks the parameter set for which the prediction is maximized subject to the model “remaining calibrated”.

Figure 13.1 also depicts contoured values of a prediction s of interest. Let it be assumed that the value of the prediction increases with increasing contour index, so that s_2 is greater than s_1 , s_3 is greater than s_2 , etc.

Suppose that we wish to establish the two-sided $(1-\alpha)$ posterior confidence interval of the prediction s . This can be done by solving a constrained maximization problem followed by a constrained minimization problem. In the first of these problems, the prediction s is maximized subject to the constraint that the objective function rises no higher than Φ_0 . In the second of these problems, the prediction is minimized subject to this same constraint. If Φ_0 marks the upper $(1-\alpha)$ confidence limit of the objective function Φ , then the difference between these two prediction values defines the $(1-\alpha)$ confidence interval of the prediction.

13.2.2 Limiting Objective Function

Formulas for calculation of Φ_0 are provided by Vechhia (1987), Christensen and Cooley (1999) and Cooley (2004). Two values for the Φ_0 constraint are provided, one being rather loose and the other bring rather tight (perhaps too tight). The first provides the “simultaneous” or “Scheffé” confidence limit of Φ , while the second provides the “individual” confidence limit of Φ . These are:

$$\Phi_0 = \Phi_{\min} \left[\frac{m}{n-m} F_{\alpha}(m, n-m) + 1 \right] \quad (13.2.1)$$

and

$$\Phi_0 = \Phi_{\min} \left[\frac{t_{\alpha/2}^2 (n-m)}{(n-m)} + 1 \right] \quad (13.2.2)$$

The difference between these two statistics is analogous to the difference between Mahalanobis distance and the alternative, tougher, statistic that are discussed in sections 3.14 and 3.15 of this book. Recall that these statistics are used to assess whether a random vector belongs to a particular probability distribution. In the present instance, the random vector of interest is comprised of model-to-measurement misfit residuals. The statistic of section 3.14 relies on summation of squares; confidence is based on the chi-squared distribution. The statistic of section 3.15 relies on conformance of individual elements of a random vector to their own sub-distributions; confidence is based on the normal distribution. The statistic provided by equation (13.2.1) relies on the ratio of chi-squared distributions, i.e. the F distribution. The statistic provided by equation (13.2.2) relies on the Student's t distribution. This is closely related to the normal distribution; however it takes account of uncertainties in the mean and standard deviation of a random variable incurred by empirical estimation of these quantities. (Recall from chapter 7 that the post-calibration reference variance is generally estimated from the minimized objective function, and that the statistics of measurement noise are therefore assumed to be known only up to a constant of proportionality prior to solution of a well-posed inverse problem.)

Christensen and Cooley (1999) and Cooley (2004) point out that a correction factor (which is not too different from unity) should be applied to the right side of equation (13.2.2) to accommodate nonlinear model behaviour.

13.2.3 Predictive Noise

There are situations where a modeller may wish to expand the predictive confidence interval that is calculated in the above manner in order to account for “noise” that is associated with a prediction. The expanded interval is referred to as the “prediction interval”.

If a prediction is similar to measurements comprising a calibration dataset, presumably the noise associated with that prediction has the same statistical characteristics as that which is associated with field measurements. If not, then the prediction may possess its own kind of noise. In either case, the prediction interval accommodates model inadequacies in making this prediction precisely. The stochastic properties of “predictive noise” must, of course, be known if the prediction interval is to be calculated.

We assume that the linearized action of a model under calibration conditions is represented by equation (7.1.6), that is

$$\mathbf{h} = \mathbf{Xp} + \boldsymbol{\varepsilon} \quad (13.2.3)$$

We denote parameters using \mathbf{p} rather than \mathbf{k} in this section to remind the reader that the constrained minimization/maximization methodology that is the focus of the present discussion is applicable to well-posed inverse problems, or to inverse problems that have been rendered well-posed through use of prior information. Meanwhile, the action of the model under predictive conditions is described by

$$s = \mathbf{y}^t \mathbf{p} \quad (13.2.4)$$

where s is the prediction made by the model, and elements of the vector \mathbf{y} contain sensitivities of the prediction to model parameters \mathbf{p} . Where a prediction is accompanied by “predictive

noise", this becomes

$$s = \mathbf{y}^t \mathbf{p} + e \quad (13.2.5)$$

We associate a weight w_e with predictive noise. This should be inversely proportional to its standard deviation. The constant of proportionality should be the same as that used in calculating elements of the weight matrix \mathbf{Q} from the inverse of measurement noise $\mathbf{C}^{-1}(\boldsymbol{\varepsilon})$.

Under these conditions, equations (13.2.1) and (13.2.2) become

$$\Phi_0 = \Phi_{\min} \left[\frac{m+1}{n-m} F_{\alpha}(m+1, n-m) + 1 \right] \quad (13.2.6)$$

and

$$\Phi_0 = \Phi_{\min} \left[\frac{t_{\alpha/2}^2 (n-m)}{(n-m)} + 1 \right] \quad (13.2.7)$$

Note that equation (13.2.7) is identical to equation (13.2.2) despite inclusion of the predictive error term in definition of predictive limits.

13.2.4 Implementation

Vecchia and Cooley (1987) show that the constrained predictive minimization or maximization problem that is depicted in two dimensions in figure 13.1 can be solved for \mathbf{p} using the following formula

$$\mathbf{p} = (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \left\{ \mathbf{X}^t \mathbf{Q} \mathbf{h} - \frac{\mathbf{y}}{2\lambda} \right\} \quad (13.2.8)$$

where

$$\left(\frac{1}{2\lambda} \right)^2 = \pm \frac{\Phi_0 - \mathbf{h}^t \mathbf{Q} \mathbf{h} + \mathbf{h}^t \mathbf{Q} \mathbf{X} (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \mathbf{h}}{\mathbf{y}^t (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{y}} \quad (13.2.9)$$

Φ_0 is specified by the user using one of the equations provided in the previous subsection. The positive or negative square root of $\left(\frac{1}{2\lambda} \right)^2$ is chosen according to whether prediction minimization or maximization is performed. The value of \mathbf{p} that is calculated using equation (13.2.8) provides the parameter coordinates of the black dot of figure 13.1, or of its minimized counterpart.

Where predictive noise is included in the analysis, equation (13.2.8) is still employed to calculate \mathbf{p} . However (13.2.9) becomes

$$\left(\frac{1}{2\lambda} \right)^2 = \pm \frac{\Phi_0 - \mathbf{h}^t \mathbf{Q} \mathbf{h} + \mathbf{h}^t \mathbf{Q} \mathbf{X} (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{X}^t \mathbf{Q} \mathbf{h}}{\mathbf{y}^t (\mathbf{X}^t \mathbf{Q} \mathbf{X})^{-1} \mathbf{y} + w_e^{-2}} \quad (13.2.10)$$

while the estimated value of predictive error e under conditions of constrained objective function minimization/maximization is

$$e = -w_e^{-2}/2\lambda \quad (13.3.11)$$

PEST solves the above equations when run in "predictive analysis" mode. For nonlinear

models, equations (13.2.8) to (13.2.10) are solved for \mathbf{p} using an iterative procedure in which sensitivities embodied in the vector \mathbf{y} and the matrix \mathbf{X} are updated during each optimization iteration in accordance with current values of \mathbf{p} . Meanwhile residuals \mathbf{r} replace \mathbf{h} in the above equations; these are updated too as parameters change from iteration to iteration. Theoretically, this iterative procedure can commence from any desired starting \mathbf{p} . Practically, a choice of starting \mathbf{p} equal to \mathbf{p}_{\min} which minimizes the objective function Φ (i.e. a choice of starting \mathbf{p} for which the objective function is equal to Φ_{\min}) can reduce the number of model runs required for solution of this constrained minimization/maximization problem.

Note that the above procedure can be employed to analyse parameter as well as predictive uncertainty. Where the uncertainty of a parameter is the subject of inquiry, the predictive sensitivity vector \mathbf{y} is replaced by a vector whose elements are all zero except for that which pertains to the parameter in question; this element of \mathbf{y} has a value of one.

13.2.5 Some Practicalities

13.2.5.1 General

While the constrained, nonlinear, predictive minimization/maximization problem through which the posterior confidence or predictive interval of a prediction can be obtained is theoretically solvable using equations provided above, use of these equations in real-world modelling contexts is often fraught with difficulties. Some of these difficulties are now briefly discussed.

13.2.5.2 Numerical Considerations

Practical experience in solving the constrained minimization/maximization problem discussed above reveals that its numerical solution may require many model runs. As for the nonlinear parameter estimation problem which must precede it (in order to evaluate Φ_{\min}), an iterative procedure is required. Each iteration of this procedure is comprised of two steps. The first step involves calculation of a Jacobian matrix that replaces \mathbf{X} in the above equations and calculation of predictive sensitivities that replace \mathbf{y} . Generally, these sensitivities are computed using finite parameter differences. The second step requires calculation of parameter upgrade vectors using different trial values of the Marquardt lambda. As in equation (7.4.2), the Marquardt lambda term is added to the $\mathbf{X}^t \mathbf{Q} \mathbf{X}$ matrix of equation (13.2.8) before it is inverted. Conceptually, the Marquardt lambda should approach zero as the iterative solution process nears completion; furthermore it should start at a low value if this process commences with \mathbf{p} equal to \mathbf{p}_{\min} . The iterative solution process is terminated when parameters change no further between iterations, and Φ is suitably close to the user-specified value of Φ_0 .

In practice, it is found that many iterations may be required to solve the nonlinear constrained minimization/maximization problem that is defined above, even where the nonlinear solution procedure commences with \mathbf{p} equal to \mathbf{p}_{\min} . Experience also shows that its solution may require a line search along different directions of the parameter upgrade vector calculated using different values of the Marquardt lambda. In contrast to filling a Jacobian matrix using independent model runs based on finite parameter differences, a line search procedure is not inherently parallelizable. This further detracts from the numerical inefficiency of the solution process.

Numerical problems are exacerbated if model outputs are contaminated by numerical granularity arising from convergence difficulties that may be experienced by its numerical solver. This can lead to degradation of the quality of the Jacobian matrix, and of the predictive

sensitivity vector \mathbf{y} . The Φ and s contours of figure 13.1 can then become serrated. If the two sets of contours intersect at shallow angles, the location of the black dot depicted in this figure may then become difficult to find; in fact contour serrations may introduce multiple black dots.

13.2.5.3 Predictive Model Runs

In many decision-support contexts, the model run time required for the making of predictions far exceeds that required for calculation of outputs corresponding to members of a calibration dataset. This is because the time period over which measurements comprising a calibration dataset were taken is often relatively small in comparison to that over which management-predictions are required.

As already stated, the \mathbf{X} matrix and \mathbf{y} vector which appear in the above equations must be replaced by a different set of sensitivities during each iteration of the nonlinear minimization/maximization process. The “model” as seen by software such as PEST which implements this procedure must therefore be run over the measured past and into the predictive future. If a model has many parameters, if a prediction of interest is far into the future, and if the model takes a long time to run, the entire nonlinear process becomes computationally expensive.

There are, of course, some modelling circumstances where this is not the case. If the prediction of interest is the value of a parameter, or the value of a model output that is different from that used in the calibration process but that pertains to the calibration period, then the calibration and predictive model runs are the same.

If a modeller is interested in defining the two-sided confidence or prediction interval of a prediction, problems incurred by long model run times are exacerbated, for two complementary constrained optimization problems must then be solved – one in which the prediction of interest is maximized, and one in which it is minimized. If, in addition to this, prediction values corresponding to a number of different confidence levels are required, the constrained optimization problem must be solved for a number of different values of Φ_0 . If confidence intervals for more than one prediction are also required, then the entire process must be repeated for the other predictions.

13.2.5.4 Limiting Objective Function

As already discussed, any of equations (13.2.1), (13.2.2), (13.2.6) or (13.2.7) can be used to define the objective function Φ_0 that constrains the prediction minimization/maximization process. Where the number of observations comprising a calibration dataset \mathbf{h} is large, it will often be found that Φ_0 is very close to Φ_{\min} , this leaving very little “parameter wiggle room” in defining a predictive confidence interval. If employed, the predictive noise term then becomes the dominant contributor to predictive variability. Where a prediction is similar in character to at least some members of a calibration dataset \mathbf{h} , assistance in estimating the magnitude of this noise term will have been provided by the calibration process. Where it is different in character from members of a calibration dataset, a modeller must guess this noise term.

Reduction in constrained parameter variability that follows from increased size of a calibration dataset is an outcome of measurement noise cancellation. Conceptually, regardless of the magnitude of random error that accompanies each measurement of system state, if enough of these measurements have been made, a good estimate of model parameters can be made (if an inverse problem is well-posed). In contexts such as that of surface water model calibration where a daily or hourly streamflow time series may comprise a calibration dataset, use of

equations (13.2.1) and (13.2.2) may suggest that posterior parameter uncertainty is very small indeed, so that posterior predictive confidence is therefore very high. Intuition suggests, however that this cannot be the case.

As has been stated often in this text, in practical environmental modelling, most model-to-measurement misfit is an outcome of structural noise rather than measurement noise. As will be discussed later in this book, structural noise does not cancel. Hence increasing the number of measurements that comprise a model calibration dataset does not result in better estimates of its parameters if use of those parameters leads to computation of model outcomes that are persistently wrong. The problem of structural noise accommodation is exacerbated by the probable non-invertibility of the covariance matrix of structural noise (see later); this makes calculation of a suitable \mathbf{Q} matrix for use in model calibration, and in constrained predictive minimization/maximization, difficult. At the same time, association of a confidence level with a specific value of Φ_0 becomes virtually impossible.

Selection of a suitable Φ_0 constraint value is therefore generally subjective. In practice, a modeller may select a value that corresponds to visible deterioration of model-to-measurement misfit from that attained through minimization of Φ – enough deterioration for the modeller to deem that his/her model is “no longer calibrated”. While mathematical purity is thereby foregone, use of a rather arbitrary objective function constraint in solution of the predictive minimization/maximization problem posed above can still provide an effective means of establishing the amount of parameter and predictive variability that is consistent with expert knowledge on the one hand and respect for the information content of a calibration dataset on the other hand, with model structural defects taken into account.

It should be noted, however, that model structural imperfections can introduce further problems to model predictive uncertainty analysis. Equations presented above assume that predictive noise is independent of noise associated with a calibration dataset. In fact, the same model imperfections that degrade calibration fit can also impair its predictions. Predictive noise is therefore likely to be statistically correlated with measurement noise. Thus, for example, a model’s error in predicting future peak river flows will be correlated with its inability to fit past flow peaks. A modeller may therefore decide to forego formal inclusion of predictive noise in the predictive minimization/maximization process, preferring instead to add a heuristic “correction term” to a minimized/maximized prediction that reflects a model’s performance in fitting outputs of similar type as discovered during calibration.

13.2.5.5 Large Numbers of Parameters

Deployment of the predictive minimization/maximization procedure that is described above assumes inverse problem well-posedness. As discussed above, this can be attained using manual regularization. It can also be attained using Tikhonov regularization. However the regularization weight factor must be known in advance.

The numerical cost of constrained predictive minimization/maximization where parameter numbers are large has already been discussed. Hence fast model execution speed is important in these circumstances. It is worthy of note that run times of the DSI statistical model that is discussed in the following chapter enable predictive uncertainty analysis through predictive minimization/maximization.

13.2.6 A Heuristic Implementation

As has already been discussed, a modeller may decide to ignore formulas presented above for computation of Φ_0 , preferring instead to select a value that is based on visual evidence of

model-to-measurement misfit.

A third option is now suggested. This makes use of concepts that underpin calculation of the above limiting objective functions; however it implements these concepts in a different way.

Suppose that a modeller has just calibrated a model. In doing so, he/she has attained a set of “calibrated parameter values” and a complementary set of “calibrated model outputs” that are calculated using these parameters. The inverse problem may, or may not, be well posed.

A new inverse problem can then be easily set up. In this new inverse problem, initial parameter values are calibrated parameter values. Meanwhile, “observations” are model outputs that correspond to these calibrated parameter values; that is, they are “calibrated model outputs”. An observation weight matrix \mathbf{Q} is calculated as the inverse of the modeller’s conception of the covariance matrix of measurement noise $\mathbf{C}(\boldsymbol{\varepsilon})$; normally $\mathbf{C}(\boldsymbol{\varepsilon})$ is assumed to be diagonal. If the inverse problem is ill posed, prior information can be employed to provide a preferred value for one or a number of model parameters. These preferred values should be initial parameter values (which are calibrated parameter values). The weight matrix for these items of prior information should be formulated as the inverse of $\mathbf{C}(\mathbf{k})$, the prior parameter covariance matrix. (PEST allows a user to provide a covariance matrix instead of measurement weights; it inverts the matrix automatically.)

Obviously, with this choice of parameters, observations and prior information equations, the initial value of the objective function is zero.

To investigate the post-calibration uncertainty of a prediction of interest, constrained predictive minimization/maximization can now be undertaken. Two options are available for formulation of the objective function constraint Φ_0 pertaining to a $(1-\alpha)$ confidence level. The first is generous, whereas the second is tight.

The first objective function constraint is calculated as $\chi^2(n+m)_\alpha$ where n is the number of observations and m is the number of prior information equations that are featured in the inverse problem. (As stated above, the latter are only necessary if the inverse problem is ill-posed.) If measurement weights are assigned in the manner described above, the objective function computed by PEST is

$$\Phi = \delta\mathbf{h}^t \mathbf{C}^{-1}(\boldsymbol{\varepsilon}) \delta\mathbf{h} + \delta\mathbf{k}^t \mathbf{C}^{-1}(\mathbf{k}) \delta\mathbf{k} \quad (13.3.12)$$

In equation (13.3.12), $\delta\mathbf{h}$ and $\delta\mathbf{k}$ are perturbations from their calibrated values. This objective function is directly comparable with $\chi^2(n+m)$, so that the constraining objective function Φ_0 can be set to $\chi^2(n+m)_\alpha$ for a $(1-\alpha)$ two-sided predictive confidence interval.

Alternatively, in accordance with considerations presented in section 3.15, Φ_0 can be set to $N^2_{\alpha/2}$, where N signifies the normal distribution. To see why, rewrite (13.3.12) as

$$\Phi = \mathbf{r}_\varepsilon^t \mathbf{r}_\varepsilon + \mathbf{r}_k^t \mathbf{r}_k \quad (13.3.13)$$

where

$$\mathbf{r}_\varepsilon = \mathbf{C}^{-1/2}(\boldsymbol{\varepsilon}) \delta\mathbf{h} \quad (13.3.14a)$$

and

$$\mathbf{r}_k = \mathbf{C}^{-1/2}(\mathbf{k}) \delta\mathbf{k} \quad (13.3.14b)$$

Each element of both \mathbf{r}_ε and \mathbf{r}_k is an independent standard normal variate. The scalar projection of $\begin{bmatrix} \mathbf{r}_\varepsilon \\ \mathbf{r}_k \end{bmatrix}$ onto any unit vector in the same space is normally distributed. The objective function is the square of this projection. Using this criterion, if $(1-\alpha)$ is set to 99.7 percent, Φ_0 should be

set to 9.0. This is because 99.7 percent of the variability of a normal distribution falls within 3 standard deviations of its mean and 9 is the square of 3. Because of the way in which the inverse problem is formulated, the mean of \mathbf{r} is zero.

13.3 Simulator-Based Hypothesis-Testing

13.3.1 Calibration as Hypothesis-Testing

Hypothesis-testing and falsifiability lie at the heart of the Popperian view of the scientific method. A hypothesis is proposed, and evidence is collected to test it. If the evidence allows rejection of the hypothesis, then something has been learned; the range of possible explanations for a scientific phenomenon has thereby been reduced. However failure to reject a hypothesis does not imply truth of the hypothesis. Maybe there are other ways to reject it. Hence acceptance of a hypothesis as an explanation of reality can only ever be provisional.

These considerations are salient to our understanding of natural systems, particularly groundwater systems. A hypothesis that explains the behaviour of a system can never be *accepted* as truth, for there may be more than one explanation for its behaviour. However it may be possible to *reject* certain hypotheses that purport to explain the cause-and-effect relationships that govern past system behaviour and will govern future system behaviour. Hypothesis rejection can occur if a hypothesis is demonstrably incompatible with observations of system state, and/or with other knowledge of a system that has been previously acquired and has been demonstrated to have integrity.

A numerical model of environmental processes constitutes a hypothetical explanation for the behaviour of an environmental system. The hypothesis that is embodied in a numerical model can be assessed by attempting to calibrate the model. If the model is able to replicate historical system behaviour “reasonably well”, using parameters that are “reasonable” from an expert knowledge point of view, then the hypothesis that the model is a useable descriptor of environmental behaviour cannot be rejected. Note that in making this (mostly subjective) assessment, reasonableness of fit may need to be traded off against reasonableness of parameter values using L-curve concepts such as are discussed in section 8.3.

A lot hangs on the definition of “reasonable”. If a certain level of model-to-measurement misfit is deemed by a modeller to be tolerable, then knowledge of the noise associated with a calibration dataset is implied. Of course, this “noise” is partly measurement noise, and partly “structural noise” that is attributable to model defects. Whatever its source, the fact that the model is retained as a possibly valid descriptor of system behaviour implies at least an approximate definition of $C(\boldsymbol{\varepsilon})$, the covariance matrix of measurement/structural noise. (We retain the use of $\boldsymbol{\varepsilon}$ rather than $\boldsymbol{\tau}$ to denote combined measurement/structural noise to preserve notational consistency, even though $\boldsymbol{\tau}$ is used on several occasions in this book to acknowledge the contribution that structural noise makes to model-to-measurement misfit.)

Exploration of parameter and predictive uncertainty requires knowledge of $C(\boldsymbol{\varepsilon})$. Normally $C(\boldsymbol{\varepsilon})$ is denoted as diagonal in spite of the fact that this is demonstrably incorrect where model-to-measurement misfit reflects structural noise. However diagonal matrices are easy to store and invert. Furthermore, knowledge of off-diagonal elements is generally scanty or absent.

“Reasonableness” of values estimated for model parameters is also required if the hypothesis that a model is an adequate descriptor of environmental processes at a particular site is not to be rejected. Obviously, the values that must be assigned to a model’s parameters in order for model outputs to fit a calibration dataset must be judged according to metrics set by expert

knowledge. If the calibration process comprises a well-posed inverse problem, or has been enabled by a suitable regularization scheme that seeks conformity of parameter values with expert expectations, then unreasonableness of estimated parameter values indicates that parameters must adopt roles for which they were not designed in order for the model to provide an acceptable fit with field measurements. Nevertheless, for reasons that are outlined in chapter 15, some degree of parameter unreasonableness may be considered acceptable if it is judged by a modeller that the surrogate roles that some parameters must play if model outputs are to fit a calibration dataset do not introduce errors to model predictions of management interest. Alternatively, a modeller may judge that estimated system properties do in fact constitute a true reflection of actual system properties, even if he/she has been surprised by some of the properties that the calibration process has revealed. Furthermore, if unexpected heterogeneity has been revealed at one location within a model domain by collocated data, perhaps it exists at other places as well. In either of these cases, the decision not to reject the hypothesis that a model is a reasonable simulator of environmental behaviour at a particular study site may oblige a modeller to accept a $C(k)$ matrix of prior parameter covariance that expresses greater parameter variability than that which would have been expected on the basis of expert knowledge alone. A modeller has thus expanded the horizons of his/her expert knowledge through the calibration process. As for the $C(\epsilon)$ matrix, knowledge of the $C(k)$ matrix is fundamental to exploring the uncertainties associated with model predictions of future system behaviour.

Previous sections of this book have already discussed whether model calibration, as a precursor to model predictive uncertainty analysis, advances the decision-support potential of numerical simulation. There can be no universal answer to this question. However there are probably many circumstances where model calibration can be viewed as an activity that can raise the integrity of ensuing uncertainty analysis for reasons discussed above. At the same time, if properly conducted, it can raise the integrity of decision-support modelling as an implementation of the scientific method.

On the one hand, as stated above, model calibration can be viewed as a hypothesis-testing device through which the reasonableness of a particular model as a simulator of local system behaviour can be assessed. However, at the same time as this assessment is being made, the metrics for the making of this assessment can be simultaneously assessed. This constitutes a significant point of departure from the nature of scientific inquiry as it takes place in many other disciplines where the metrics for hypothesis rejection do not necessarily alter as hypotheses are being tested. However this departure is unavoidable given the layers of uncertainty that are associated with scientific knowledge of subsurface processes and properties, and given the simplicity of any environmental model compared with the system that it purports to simulate.

Where a model will be used to assist environmental management, the model calibration process enables a modeller to assess not just the strengths, but the weaknesses of his/her model as well. If the latter are not too great, and if they can be accounted for when a model is deployed in support of decision-making, then, as will be discussed, the model can be deployed in ways that enable the testing of hypotheses pertaining to future system behaviour. However its capacity to reject decision-pertinent hypotheses about the future must accommodate lessons that are learned by simulating the past.

The outcomes of a carefully-designed model calibration process are thus fourfold. First, it establishes whether a model can be viewed as a provisional descriptor of environmental processes at a particular study site; this view of the model is acceptable if the hypothesis that

the model constitutes an inadequate description of environmental behaviour at that site cannot be rejected. Second, model calibration awards values to parameters that are hopefully of minimized error variance, though probably far from correct. Third and fourth, the calibration process allows a modeller to acquire important knowledge of $C(\epsilon)$ and $C(k)$ that take into account not only measurement noise on the one hand and expert knowledge on the other hand, but (just as importantly) the imperfections of the model as a simulator of local environmental processes. Future decision-support usage of the model benefits from all four of these calibration outcomes.

Even if not formally reconstructed, knowledge of possible modifications to $C(\epsilon)$ and $C(k)$ that may be acquired through the calibration process are as significant as any other outcomes of this process. These (actually or notionally modified) matrices are now an integral part of the hypothesis that the calibration process has failed to reject. Specifically, the hypothesis that has been retained through calibrating the model is that the model constitutes a simulator of system behaviour to within a misfit tolerance that has been learned from this process (i.e. an actual or notional $C(\epsilon)$), and to within a parameter variability tolerance that has also been learned from this process (i.e. an actual or notional $C(k)$).

It is worth noting that in most practical modelling circumstances $C(k)$ is a simplistic descriptor of far more complex real-world system property variability. It follows that values assigned to elements of $C(k)$ should accommodate limitations in its ability to characterize true hydraulic property variability, at the same time as they may reflect the need for some parameters to adopt roles that compensate for model structural and parametric defects. Hence some elements of $C(k)$ may overstate the true variability of those system properties that the corresponding elements of k purport to represent. This is a necessary part of using $C(k)$ in conjunction with a numerical model. After all, there is no model without parameters; and there are no parameters without $C(k)$. Both the model and $C(k)$ are imperfect reflections of reality. However for them to support decision-making, they must complement each other in ways that accommodate their mutual imperfections.

13.3.2 Simulator-Aided Decision-Making

In most decision-making contexts, models are built to make predictions of future environmental behaviour under a management regime which will be different from that which prevailed in the past. Looming large in most decision-making contexts is the potential for something to go wrong following implementation of a new management regime. It follows that, for it to be considered as successful, future management of a system must preclude the occurrence of unwanted eventualities. These unwanted eventualities are normally well defined and acknowledged by all stakeholders.

Uncertainties that are associated with predictions of the environmental future (especially where the future involves stresses to which a system has not yet been subjected) are such that it may not be possible to preclude the occurrence of a specific unwanted event, or “bad thing”. Hence environmental management is accompanied by risk. Risk is often defined as the probability of an unwanted occurrence times the cost of its occurrence. Where the cost of something going wrong is low, then decision-makers can accept the fact that considerable uncertainty may be associated with predictions that it will not eventuate. In contrast, where the cost of something going wrong is high, good environmental management demands that the certainty of its non-occurrence be correspondingly high.

Management of risk is therefore integral to decision-making. Ideally, environmental modelling assists environmental management by associating probabilities with unwanted events. See

Freeze et al (1990) for an illuminating discussion of this issue.

Where decision-making is scientifically based, an unwanted event constitutes a hypothesis that a decision-maker would like to be able to reject - at a certain level of confidence. The unwanted event may be a "bad thing" from a financial point of view; for example a proposed mine dewatering system may allow more inflow to a mine than can be removed by installed pumps. Alternatively, an unwanted event may be a "bad thing" from an environmental point of view; for example, a proposed development plan may result in depletion of streamflow and the consequential creation of conditions that are harmful to stream biota. In another management context a "bad thing" may be that system managers have made the wrong decision by adopting management plan B over management plan A when in fact management plan A would have provided a better outcome according to prevailing decision criteria.

Unwanted events can be defined in arbitrarily complex ways that reflect the specific demands of environmental decision-making, and the uncertainties that are associated with environmental management, at a particular study site. For example, it may not be possible to reject the hypothesis that a certain development plan will have an unwanted environmental consequence. However it may be possible to reject the hypothesis that if development is accompanied by a suitable monitoring network to which developers will immediately respond by curtailing their activities if monitoring thresholds are crossed, then the bad thing will not happen. Rejection of this hypothesis may then allow development to proceed with the monitoring network in place. Meanwhile modelling must provide the assurance that the monitoring network is adequate by specifically testing and rejecting the hypothesis of its inadequacy.

With decision-making viewed in this way, the contribution that environmental modelling can make to the decision-making process is aligned with implementation of the scientific method. The task of modelling is to test the hypothesis that if a certain course of management action is followed, then the eventuality of a bad thing can be rejected at a certain level of confidence. The higher is the cost associated with the eventuality of the bad thing, the higher does this level of confidence need to be.

In the environmental management context (and in many other contexts) a hypothesis can be rejected if it is demonstrated to be incompatible with either or both of expert knowledge, and/or with the historical behaviour of the system whose management is being considered. The environmental model constitutes a unique scientific instrument for the testing of environmental hypotheses because of its ability to provide receptacles for these two types of information. Expert knowledge resides in the $C(\mathbf{k})$ matrix (or a more complex statistical descriptor of system property variability) that is used by a model. Information pertaining to the historical behaviour of a system is embodied in the parameter set \mathbf{k} achieved through the calibration process; as has already been discussed, \mathbf{k} embodies "the heterogeneity that *must* exist", given what is known about a system, and given historical measurements of that system's behaviour. The basis for rejection of a hypothesis pertaining to the occurrence of a future unwanted event therefore lies in demonstrating that this occurrence requires that the model be endowed with a parameter field \mathbf{k} that is not in accordance with either \mathbf{k} or $C(\mathbf{k})$, and/or that the parameter field that makes a bad thing possible makes it impossible for the model to replicate, to within a level set by $C(\epsilon)$, historical system behaviour. In other words, this occurs if the occurrence of a hypothesized bad thing demands that the model adopt values of $\mathbf{k} - \mathbf{k}$ that are too large (as judged by the $C(\mathbf{k})$ metric), or for which $\mathbf{k} - \mathbf{k}$ has too great a projection onto the calibration solution space (as judged by the $C(\epsilon)$ metric).

As $C(\mathbf{k})$ and $C(\epsilon)$ are both stochastic quantities it should be possible, in theory at least, to associate a level of confidence with rejection of a hypothesis that an unwanted event will

occur. In practice, however, because $C(\mathbf{k})$ and $C(\boldsymbol{\varepsilon})$ are approximations of more complex (but undefined) statistical entities, and because a modeller may have accepted the introduction of further complexities to these quantities following calibration of a model, the assignment of a low probability to the occurrence of a bad thing may need to be qualitative rather than quantitative. This cannot be avoided.

13.3.3 Configuring a Model for Hypothesis-Testing

If a model is tasked with testing the hypothesis that a future unwanted event will not occur under a proposed management plan, then it must be provided with stresses that allow it to simulate the outcomes of that plan. In other words it must be configured to predict values associated with the unwanted event. Presumably, for a particular prediction of management interest, a certain range of predictive values are “unwanted” and therefore constitute a “bad thing”, while values outside this range are considered to be benign.

As stated above, the hypothesis that a bad thing will occur can be rejected if its occurrence is demonstrated to be incompatible with past behaviour of a system, or requires that the model be endowed with unrealistic system properties. $C(\boldsymbol{\varepsilon})$ (actual or notional) constitutes the metric for the first comparison while $C(\mathbf{k})$ (actual or notional) constitutes the metric for the second comparison.

If the model prediction as made using the calibrated parameter set \mathbf{k} is bad, then the hypothesis that the bad thing will happen cannot be rejected (of course). In contrast, if the prediction made using the calibrated parameter set \mathbf{k} is benign, this does not constitute grounds for rejection of the hypothesis that the bad thing will occur. As has been extensively discussed, there is no reason to presume that a prediction made using the parameter set \mathbf{k} is correct; all that the calibration process can promise is that this prediction is minimally wrong. This does not mean, however, that the margin of potential wrongness is not high.

Rejection of the hypothesis that a bad thing will happen requires demonstration that the particular model prediction that corresponds to the bad thing is benign for any set of parameters that calibrate the model to within a level set by $C(\boldsymbol{\varepsilon})$, and that are reasonable from an expert knowledge point of view to a level set by $C(\mathbf{k})$.

Hypothesis-testing of this kind can be undertaken using Bayesian methods discussed in the previous chapter. In implementation of these methods, the pertinent model prediction is made using many realizations of \mathbf{k} drawn from the posterior parameter probability distribution. If all of these predictions are benign, the hypothesis that a bad thing will happen can be rejected. If only a handful of predictions are bad, the hypothesis that the bad thing will happen cannot be rejected; however it can be relegated a low level of likelihood of occurrence.

While this constitutes a conceptually correct strategy for testing the hypothesis that a bad thing will happen, it has a number of drawbacks. The first drawback is one that accompanies all Monte Carlo methods. Rejection of a hypothesis at a high level of confidence requires that many model run be undertaken, each based on different samples of the posterior parameter probability distribution. The higher the level of confidence for which hypothesis rejection is required, the greater is the number of model runs required to be certain of the confidence level. This number increases dramatically as parameter numbers increase and as the desired confidence level increases.

The second problem is that parameter fields used by ensemble methods often embody simplified versions of real-world heterogeneity. As has been discussed, this is an outcome of the requirement that parameters be adjustable, and that model outputs be continuous functions

of parameters. This may render history-match-constrained representation of structural and alluvial features that embody connected permeability difficult. False rejection of a bad thing hypothesis may therefore occur.

A third problem is the credibility of parameter realizations yielded by ensemble methods such as IES as samples of the posterior parameter probability distribution. Model-to-measurement fits associated with some of these samples may not be good; posterior predictive uncertainty may therefore be overstated.

These problems can be partially overcome by more direct implementation of simulation-based hypothesis-testing. The strategy that is now explained resembles, to some extent, calibration-constrained predictive minimization/maximization that is described earlier in this chapter.

Direct predictive hypothesis-testing is undertaken by asking a package such as PEST that is capable of highly-parameterized, regularized inversion to run a simulator under both calibration and predictive conditions. The normal calibration dataset is then expanded to include just one extra “observation”. This is the observation that a bad thing actually happens. The inversion engine is then asked to fit the calibration dataset (as it did when the model was calibrated) at the same time as it also fits the new “observation”. If the inversion engine is able to derive a parameter set that maintains a reasonable fit with the calibration dataset (as judged according to a formal or informal $C(\epsilon)$ metric), while maintaining parameter reasonableness (as judged by a formal or informal $C(k)$ metric), then the hypothesis that the bad thing will happen cannot be rejected.

This hypothesis-testing inversion process should use previously estimated parameter values k as starting parameters. Tikhonov constraints may be used (in conjunction with a suitable $C(k)$ matrix) to minimize parameter departures from these values as the expanded calibration dataset is fitted. The use of a large number of parameters, particularly in parts of the model domain that are most salient to the prediction of interest, allows heterogeneity to arise if it needs to arise for the bad thing to happen. Thus the occurrence of the bad thing is not precluded by inappropriate parameter parsimony in critical parts of the model domain. At the same time, exploration of an unwanted predictive occurrence through regularized inversion allows emergence of additional hydraulic property heterogeneity from that which is featured in the calibrated model only where necessary for its occurrence, and only to the minimum amount required for its occurrence. This reduces the possibility of false rejection of a bad thing hypothesis because inappropriate heterogeneity is construed as being required for its occurrence.

If the model is indeed able to simulate the occurrence of the bad thing, then assessment of the likelihood of its occurrence may be made either subjectively or in a statistically formal manner; this is the modeller’s choice. However, in many cases subjectivity is mandatory, this arising from the inadequate nature of any $C(\epsilon)$ as a descriptor of model-to-measurement misfit, and of any $C(k)$ as a descriptor of complex system hydraulic property heterogeneity. Meanwhile, the previous calibration process has taught the modeller much about what he/she will tolerate in the way of model-to-measurement misfit, and in the shapes and magnitudes of emergent heterogeneity.

By conducting a numerical experiment that is specifically designed to expose management failure, a modeller also learns much about possible failure mechanisms. Exposure of these mechanisms through failure-targeted inversion may change his/her mind about what comprises likely and unlikely patterns of heterogeneity, as he/she is not basing his assessment of failure likelihood on an ensemble of abstract expressions of a somewhat arbitrary, inflexible, and

possibly inappropriate heterogeneity based on a $C(\mathbf{k})$ matrix. Heterogeneity that is introduced through the hypothesis-testing process is failure-salient. Whether, or not, a failure hypothesis is deemed to be unlikely, emergent patterns of heterogeneity can suggest appropriate locations for monitoring points that can provide early warning of failure, should the worst occur.

If a more formal, statistically-based, criterion is sought for rejection of a failure hypothesis, theory provided in section 13.2.6 may prove useful. See also Moore et al (2010). It is the author's opinion however that the inadequacies of $C(\epsilon)$ and $C(\mathbf{k})$ as statistical descriptors of model-to-measurement misfit on the one hand, and expert knowledge of hydraulic properties on the other hand, makes subjective assessment inescapable.

Implementation of the direct hypothesis-testing methodology described in this section requires that a weight be assigned to the "bad thing observation" in the inversion process through which it is tested. Some trial and error may be required in assigning this weight. So too may some trial and error be required in assigning an "observed bad value" to the observation that is added to the calibration dataset. Ideally, the weight assigned to this observation should provide a strong incentive for the inversion process to actually make the bad thing happen. However if this weight is too large, then perhaps the inversion process may ignore parts of the remainder of a calibration dataset.

Some of these problems can be circumvented if PEST is run in "Pareto" mode. The weight ascribed to the observed bad thing is then gradually applied. At first this weight can be zero (it is the modeller's choice). If the initial parameter set is equal to the parameter set $\underline{\mathbf{k}}$ achieved through previous model calibration, then the hypothesis-testing process will commence with the model calibrated. The weight assigned to the observed bad thing is then gradually increased. As this happens, model parameters start to depart from $\underline{\mathbf{k}}$; however departures will be kept to a minimum, particular if prior information which limits perturbations of \mathbf{k} from $\underline{\mathbf{k}}$ is included in the hypothesis-testing inversion procedure. As the weight ascribed to the prediction is increased even further, the calibration component of the objective function may (or may not) start to rise as parameters depart even further from their preferred $\underline{\mathbf{k}}$ values. At some stage a modeller may judge that parameter or calibration credibility has been strained to breaking point. The process can then be halted. The hypothesis that the bad thing will happen can be rejected.

If PEST is used to traverse the so-called "Pareto front" in this manner, it will record parameter values and corresponding model outputs at all points along its journey. A modeller can therefore associate credibility with each parameter set on the basis of a supplied or implied $C(\mathbf{k})$, and with each set of calibration-pertinent model outputs on the basis of a supplied or implied $C(\epsilon)$ as he/she sees fit. On this basis he/she may therefore attempt to establish the value of the unwanted prediction at which its occurrence becomes unlikely.

14. Data Space Inversion

14.1 Introduction

Like all of the methods that have been discussed so far in this book, the purpose of data space inversion (DSI) is to support the use of numerical simulation in environmental decision-making. As has been much discussed herein, numerical simulation cannot support environmental decision-making on its own. Simulation must be part of an enterprise whose intention is to harvest information from data (thereby reducing predictive uncertainty) and to quantify the predictive repercussions of information insufficiency (that is, the uncertainty that remains after harvesting of all available information). Numerical simulation is a necessary part of this enterprise because it provides repositories for harvested information. These repositories are a model's parameters.

DSI is different from other data assimilation methods that have been discussed so far because it does not use the parameters of a numerical simulator to host data-harvested information. Instead, it uses the parameters of a surrogate statistical model. This makes data assimilation and uncertainty quantification very fast indeed. What makes DSI even more useful is that this statistical model embodies direct linkages between measurements of past system behaviour and predictions of future system behaviour. This is tremendously liberating because the numerical model that is used to build these linkages does not need to be equipped with "parameters". Recall that parameters are, by definition, adjustable; hence they often comprise simplistic representations of subsurface heterogeneity. Instead, the numerical model that is used to build the DSI statistical model can be equipped with realizations of subsurface properties that are of arbitrary geostatistical complexity. These realizations can be continuous or categorical, or any combination of the two. At the same time, uncertainties of geostatistical hyperparameters can also be accommodated; different realizations of hydraulic properties can be generated using different geostatistical hyperparameters.

Of course, hydraulic properties are rarely the only aspects of a model's representation of an environmental system that are uncertain. Details of boundary conditions and historical stresses may also be incompletely known. In conventional model history-matching these can also be parameterized so that they can be calibrated and/or their posterior probability distributions sampled during history-matching. Their prior uncertainties can also be accommodated when DSI is employed for data assimilation. However the data assimilation process becomes easier because the probability distributions that are sampled to generate different realizations of these model attributes can be of arbitrary complexity. Meanwhile, no aspect of the numerical model is actually adjusted; however any model feature that is uncertain can be sampled and represented stochastically in the model.

DSI is an ensemble method. As such, it has some of the advantages and some of the disadvantages of other ensemble methods. Integrity of DSI outcomes requires that ensembles be populated with sufficient realizations to ensure that empirical covariance matrices are not contaminated by spurious correlations. However the number of realizations required to ensure correlation integrity may not be as large as for methods such as IES. This depends on the predictions which are the focus of management interest. If these reflect only some aspects of a system's behaviour, then it may be possible to construct empirical past-to-future covariance matrices using fewer model runs than IES may require to adjust an entire model parameter field.

Other factors also contribute to the light numerical burden of DSI deployment. Covariance matrices that link the measured past to the predicted future are filled only once. Their construction does not need to be repeated during an iterative parameter adjustment process. The numerical burden of DSI is therefore the lightest of any of the data assimilation methods that are discussed in this book.

Use of DSI is accompanied by other advantages. There are many circumstances where the integrity of DSI-calculated posterior predictive confidence limits is likely to exceed that of other methods that are described in this book. Because a complex simulator does not need to be equipped with parameters that must undergo adjustment in order to assimilate information, fewer compromises in geostatistical representation of subsurface hydraulic properties are required. Meanwhile, because the number of numerical model runs that is required for implementation of DSI is fewer than that required by other methods, a modeller has the option of using a complex, slower-running simulator if he/she feels that this is necessary to ensure the integrity of model predictions. Another advantage of DSI is that in many decision-support circumstances, statistical linkages between the past and the future are more linear than those between field measurements and parameters on the one hand, and parameters and predictions of future system behaviour on the other hand. The role of model nonlinearity in eroding the Bayesian credentials of ensemble-based parameter adjustment has been discussed in previous chapters.

Of course, advantages come with disadvantages. DSI theory rests on the assumption that statistical relationships between the past and the future can be encapsulated in a multiGaussian probability distribution; implied in this is a linear relationship between the past and the future. As we shall see, steps can be taken to accommodate violation of this assumption. However the success of these ameliorative actions may be limited under circumstances of grossly nonlinear past-to-future linkages that may precipitate multi-modal predictive probability distributions.

The performance of DSI also suffers if there is prior-data conflict. This is because the integrity of DSI analysis rests on the premise that the diversity of numerical model outputs, calculated using many different stochastic realizations of system hydraulic properties, is sufficient to span the true behaviour of an investigated environmental system. As is described at the end of this chapter, analysis of prior-data conflict is easily undertaken using the same input dataset as that which is required for implementation of DSI.

14.2 Linear DSI

14.2.1 General

We commence our investigation of DSI by considering a linear variant of the methodology. In the next subsection, complexities are added that assist in accommodation of more complex system behaviour.

14.2.2 Concepts

As usual, our analysis begins with the same equation that has already appeared many times in this book. It describes the action of a linearized model under calibration conditions.

$$\mathbf{h} = \mathbf{o} + \boldsymbol{\varepsilon} = \mathbf{Zk} + \boldsymbol{\varepsilon} \quad (14.2.1)$$

In equation (14.2.1), the vector \mathbf{h} is comprised of field measurements of system behaviour while the vector \mathbf{o} is comprised of model-calculated counterparts to these measurements. The matrix \mathbf{Z} embodies the action of a model on its parameters \mathbf{k} , while the vector $\boldsymbol{\varepsilon}$ contains an

unknown realization of measurement noise that accompanies the particular measured \mathbf{h} that comprises the history-matching dataset.

Let the vector \mathbf{s} represent a set of predictions made by the model. In earlier sections of this book, the following equation has been used to represent a single prediction.

$$s = \mathbf{y}^t \mathbf{k} \quad (14.2.2a)$$

For multiple predictions, the scalar s becomes a vector, and the vector \mathbf{y} becomes a matrix, so that

$$\mathbf{s} = \mathbf{Y}^t \mathbf{k} \quad (14.2.2b)$$

The vector \mathbf{k} hosts a model's parameters. These can be conceptualized as being any aspect of a simulated system that is incompletely known, and hence must be expressed probabilistically. These include hydraulic properties, boundary conditions, historical stresses, and possibly even structural features such as the existence, or otherwise, of an aquitard. In the present circumstances, parameters do not need to be adjustable. They can be categorical in nature. They can acknowledge the fact that existence, or nonexistence, of certain subsurface features is a matter of conjecture. Nor is perfect knowledge of the prior parameter probability distribution required. Hence geostatistical and other hyperparameters can be included in \mathbf{k} in order to allow stochastic representation of uncertainty in prior uncertainty.

Random samples of \mathbf{k} are now generated. We denote a random realization of the vector \mathbf{k} as \mathbf{k}_i . For each realization of \mathbf{k} , the model is run in order to generate a realization \mathbf{o}_i of model outputs \mathbf{o} . A realization of measurement noise $\mathbf{\epsilon}_i$ is then added to \mathbf{o}_i to generate a realization of field measurements \mathbf{h}_i . After this has been done many times, enough samples are available to build some empirical covariance matrices. Construction of empirical covariance matrices is described in section 3.5.2. We wish to build a covariance matrix that links model-generated field measurements \mathbf{h} to model predictions \mathbf{s} .

Let columns of the matrix \mathbf{H} comprise realizations of model-generated field measurements \mathbf{h}_i obtained in the manner described above. Now define $\underline{\mathbf{H}}$ as follows (where N is the number of realizations)

$$\underline{\mathbf{H}} = \frac{\mathbf{H}^1}{N} \quad (14.2.3a)$$

That is, each column $\underline{\mathbf{h}}$ of $\underline{\mathbf{H}}$ is identical to every other column $\underline{\mathbf{h}}$ of $\underline{\mathbf{H}}$. The i 'th element of each column of $\underline{\mathbf{H}}$ is the average of the i 'th element over all columns of \mathbf{H} . That is

$$\underline{h}_i = \frac{\sum_{j=1}^N h_{ij}}{N} \quad (14.2.3b)$$

We define the matrix $\underline{\mathbf{S}}$ in a similar way

$$\underline{\mathbf{S}} = \frac{\mathbf{S}^1}{N} \quad (14.2.4)$$

Empirical covariance matrices can then be calculated as follows

$$C_{hh}^e = \frac{(\mathbf{H} - \underline{\mathbf{H}})(\mathbf{H} - \underline{\mathbf{H}})^t}{N-1} \quad (14.2.5a)$$

$$C_{hs}^e = \frac{(\mathbf{H} - \underline{\mathbf{H}})(\mathbf{s} - \underline{\mathbf{S}})^t}{N-1} \quad (14.2.5b)$$

$$C_{sh}^e = \frac{(\mathbf{s} - \underline{\mathbf{S}})(\mathbf{H} - \underline{\mathbf{H}})^t}{N-1} \quad (14.2.5c)$$

$$C_{ss}^e = \frac{(\mathbf{s} - \underline{\mathbf{s}})(\mathbf{s} - \underline{\mathbf{s}})^t}{N-1} \quad (14.2.5d)$$

These are submatrices of a partitioned covariance matrix defined as follows

$$C \begin{pmatrix} \mathbf{h} \\ \mathbf{s} \end{pmatrix} = \begin{bmatrix} C_{hh} & C_{hs} \\ C_{sh} & C_{ss} \end{bmatrix} \quad (14.2.6)$$

We now condition this covariance matrix based on what we know. Conditioning equations are presented in section 3.8. Direct application of equations (3.8.4) and (3.8.5) provides estimates of the conditional (i.e. posterior) mean of predictions \mathbf{s} and the posterior covariance matrix of these same predictions.

Define $\underline{\mathbf{h}}$ as the prior mean of \mathbf{h} (this vector occupies all columns of $\underline{\mathbf{H}}$) and $\underline{\mathbf{s}}$ as the prior mean of \mathbf{s} (this vector occupies all columns of $\underline{\mathbf{S}}$). Let the vector \mathbf{h}_m host field measurements. The posterior mean of \mathbf{s} is given by

$$\underline{\mathbf{s}}' = \underline{\mathbf{s}} + C_{sh}^e C_{hh}^{e-1} (\mathbf{h}_m - \underline{\mathbf{h}}) \quad (14.2.7)$$

The posterior covariance matrix of predictions is given by

$$C'_{ss} = C_{ss}^e - C_{sh}^e C_{hh}^{e-1} C_{hs}^e \quad (14.2.8a)$$

For an individual prediction s this becomes

$$\sigma_s^{2'} = \sigma_s^2 - C_{sh}^e C_{hh}^{e-1} C_{hs}^e \quad (14.2.8b)$$

Note that in equation (14.2.8b), the “ s ” subscript in C_{sh}^e is italicized. This matrix holds covariances between the individual prediction s and model-calculated observations \mathbf{h} .

It is apparent from the above equations that with a few hundred runs of an arbitrary complex model (enough to build an empirical covariance matrix that links the measured past to the predicted future), it is possible to make data-informed predictions of future system behaviour, and to quantify the uncertainties of these predictions. No parameter adjustment is necessary.

Equations (14.2.7) and (14.2.8) assume a multiGaussian distribution. This assumption is generally violated in real-world modelling practice. However, Gaussian behaviour can be approximated if individual model-generated \mathbf{h} ’s and \mathbf{s} ’s are subjected to normal score transformation before being used in the above equations. Post-transformed elements of model-generated \mathbf{h} and \mathbf{s} then have distributions that approach $N(0,1)$. Normal score transformation is schematized in Figure 14.1.

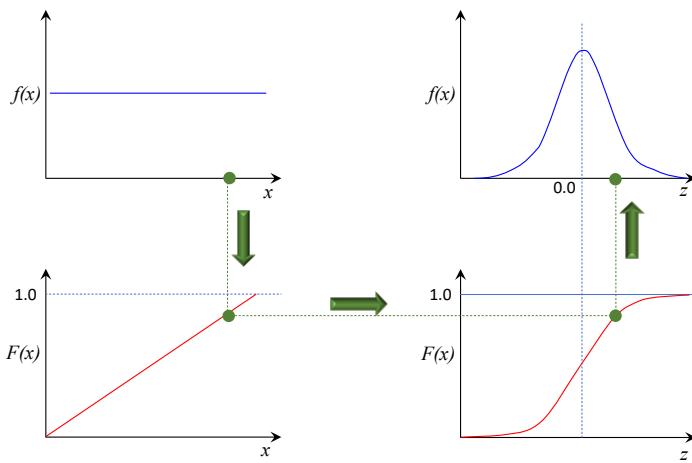


Figure 14.1 Schematic of normal score transformation of a uniformly distributed variable x to a standard normal variate z .

14.2.3 Comparison with other Methods

It is instructive to compare the outcomes of equations (14.2.7) and (14.2.8) with other data assimilation and predictive uncertainty quantification methods discussed in this book that do not bypass consideration of model parameters \mathbf{k} .

Equations (14.2.1) and (14.2.2b) can be written as a single equation.

$$\begin{bmatrix} \mathbf{h} \\ \mathbf{s} \end{bmatrix} = \begin{bmatrix} \mathbf{Z} & \mathbf{I} \\ \mathbf{Y}^t & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{k} \\ \boldsymbol{\varepsilon} \end{bmatrix} \quad (14.2.9)$$

Using the standard formula for propagation of variance (3.9.2), the joint covariance matrix of \mathbf{h} and \mathbf{s} can be written as

$$\mathbf{C} \left(\begin{bmatrix} \mathbf{h} \\ \mathbf{s} \end{bmatrix} \right) = \begin{bmatrix} \mathbf{Z} & \mathbf{I} \\ \mathbf{Y}^t & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{C}(\mathbf{k}) & \mathbf{0} \\ \mathbf{0} & \mathbf{C}(\boldsymbol{\varepsilon}) \end{bmatrix} \begin{bmatrix} \mathbf{Z}^t & \mathbf{Y} \\ \mathbf{I} & \mathbf{0} \end{bmatrix} \quad (14.2.10)$$

After multiplying out, this becomes

$$\mathbf{C} \left(\begin{bmatrix} \mathbf{h} \\ \mathbf{s} \end{bmatrix} \right) = \begin{bmatrix} \mathbf{ZC}(\mathbf{k})\mathbf{Z}^t + \mathbf{C}(\boldsymbol{\varepsilon}) & \mathbf{ZC}(\mathbf{k})\mathbf{Y} \\ \mathbf{Y}^t\mathbf{C}(\mathbf{k})\mathbf{Z}^t & \mathbf{Y}^t\mathbf{C}(\mathbf{k})\mathbf{Y} \end{bmatrix} \quad (14.2.11)$$

Applying equation (14.2.7) to this matrix we find

$$\underline{\mathbf{s}}' = \underline{\mathbf{s}} + \mathbf{Y}^t\mathbf{C}(\mathbf{k})\mathbf{Z}^t[\mathbf{ZC}(\mathbf{k})\mathbf{Z}^t + \mathbf{C}(\boldsymbol{\varepsilon})]^{-1}(\mathbf{h}_m - \underline{\mathbf{h}}) \quad (14.2.12a)$$

If the prior predictive means of \mathbf{s} and \mathbf{h} are set to $\mathbf{0}$ so that these vectors express perturbations from prior means (as is done throughout this book) equation (14.2.12a) becomes

$$\underline{\mathbf{s}}' = \mathbf{Y}^t\mathbf{C}(\mathbf{k})\mathbf{Z}^t[\mathbf{ZC}(\mathbf{k})\mathbf{Z}^t + \mathbf{C}(\boldsymbol{\varepsilon})]^{-1}\mathbf{h}_m \quad (14.2.12b)$$

Application of (14.2.8) yields

$$\mathbf{C}'_{ss} = \mathbf{C}_{ss} - \mathbf{Y}^t\mathbf{C}(\mathbf{k})\mathbf{Z}^t[\mathbf{ZC}(\mathbf{k})\mathbf{Z}^t + \mathbf{C}(\boldsymbol{\varepsilon})]^{-1}\mathbf{ZC}(\mathbf{k})\mathbf{Y} \quad (14.2.13a)$$

From (14.2.6), the prior predictive covariance matrix \mathbf{C}_{ss} is given by the matrix that occupies the bottom right partition of the matrix on the right side of equation (14.2.11). Hence

$$\mathbf{C}'_{ss} = \mathbf{Y}^t\mathbf{C}(\mathbf{k})\mathbf{Y} - \mathbf{Y}^t\mathbf{C}(\mathbf{k})\mathbf{Z}^t[\mathbf{ZC}(\mathbf{k})\mathbf{Z}^t + \mathbf{C}(\boldsymbol{\varepsilon})]^{-1}\mathbf{ZC}(\mathbf{k})\mathbf{Y} \quad (14.2.13b)$$

Equation (14.2.12b) is the same as (10.3.6) if the latter is pre-multiplied by predictive

sensitivities. Equation (14.2.13b) is the same as (10.4.1). It follows that, conceptually at least, conditioning of empirical covariance matrices that directly link measurements of the past to predictions of the future yields the same results as those obtained using equations that do not bypass parameters.

Where parameters are many, model run times are long, and where sensitivities are calculated using finite differences, the filling of matrices that appear in equations (14.2.12) and (14.2.13) is expensive. Additionally, use of these equations requires that parameters be actually defined, and that model outputs with respect to parameters be continuous. In contrast, deployment of equations (14.2.7) and (14.2.8) has no such limitations. The numerical cost of their use is comprised of model runs that are required to calculate empirical covariance matrices on which they are based, regardless of the number of parameters with which a model is endowed. Furthermore, hydraulic properties can be represented nonparametrically using sophisticated geostatistical realizations; as stated above, uncertainties in prior uncertainties can be echoed in these realizations.

14.2.4 Data Worth Analysis

A notable feature of equation (14.2.8) is that it contains the values of neither parameters nor model outputs. It can therefore be used to estimate post-history-matching predictive uncertainty, regardless of whether history-matching has actually been undertaken, and regardless of whether field measurements that comprise some or all elements of \mathbf{h} have actually been made. In doing so, it can be used to assess the ability of data that have not yet been acquired to reduce the uncertainties of decision-critical model predictions. To do this, equation (14.2.8) is used to evaluate the uncertainty of a key prediction with and without inclusion of various members of \mathbf{h} in this calculation. The difference in calculated uncertainty of a prediction of interest is a measure of the information content of these members of \mathbf{h} with respect to the prediction.

Note that apparent lack of prediction-pertinent information in a subset of \mathbf{h} may be an outcome of either of two factors. Perhaps this subset of \mathbf{h} is indeed information-free with respect to the prediction of interest. Alternatively, its information content with respect to the prediction may duplicate that of other subsets of \mathbf{h} .

An alternative assessment of data worth can be made using equation (14.2.7). The expected value of a prediction of interest can be made with and without inclusion of various members of \mathbf{h}_m in a history matching dataset. These measurements may be real, or they may be model-generated. The difference in the expected value of a prediction that is evaluated in this way is a measure of the ability of information that is resident in the tested values of \mathbf{h} to alter the value of that prediction from its prior mean. Where \mathbf{h}_m does not yet exist, then these calculations should be repeated many times using different stochastic realizations of \mathbf{h}_m calculated by the model. These “measurements” can be gathered from the outputs of numerical model runs that were used to build the covariance matrices that appear in preceding equations.

Regardless of which of the above options for data worth assessment is employed, the numerical cost of this assessment is remarkably cheap. All that is required are that covariances be calculated between model outputs that correspond to data whose worth is being assessed on the one hand, and predictions of management interest on the other hand. As has already been discussed, this may require only a few hundred runs of a complex simulator, regardless of hydraulic property complexity that is represented in that simulator. Unlike linear data worth analysis that is discussed in section 10.5, sensitivities of model outputs with respect to model parametric detail are not required. The cost of this latter type of linear analysis increases with

the level of detail that must be represented, and hence with the number of parameters that are assigned to a model in order to embody this detail.

Use of equations (14.2.7) and (14.2.8) in data worth analysis is discussed by He et al (2018). These authors urge caution in choosing the number of realizations used in evaluation of covariance matrices that feature in these equations. They point out that the integrity of data worth assessment can be checked by increasing the number of realizations and monitoring for convergence of assessed posterior predictive uncertainties as calculations are repeated.

14.3 Nonlinear DSI

14.3.1 General

Deployment of concepts that are discussed in the preceding section can be expanded if some alterations are made to formulation of the governing equations. Workflows that result from use of these modified equations can tolerate greater departures from the assumption of Gaussianity than is tolerable with direct use of equations (4.2.7) and (4.2.8). The methodology that is now discussed, as well as variations of it, are described and demonstrated by Sun and Durlofsky (2017), Lima et al (2020), Liu et al (2021), Delottier et al (2023), Kitlasten et al (2025), Doherty and Lotti (2025) and papers cited therein.

14.3.2 Statistical Model

14.3.2.1 Defining Equations

Alternative formulation of a DSI workflow begins by building a statistical model emulator.

In the previous section of this chapter it was shown how an empirical covariance matrix $C\left(\begin{bmatrix} \mathbf{h} \\ \mathbf{s} \end{bmatrix}\right)$ can be filled from the outcomes of an ensemble of numerical model runs. Recall that realizations of measurement noise $\boldsymbol{\epsilon}$ are added to model outputs \mathbf{o} in order to produce realizations of \mathbf{h} . In the present workflow this step is bypassed as measurement noise is accommodated later. So the workflow commences by building an empirical covariance matrix $C\left(\begin{bmatrix} \mathbf{o} \\ \mathbf{s} \end{bmatrix}\right)$ that links model outputs that correspond to field measurements to model predictions of management interest. Otherwise, steps are the same as for the previous workflow.

The data assimilation problem is then formulated as follows.

$$\begin{bmatrix} \mathbf{h} \\ \mathbf{s} \end{bmatrix} = C^{1/2} \left(\begin{bmatrix} \mathbf{o} \\ \mathbf{s} \end{bmatrix} \right) \mathbf{x} + \begin{bmatrix} \mathbf{o} \\ \mathbf{s} \end{bmatrix} + \begin{bmatrix} \boldsymbol{\epsilon} \\ \mathbf{0} \end{bmatrix} \quad (14.3.1)$$

The “model” of equation (14.3.1) is $C^{1/2} \left(\begin{bmatrix} \mathbf{o} \\ \mathbf{s} \end{bmatrix} \right)$. This model has its own set of parameters. These are encapsulated in the vector \mathbf{x} . The prior mean and covariance matrix of \mathbf{x} are

$$\underline{\mathbf{x}} = \mathbf{0} \quad (14.3.2a)$$

$$C(\mathbf{x}) = \mathbf{I} \quad (14.3.2b)$$

The statistics of $\boldsymbol{\epsilon}$ are unchanged from previous formulations of the inverse problem. That is, $\boldsymbol{\epsilon}$ has a mean of $\mathbf{0}$ and a covariance matrix of $C(\boldsymbol{\epsilon})$.

It is easy to show that the model described by equation (14.3.1) is a useful replacement for the model described by equations (14.2.1) and (14.2.2b). To show this, we first set \mathbf{x} and $\boldsymbol{\epsilon}$ to their expected values (i.e. their prior mean values) of $\mathbf{0}$ to obtain

$$E \begin{pmatrix} \mathbf{h} \\ \mathbf{s} \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ \mathbf{s} \end{pmatrix} \quad (14.3.3)$$

This is obviously correct. Now let us determine the covariance matrix on the left side of equation (14.3.1) using the propagation of variance formula (3.9.2) (while noting that variances of mean values are $\mathbf{0}$ and that the square root of a covariance matrix is symmetrical).

$$C \begin{pmatrix} \mathbf{h} \\ \mathbf{s} \end{pmatrix} = C^{e/2} \begin{pmatrix} \mathbf{0} \\ \mathbf{s} \end{pmatrix} C(\mathbf{x}) C^{e/2} \begin{pmatrix} \mathbf{0} \\ \mathbf{s} \end{pmatrix} + \begin{pmatrix} C(\boldsymbol{\varepsilon}) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} \quad (14.3.4)$$

From (14.3.2b) this becomes

$$C \begin{pmatrix} \mathbf{h} \\ \mathbf{s} \end{pmatrix} = C^e \begin{pmatrix} \mathbf{0} \\ \mathbf{s} \end{pmatrix} + \begin{pmatrix} C(\boldsymbol{\varepsilon}) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} \quad (14.3.5)$$

This equation too is obviously correct to the extent that $C^e \begin{pmatrix} \mathbf{0} \\ \mathbf{s} \end{pmatrix}$ is a good representation of $C \begin{pmatrix} \mathbf{0} \\ \mathbf{s} \end{pmatrix}$, that is to the extent that the empirical covariance matrix that links model outputs corresponding to field measurements to model predictions reflects the actual covariance of these quantities. The integrity of $C^e \begin{pmatrix} \mathbf{0} \\ \mathbf{s} \end{pmatrix}$ increases with the number of numerical model runs on which this empirical covariance matrix is based. However, regardless of the number of model runs used in its computation, its integrity is questionable if the relationship between \mathbf{o} and \mathbf{s} is highly nonlinear.

Use of equation (14.3.1) raises the question of how the square root of an empirical covariance matrix should be calculated.

First $C^e \begin{pmatrix} \mathbf{0} \\ \mathbf{s} \end{pmatrix}$ must be filled. This is done by running the model many times using different realizations of hydraulic properties (and all other facets of the model that are incompletely known, and hence require stochastic characterization). The empirical covariance matrix is then calculated as

$$C^e \begin{pmatrix} \mathbf{0} \\ \mathbf{s} \end{pmatrix} = \frac{1}{N-1} \begin{bmatrix} (\mathbf{o} - \underline{\mathbf{o}})(\mathbf{o} - \underline{\mathbf{o}})^t & (\mathbf{o} - \underline{\mathbf{o}})(\mathbf{s} - \underline{\mathbf{s}})^t \\ (\mathbf{s} - \underline{\mathbf{s}})(\mathbf{o} - \underline{\mathbf{o}})^t & (\mathbf{s} - \underline{\mathbf{s}})(\mathbf{s} - \underline{\mathbf{s}})^t \end{bmatrix} \quad (14.3.6)$$

In equation (14.3.6), matrices with underscores have the same significance as described above; their columns are means over all realizations. Now subject $C^e \begin{pmatrix} \mathbf{0} \\ \mathbf{s} \end{pmatrix}$ to singular value decomposition so that

$$C^e \begin{pmatrix} \mathbf{0} \\ \mathbf{s} \end{pmatrix} = \mathbf{E} \mathbf{F} \mathbf{E}^t \quad (14.3.7)$$

The square root of this matrix is easily formulated as

$$C^{e/2} \begin{pmatrix} \mathbf{0} \\ \mathbf{s} \end{pmatrix} = \mathbf{E} \mathbf{F}^{1/2} \mathbf{E}^t \quad (14.3.8)$$

The matrix \mathbf{F} is diagonal. It contains positive or zero singular values. Formulation of $\mathbf{F}^{1/2}$ requires that the square root be taken of these diagonal elements. It is easily verified that if the right side of (14.3.8) is multiplied by itself, then the right side of (14.3.7) is obtained.

It should be noted that a real (i.e. non-empirical) covariance matrix is positive definite. Therefore, all of its singular values are positive. The same does not apply to an empirical covariance matrix where the number of realizations used in its calculation is less than the number of elements of the vector whose stochasticity it is meant to describe. In this case, the

empirical covariance matrix is of diminished rank. Hence some of its singular values are zero; it can only therefore be positive semidefinite. Note that singular values are arranged from highest to lowest down the diagonal of the \mathbf{F} matrix.

In practice, a modeller may decide to truncate singular values before they become zero. This protects an empirical covariance matrix from expressing spurious correlations that arise from limited sampling. Often the concept of “energy” is used to determine the truncation point. The total “energy” of a covariance matrix can be conceptualized as the sum of its squared singular values. A modeller may decide to retain, for example, 95% of the total energy of a covariance matrix. In this case, only enough singular values that are close to zero are declared to be zero in order for this level of energy to be retained; those of lowest value are truncated first.

14.3.2.2 Using the Statistical Model

Integrity of the statistical model emulator of equation (14.3.1) can often be improved if numerical model outputs are subjected to normal score transformation before calculating an empirical covariance matrix from these outputs. Outputs of the statistical model are then back-transformed to native model output space before being compared with field measurements, or before being used to make predictions. However, while normal score transformation promotes Gaussianity of individual model outputs, it does not necessarily achieve Gaussianity of cross-covariance terms. This can be accommodated by subjecting the model of (14.3.1) to calibration and/or Bayesian history-matching when deploying it in place of the numerical model.

Bayesian history-matching of the surrogate DSI statistical model yields the posterior probability distribution of its parameters \mathbf{x} . History-matching is achieved by comparing its outputs to field measurements \mathbf{h}_m , just as if it were a numerical model.

Actually, history-matching of the DSI surrogate model can be conducted in a number of ways. Because the surrogate model runs so fast (generally a small fraction of a second, as it only requires that a vector be multiplied by a matrix) Markov chain Monte Carlo is a history-matching option. Alternatively, Bayesian history-matching can be undertaken using the PESTPP-IES ensemble smoother.

Yet another alternative that is provided by the PEST suite is to first calibrate the DSI statistical model using regularized inversion. This provides minimum error variance estimates of statistical model parameters \mathbf{x} and (more importantly) minimum error variance statistical model predictions \mathbf{o} . DSI support software provided with the PEST suite then enables linear predictive uncertainty quantification of DSI statistical model parameters and outputs using equations that are presented in chapter 10. This process is implemented in Gaussian-transformed space, where use of these equations incurs little or no error. Confidence intervals are then back-transformed to native parameter and predictive space. This approach to DSI statistical model history-matching is fast and stable. Furthermore, it can mitigate the introduction of artificial temporal and spatial irregularities to DSI statistical model predictions (see below). See PEST documentation of the DSIMOD and POSTDSIMOD utilities for more details.

When history-matching a DSI statistical model, the same considerations apply to choice of observation weights as those that apply when history-matching a numerical model. As is discussed in section 12.4.2, and as is further discussed in the following chapter, a distinction may need to be made between the role of weights and the way in which history-matching accommodates measurement noise. Where a model is perfect, no such distinction is necessary. In this case, measurement weights should be inversely proportional to the standard deviations of measurement noise. Where a model is imperfect, weights should be chosen to ensure

visibility of different groups of observations in the overall objective function that is reduced through history-matching. This ensures harvesting of information that each of these groups contain. Meanwhile, noise that is associated with field measurements must be characterized so that realizations of it can be added to field measurements \mathbf{h}_m (if using IES), or so that it can be used in linear uncertainty calculations that are performed in Gaussian space by PEST DSi support software. Both PEST and PESTPP-IES allow a modeller to supply measurement noise standard deviations separately from weights when implementing DSi-based data assimilation so that their separate roles can be respected.

Because a DSi statistical model runs so fast, other modes of usage are possible. Following calibration of a DSi statistical model using regularized inversion, confidence intervals of individual predictions can be explored using the predictive minimization/maximization process that is described in section 13.2. The many model runs that may be required by this process incur little numerical burden because of the trivial run time of the DSi statistical model.

14.3.3 Data Worth Analysis

Section 14.2.4 discusses DSi-based data worth analysis based on equations (14.2.7) and (14.2.8). There it is pointed out that a benefit that accompanies use of equation (14.2.8) is that the values of observations do not need to be known when testing their ability to reduce the uncertainty of a prediction of interest. The same does not apply when using the DSi statistical model that is embodied in equation (14.3.1). However, this equation is able to accommodate more complex interactions between observations of the past and predictions of the future than equation (14.2.8). Hence data worth analysis that is based on this equation may possess greater integrity.

Suppose that a modeller wishes to introduce a new observation into a history-matching dataset \mathbf{h} in order to test its ability to reduce the uncertainty of a prediction of interest. We will refer to this new observation as h_n (where “n” stands for “new”). Presumably, when building the surrogate DSi statistical model, the complex numerical model has been asked to calculate outputs that correspond to this observation; it will have calculated this output using many realizations of complex model hydraulic property fields. Hence many values of h_n are available.

Ideally, to obtain a robust assessment of the worth of h_n , the following steps can be taken.

1. In the manner described above, build a DSi surrogate model based on N realizations of complex model hydraulic properties. Include o_n , the model output that corresponds to h_n , in the $\begin{bmatrix} \mathbf{0} \\ \mathbf{s} \end{bmatrix}$ vector on which the DSi statistical model is based.
2. Make model predictions of interest \mathbf{s} , and quantify their uncertainties, by subjecting the DSi statistical model to Bayesian history-matching with the weight for h_n set to 0.
3. Create a value of h_n by adding a realization of measurement noise to one of the N values of o_n calculated by the complex model.
4. Subject the DSi statistical model to Bayesian history-matching again, this time with h_n included in the history-matching dataset.
5. Record alterations to posterior predictive mean values. Record reductions in posterior predictive uncertainties.
6. Repeat steps 3 to 5 for many different realizations of h_n .

It is apparent that this workflow is a lot more complicated than that described in the previous section, for it requires that the DSi statistical model be history-matched many times. However the outcomes of data worth analysis that is implemented in this way are likely to be more robust. This more complex data worth assessment workflow is also likely to provide greater

insights into the information content of h_n , as many different possible values of h_n are employed.

14.3.4 Strengths and Weaknesses of DSI

14.3.4.1 Strengths

DSI has many attractions.

Foremost among these attractions is the speed with which it can be implemented. The largest cost of any kind of simulator-based data assimilation is that which is incurred by repeated running of a numerical model. There is no model-based data-assimilation methodology that requires fewer model runs than data space inversion. Generally, training of the DSI statistical model requires only a few hundred numerical model runs. As implementation of DSI requires no adjustment of numerical model parameters, these runs do not need to be repeated during successive iterations of regularized inversion or ensemble-based history-matching.

While a DSI statistical model embodies its own approximations, the numerical model on which it is based does not require approximations that are needed for implementation of other history-matching methods that are described herein. Limits must be imposed on the execution speed (and hence complexity) of a numerical model that is subjected to regularized inversion or ensemble-based history-matching, as many model runs are required for implementation of these processes. In contrast, because the number of model runs required to train a DSI statistical model is generally commensurate with the number of model runs that is required for just one iteration of an IES history-matching process, model run times can be much longer. Model complexity can therefore be greater if this is deemed to be necessary for attainment of predictive integrity, and for representation of the full uncertainty range of model predictions.

As has also been discussed, parameter adjustment requires construction of a model parameterization scheme that, on the one hand, is complex enough to represent real-world patterns of hydraulic property heterogeneity, but on the other hand facilitates reassignment of values to parameters as the history-matching process progresses. This requires a level of abstraction that may erode the ability of a parameterization scheme to represent, and preserve, hydraulic property connectedness as parameters are adjusted. Difficulties may also be encountered in representing “uncertainty in prior uncertainty”, and in implementing adjustable representation of other aspects of a complex subsurface system whose details and historical stresses are only partly known. In contrast, the numerical model that is used for construction of a DSI statistical surrogate model can be endowed with categorical and discontinuous representations of hydraulic properties and other aspects of an environmental system that are imperfectly known. Meanwhile, statistical hyperparameters that are used for generation of these features can be altered from realization to realization in recognition of their own uncertainties.

A philosophically appealing feature of DSI is that the future is considered at the same time as the past when running a numerical model prior to construction of its DSI surrogate model. The suitability of the numerical model for the making of decision-pertinent predictions can therefore be assessed. A modeller can assure him/herself that the numerical model contains all features that are necessary for the making of these predictions, and that the model is capable of simulating the range of predictive behaviours that are compatible with current knowledge of a study site. This allows a modeller to optimize the design and capabilities of his/her numerical model as far as the making of predictions of management interest is concerned.

The prediction-specific nature of the DSI process is actually built into DSI statistical model

parameters.

Consider a complex numerical model that is used to train a DSI statistical model. Let us represent its hydraulic properties using the letter \mathbf{k} , even though these do not need to be adjustable. If we assume linearity, and also assume that prior parameter means are zero, then numerical model outputs are calculated from parameters \mathbf{k} using the equation

$$\begin{bmatrix} \mathbf{o} \\ \mathbf{s} \end{bmatrix} = \begin{bmatrix} \mathbf{Z} \\ \mathbf{Y}^t \end{bmatrix} \mathbf{k} \quad (14.3.9)$$

Meanwhile, DSI statistical model outputs are calculated from its iid parameters using the equation

$$\begin{bmatrix} \mathbf{o} \\ \mathbf{s} \end{bmatrix} = \mathbf{C}^{e^{-1/2}} \left(\begin{bmatrix} \mathbf{o} \\ \mathbf{s} \end{bmatrix} \right) \mathbf{x} \quad (14.3.10)$$

The “meaning” of statistical model \mathbf{x} parameters can be conceptualized by equating the right sides of the above two equations.

$$\mathbf{x} = \mathbf{C}^{e^{-1/2}} \left(\begin{bmatrix} \mathbf{o} \\ \mathbf{s} \end{bmatrix} \right) \begin{bmatrix} \mathbf{Z} \\ \mathbf{Y}^t \end{bmatrix} \mathbf{k} \quad (14.3.11)$$

It is apparent from equation (14.3.11) that the independent normal variates that comprise DSI statistical model parameters \mathbf{x} are ultimately derivable from the numerical model’s representation of real-world hydraulic properties embodied in the vector \mathbf{k} . However they are derived in a way that accounts for their effects on model outputs that are used in history-matching, and model outputs that comprise management-salient predictions. The whole purpose of decision-support modelling, which is to harvest information from the former and deliver it to the latter, is built into the very fabric of DSI statistical model \mathbf{x} parameters.

Another beneficial aspect of DSI is its convenience. If a modeller has prepared a dataset for IES-based history-matching, then he/she has also prepared a dataset for DSI-based history-matching. Both of these processes require an ensemble of parameter sets sampled from the prior probability distribution of numerical model parameters. If the numerical model is configured to make predictions, then this prior ensemble can be used to quantify the prior uncertainty of these predictions. If using IES, ensemble members are then adjusted so that the model can replicate the past, and therefore sample the posterior predictive probability distribution. As has been discussed in other chapters of this book, history-matching methodologies such as IES and ENSI which gain numerical efficiencies through ensemble-based dimensional reduction are vulnerable to parameter and predictive bias as parameters are adjusted. Bias can be exacerbated where the relationships between model parameters and model outputs are nonlinear. Use of DSI to make these same predictions, and quantify these same uncertainties at no extra numerical cost based on the same input dataset can serve as a check on the integrity of these methods as they are deployed at a particular study site.

14.3.4.2 Weaknesses

DSI is particularly susceptible to prior-data conflict. Suppose, for example, that model outputs \mathbf{o} calculated under historical conditions using a suite of hydraulic property realizations do not fully span all real-world observations encapsulated in a measurement dataset. Then some measurements that are represented in this dataset cannot be used to condition the \mathbf{o} -to- \mathbf{s} covariance matrix that comprises the DSI statistical model. Actually, prior-data conflict can be more subtle than this; see later in this chapter. Fortunately, the existence (or otherwise) of prior-data conflict is easily tested using the same dataset that is required for construction of a DSI

statistical model.

Another, more fundamental, problem with DSI is that relationships between the measured past and the predicted future that are highly nonlinear cannot be characterized using a covariance matrix. Examples of predictions that cannot be characterized in this manner are those that have natural upper or lower limits. These include contaminant and heat breakthrough curves. These types of system outcomes also pose challenges for normal score transformation as many predictions have the same value. Fortunately, problems incurred by natural predictive limits can be overcome using the methodology described by Doherty and Lotti (2025). DSI model training may then need to be prediction-specific. Training of a DSI statistical model to make a specific prediction can then be restricted to the range of values over which the prediction is able to vary.

Sometimes a DSI statistical model introduces temporal or spatial noise into some of its predictions. For example, a contaminant breakthrough curve predicted by a DSI statistical model may exhibit temporary non-monotonic behaviour. Predictions that are ill-informed by past measurements of system behaviour can be particularly susceptible to this effect. This kind of “predictive noise” can result from spurious correlations arising from use of a finite number of realizations for construction of a DSI statistical model. This problem must be seen in context, however. Noise of this type is generally considerably smaller than the uncertainties of predictions that they afflict. However the optics are not good.

Experience suggests that, to some extent, “predictive noise” of this type can be reduced by employing Tikhonov regularized inversion rather than MCMC or IES to history-match a DSI statistical model. When implementing Tikhonov regularization, the target measurement objective function should be set somewhat higher than the best fit that can actually be attained with a calibration dataset. Uncertainty intervals can then be quantified in the manner discussed above. That is, linear intervals can be calculated in Gaussian space, and then back-transformed to predictive space. Experience demonstrates that adoption of this strategy can often do much to “tidy up” DSI model outputs. This uncertainty analysis strategy is implemented by the PEST POSTDSIMOD utility program.

The last DSI problem that we mention here is not a numerical weakness. Rather it is an inherent aspect of the DSI methodology itself. It may be found that the breadth of a particular DSI-quantified uncertainty interval for a particular, management-salient prediction suggests the possibility of undesirable consequences accompanying a contemplated course of management action. However DSI cannot display the parameter field that precipitates these consequences; it can only demonstrate the possibility of their occurrence. Decision-makers may find this frustrating, or even decide that they cannot believe the possibility of management failure unless they “see it with their own eyes”.

In a situation such as this, a modeller may decide to embark on parameter-based data-assimilation. In order to maximize the efficiency of this process, he/she may include the value of the unwanted management outcome in a “prediction-specific measurement dataset”. He/she may calibrate the model so that it achieves this predictive outcome at the same time as it matches the historical behaviour of the system; see section 13.3. Use of a dimensional reduction method such as ENSI when implementing this process may reduce its numerical burden while allowing simultaneous adjustment of model parameters and geostatistical hyperparameters. The parameter field that emerges from this process will explain how the unwanted management outcome can occur.

14.4 DSIVC

14.4.1 Introduction

“DSIVC” stands for “data space inversion with variable control”. It uses data space inversion as a surrogate model when undertaking optimization under uncertainty.

Management optimization does not receive detailed treatment in this book. Nevertheless, it is discussed briefly in this section in order to give some context to an equally brief discussion of DSIVC. Readers should note, however, that the PEST++ and PyEMU suites provide a variety of workflows that implement optimization under uncertainty.

14.4.2 Optimization

14.4.2.1 Optimization and Decision Variables

The literature on optimization is vast. It is integral to management of any kind, particularly economic measurement. The term “optimization” implies minimization or maximization of some quantity. Less obvious, though equally important, is the notion of constraints that must be imposed on the quantity that is optimized.

In the broadest sense of the word, PEST performs optimization. After all, minimization of an objective function is a kind of optimization problem. In solving this problem, parameters are adjusted in order to reduce an objective function to its lowest value. Constraints on parameters may exist in the form of bounds whose enforcement maintains parameter value integrity. Ideally, however, the minimum value of an objective function is not adjacent to a parameter bound, for this implies that the value of at least one parameter is worryingly low or high, and that the objective function would fall even further if this constraint did not exist. This suggests that at least one parameter is playing a compensatory role for model imperfections.

The situation is different where optimization is undertaken to support management. In this case it is rare for the quantity that is being minimized/maximized (i.e. the objective function) to possess a “natural” optimum, such as that possessed by a sum of squared residuals. Instead, it is often cost that is minimized or profit that is maximized. Furthermore, calculation of these quantities may involve far more than simply summing squared residuals. Meanwhile, minimization of cost or maximization of profit is limited by constraints. Respect for these constraints becomes an essential element of the optimization algorithm.

An example may help.

Design of a groundwater remediation system may require removal of dissolved contaminant from under a certain site. This management goal may be pursued using a pump-and-treat system. Design of the remediation system requires that wells be emplaced and pumping rates be set. Use of fewer wells may require that more water be extracted from each well (and then disposed of). Under these circumstances, the ratio of contaminated to non-contaminated extracted water may be smaller than if water were extracted from many wells, as large local drawdowns may draw water from considerable distances. However installation costs may be lower than for a multiplicity of wells. Trade-offs must therefore be made. Optimization explores these trade-offs in an effort to minimize costs. The objective function that it minimizes must therefore include all contributors to cost, some of which are once-only (installation costs) and some of which are continuous (pumping costs).

From consideration of cost alone, the simplest solution to this optimization problem is not to pump at all in order to minimize cost. But this does not enable cleanup. That is, it does not

satisfy constraints that are an integral part of the optimization problem. In this particular example, the principal constraint may be that predicted contaminant levels in specific monitoring wells be zero, or below a certain threshold, after a certain amount of time has elapsed. Other constraints may pertain to drawdowns that are experienced at impacted off-site wells; perhaps these must not exceed certain regulatory thresholds.

Items that are varied in order to minimize or maximize a management objective function are referred to as “decision variables”. Well coordinates and pumping rates are decision variables in the present example. The optimization process must discover a set of decision variables for which the objective function is minimized subject to respect for all constraints. Usually, the objective function minimum abuts one or more of these constraints; this is how its value is minimized. For example, regardless of how pumping is distributed between wells, and regardless of the number and locations of these wells, pumping cannot be reduced below a level at which the cleanup constraint is violated.

Because the unconstrained minimum of a management objective function generally falls outside constraint surfaces, management optimization is very different from optimization that is undertaken by PEST. A linear version of this problem is schematized in figure 14.2.

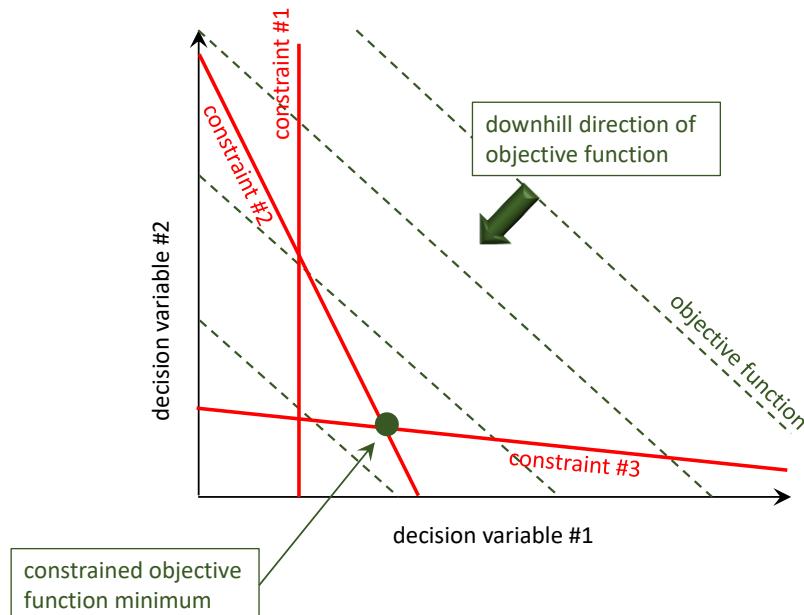


Figure 14.2. A linear version of a constrained minimization problem.

Problems such as that which is depicted in figure 14.2 can be solved using so-called “linear programming” algorithms. Where objective function contours and constraint surfaces are curved, the optimization problem must be solved iteratively based on a succession of problem linearizations; algorithms such as sequential linear programming are used for this task. In both cases, model runs are undertaken to calculate derivatives of model outputs with respect to decision variables. Outputs that are calculated by a numerical simulator generally pertain to constraints; costs are generally calculated using times and materials that do not feature in the environmental simulation process.

Algorithms such as sequential linear program can handle large numbers of decision-variables with ease. More difficult optimization problems involving highly nonlinear objective functions (possibly with multiple optima) and irregular constraint surfaces may require the use of so-called “global optimizers”. Algorithms that support global optimization are generally much

slower and require many more iterations, and hence many more model runs.

The principle numerical cost incurred by management optimization is the same as that incurred by calibration and uncertainty analysis. It is the cost of repeated running of a numerical simulator. As stated above, generally the simulator must be run in order to calculate quantities that have the potential to violate constraints. Calculation of the objective function is generally a relatively simple matter.

14.4.2.2 Optimization under Uncertainty

Numbers calculated by environmental models (particularly groundwater models) are uncertain for reasons that have been discussed extensively in this book. Where a model is used to test whether a particular set of decision variables violates a set of management constraints, this assessment can only therefore be probabilistic. Constraints themselves may be expressed in probabilistic terms. For example, it may be specified that “drawdown at point A must not exceed 2 m at a 95% level of confidence” or that “the contaminant concentration at well B must be less than 1 µg/L at a 98% confidence level”.

In a linear world, the optimization-under-uncertainty problem of figure 14.2 must now be cast as in figure 14.3. In this figure, constraint surfaces are surrounded by their uncertainty intervals. The width of each interval depends on the level of confidence with which the interval is associated. Two solutions to the optimization-under-uncertainty problem are shown in this figure. The first solution expresses risk aversion while the second solution expresses tolerance of risk. Normally risk reduction is the priority. In this schematic, the risk-averse minimized cost is higher than the minimized cost in figure 14.2 as it acknowledges uncertainties in calculation of whether (or not) management constraints are respected.

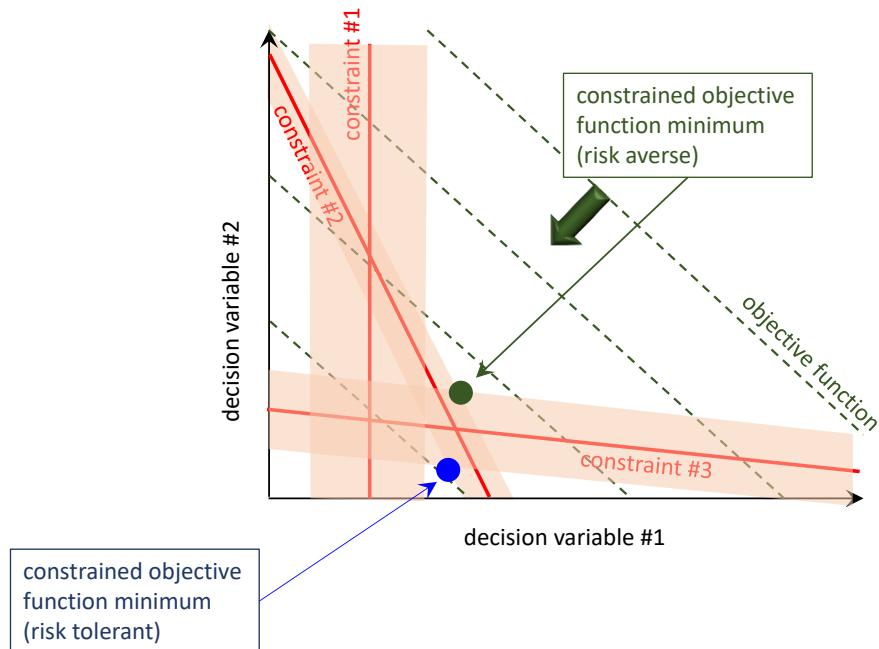


Figure 14.3. Optimization under uncertainty. The green dot is the risk-averse solution to the optimization-under-uncertainty problem while the blue dot is the risk tolerant solution. The solution shown in Figure 14.2 is the risk neutral solution.

Obviously, solution of a constrained optimization-under-uncertainty problem requires that uncertainties associated with constraint-pertinent model outputs be quantified. This can

significantly increase the numerical burden of management optimization. However if uncertainties are independent (or nearly independent) of the values of decision variables (as is depicted in Figure 14.3), then they need to be calculated for only one set of decision variables. The numerical burden of optimization-under-uncertainty is thereby somewhat relieved.

In general, however, model predictive uncertainties are nonlinear functions of decision variables. This means that whenever a numerical model is run in order to test whether use of a set of decision variables that is selected by an optimization package violates constraints at a certain level of confidence, then model output uncertainties must be re-evaluated. The “model” that the optimizer must run therefore becomes an ensemble, or “stack”, of models. All members of this ensemble employ the same set of decision variables; however they use different parameter fields sampled from the posterior parameter probability distribution to calculate whether, or not, constraints are violated. These model runs produce an ensemble of model outputs. The output that is used in the optimization process is that which corresponds to the level of confidence with which it is decreed that a particular constraint must be respected. In extremely conservative settings, this may be the most pessimistic model output as far as respect for the constraint is concerned.

The numerical burden of optimization under uncertainty becomes greater with the level of model nonlinearity. As stated above, where an optimization problem is highly nonlinear, global optimizers must be used. The model run cost of using a global optimizer is high, even where uncertainty is not taken into account. Where each “model run” becomes an ensemble of model runs, the numerical cost becomes prohibitive.

Strategies through which the number of model runs required for probabilistic assessment of respect for constraints can be reduced have been published in the literature. In risk averse decision contexts, it may be possible to reduce the size of the parameter stack to those parameter realizations that engender pessimistic predictions pertaining to specific constraints. Even where strategies such as these are adopted, however, the model run cost of solving an optimization-under-uncertainty problem may remain unduly high. The only viable option for solution of such a problem may therefore be to employ a surrogate model in place of the actual numerical model when testing constraints. Optimization strategies provided by the PEST++ suite, supported by the PyEMU suite, provide these options. Here we discuss only one option, this being DSIVC.

14.4.3 How DSIVC Works

DSIVC is explained in detail by Jiang et al (2020). It is briefly explained below.

The description of DSIVC provided earlier in this section shows how a complex numerical simulator of arbitrary structural and parametric complexity can be replaced by a simple statistical model whose parameters reside in a vector \mathbf{x} . Prior to conditioning by historical measurements of system behaviour, each of these \mathbf{x} parameters is an iid; that is, it is an identically and independently distributed standard normal variate. Ensemble-based or regularized conditioning of the DSIVC statistical model takes only a few seconds. Conditioning provides samples of the posterior probability distribution of parameter \mathbf{x} and samples of the posterior probability distributions of predictions of interest. Alternatively, depending on how it is implemented, conditioning may provide posterior mean predictions together with predictive confidence limits.

The attractiveness of DSIVC rests on its ability to provide a direct statistical linkage between the measured past and the managed future. This linkage can accommodate parametric and

structural uncertainty. What it does not accommodate, however, are changes to future stresses such as pumping rates. The making of new predictions requires construction of a new statistical model. Unfortunately, however, future system stresses must undergo continual alteration as they are optimized.

Jiang et al (2020) show how the DSI process can be modified to accommodate alterations to future system stresses. These modifications must be built into the DSI model construction and history-matching process. The algorithm is now briefly described. Refer to the original paper for more details.

Suppose that it is intended that a DSI surrogate model be used in place of a complex numerical simulator to solve a management optimization problem. For the moment we disregard predictive uncertainty. Hence a risk-neutral solution is sought to the optimization problem.

Before the optimization process can be initiated, a suitable DSI model must be trained. So the numerical model is run over the past and into the future N different times, where N may range from a few hundred to a few thousand. On each occasion that it is run, the numerical simulator employs a realization of an arbitrarily complex hydraulic property field. *It also employs a realization of future stresses.* For the groundwater remediation example that is discussed above, this realization of future stresses embodies random pumping rates at wells that may (or may not) ultimately be used in an optimized pump-and-treat remediation system. On each occasion that it is run, model outputs of interest are those that correspond to historical field measurements of system behaviour, and those that pertain to constraints that must be respected as remediation takes place.

Once these N runs of the numerically expensive simulator have been undertaken, the optimization process can begin. An optimization package may be asked to minimize remediation cost subject to the meeting of constraints pertaining to cleanup and environmental impact such as those discussed above. During the optimization process, the optimization package commissions many model runs. On each occasion that a model run is commissioned, it is actually the DSI surrogate model that is run. However the “running” of this model requires its recalibration, not just to historical measurements of system behaviour, but to pumping rates that the optimizer requests. The calibrated DSI statistical model then makes minimum error variance (i.e. posterior mean) predictions of contaminant concentration. Repeated recalibration of the statistical model may sound expensive; however each recalibration exercise can normally be accomplished in just a few seconds. The “run time” of the DSI statistical model recalibration process (which is also the DSI model prediction process) is generally much shorter than the run time of a numerical simulator.

Where uncertainty is sought in model predictions, the process becomes only slightly more complicated. Two options are possible. Instead of recalibrating the DSI statistical model on each occasion that a new set of pumping rates is tested, the DSI statistical model can be subjected to ensemble-based history-matching. The history-matching dataset is the same as before, namely historical measurements of system behaviour and the set of pumping rates that are requested by the optimizer. Constraint-pertinent predictions are now made using an ensemble of DSI statistical model \mathbf{x} parameters. Model outputs of interest, as far as the optimizer is concerned, are those that pertain to the confidence level with which remediation constraints must be respected. For example, in an extremely risk-averse optimization procedure, the most pessimistic model output pertaining to each constraint is used by the optimizer.

This procedure can be made slightly more efficient using functionality that has already been

discussed that is available through PEST's implementation of the DSI model construction and deployment process. This requires calibration rather than ensemble-based history-matching of the DSI statistical model. The POSTDSIMOD utility yields DSI model posterior mean predictions together with associated predictive confidence intervals. Model outputs of interest as far as the optimizer is concerned are the values of constraint-salient predictions pertaining to confidence limits that are adopted by the optimization-under-uncertainty process.

As stated above, the cost of DSI model recalibration and uncertainty interval assessment is only a few seconds. Over the course of an optimization-under-uncertainty process, this “model” may need to be run thousands of times. However the procedure can be supervised by different agents in a parallelized computing environment. Meanwhile, its numerical cost is far smaller than that required for repeated running of a numerical simulator in an optimization-under-uncertainty process in which each “model run” requires that numerical simulation be undertaken for an ensemble of parameter fields that sample the posterior parameter probability distribution.

14.5 Prior-Data Conflict

14.5.1 Background

Ensembles form the basis for two methodologies that have received a considerable amount of attention in this book; these are the iterative ensemble smoother (IES) and data space inversion (DSI). At the commencement of both of these processes, a numerical simulator is run many times using different realizations of hydraulic conditions (properties, boundary conditions and possibly past system stresses) that are sampled from a prior probability distribution. This distribution is designed by a modeller to embody that which is known about a groundwater system (for example, the range of realistic values that hydraulic properties may take and the degree of connectedness of these properties) and that which is unknown (for example, the actual values of hydraulic properties).

Once model runs have been undertaken using these realizations, a modeller is provided with the opportunity to subject his/her formulation of the prior probability distribution of hydraulic properties to a reality check. If, for example, the range of modelled heads at a particular site on a particular date that are calculated using all hydraulic property realizations does not span the head that was actually observed at that site on that date, then something is wrong. The suite of ensemble-based model runs has revealed that the model cannot replicate this particular aspect of system behaviour. Perhaps the model must be reconfigured. Alternatively, perhaps the hydraulic properties with which the model was supplied have hindered its ability to replicate the measured head. Perhaps the prior probability distribution of these properties must be modified in order to encompass a greater range of variability, and/or to express a greater or smaller degree of connectedness in some or many parts of the model domain.

Whatever the reason for the problem, it requires rectification before ensemble-based history-matching, or construction of a DSI statistical model is undertaken. An erroneous prior will lead to an erroneous posterior. Furthermore, both of the IES and DSI history-matching processes may imbue parameters with bias if they must adopt roles that compensate for inappropriately-expressed expert knowledge in order for model outputs to match field observations. Where realizations of model outputs do not span observations, then severe difficulties may be encountered in DSI history-matching where normal score transformation of model outputs precedes DSI model construction. Implied in this procedure is that observations of system behaviour are “covered” by the ensemble of complementary model outputs.

The situation that has just been described is referred to as “prior-data conflict”. Much can be

learned about the integrity of concepts on which model construction rests by maintaining a careful watch for its occurrence, and then taking remedial action where necessary. A benefit that accompanies use of ensemble methods is that its detection is readily enabled by their use. DSI support programs provided with the PEST suite, as well as the PESTPP-IES ensemble smoother, report incidences of field measurements being unspanned by ensembles of complementary model outputs.

Unfortunately, however, prior-data conflict may express itself in ways that are more subtle than this. Ensuring that system behaviour is compatible with the prior probability distribution of hydraulic properties requires more than simply ensuring that no measurement of system state lies outside the band of complementary model-calculated system states where these states are calculated using multiple samples of the prior. The model must also be capable of replicating temporal and spatial patterns that are exhibited by field measurements. This comparison is more difficult to make. Nevertheless, it is far from impossible.

The PEST suite includes a program named PDC1. This program analyses prior-data conflict using a methodology described by Alfonzo and Oliver (2019). The dataset required for prior-data conflict analysis is exactly the same as that required for ensemble-based history-matching or data space inversion. Hence it can be undertaken as a complementary activity to either of these data assimilation methodologies, thereby ensuring that the preconditions for their successful implementation are in place.

14.5.2 The Alfonzo and Oliver Algorithm

14.5.2.1 Mahalanobis Distance

Central to the Alfonzo and Oliver algorithm is the concept of Mahalanobis distance. This is described in section 3.14. The contents of that section are now briefly reviewed.

Let \mathbf{x} be an n -dimensional random vector whose statistical properties are described by its mean $\underline{\mathbf{x}}$ and covariance matrix $C(\mathbf{x})$. The squared Mahalanobis distance between an arbitrary vector \mathbf{z} and the mean of \mathbf{x} is calculated as

$$d_m^2 = D_m = (\mathbf{z} - \underline{\mathbf{x}})^t C^{-1}(\mathbf{z} - \underline{\mathbf{x}}) \quad (14.5.1)$$

It is easily shown that if $\mathbf{z} \sim N(\underline{\mathbf{x}}, C(\mathbf{x}))$ (i.e. if \mathbf{z} shares the same probability distribution as \mathbf{x} , and if that probability distribution is multiGaussian) then D_m possesses a chi squared distribution with the same number of degrees of freedom as the number of elements of \mathbf{x} (and hence \mathbf{z}), i.e. n . Recall from section 3.2.3 that a chi squared distribution has a mean of n and a standard deviation of $\sqrt{2n}$.

Mahalanobis distance can provide a basis for testing whether a particular vector \mathbf{z} is a sample of the same probability distribution as that to which \mathbf{x} belongs. If \mathbf{z} and \mathbf{x} have spatial/temporal connotations (and hence may display spatial/temporal correlation), then inclusion of $C(\mathbf{x})$ in calculation of the statistical distance between \mathbf{z} from the mean of \mathbf{x} takes their spatial/temporal characteristics into account.

14.5.2.2 Measuring Prior-Data Conflict

When analysing prior-data conflict, the vector \mathbf{z} of equation (14.5.1) encapsulates a dataset comprised of field measurements of system behaviour, while \mathbf{x} represents measurement-noise-contaminated, model-generated counterparts of these measurements. We denote the former using the symbol \mathbf{h}_m and the latter using the symbol \mathbf{h} .

Before the test embodied in equation (14.5.1) can be applied in practice, $\underline{\mathbf{h}}$ and $\mathbf{C}(\mathbf{h})$ must be determined. This is done empirically using stochastic realizations of \mathbf{h} . Realizations of \mathbf{h} are generated by running a model using random realizations of hydraulic properties, just as is done prior to implementation of DSI. Each \mathbf{h} that is generated in this way is then sprinkled with a random realization of measurement noise.

If the number of elements of \mathbf{h} (i.e. n) exceeds the number of realizations that are used for empirical evaluation of $\mathbf{C}(\mathbf{h})$ (we will denote the latter as N), we now have a problem. In environmental modelling practice, this problem is not uncommon, as the number of observations that comprises a history-matching dataset often exceeds the number of realizations that are used in ensemble-based history-matching, or in construction of a DSI surrogate statistical model. Where N is less than n , an empirical $\mathbf{C}^e(\mathbf{h})$ matrix that is built from noise-enhanced model outputs that correspond to field measurements is rank-deficient and therefore noninvertible. Hence another matrix must be inverted in its place. Ideally this matrix is “close to” $\mathbf{C}^e(\mathbf{h})$ but is not singular. Alfonzo and Oliver (2019) describe such a matrix (see below). They claim that this matrix provides superior performance to an approximate inverse matrix obtained by first subjecting $\mathbf{C}^e(\mathbf{h})$ to singular value decomposition and then taking the reciprocal of non-zero singular values.

The Alfonzo and Oliver methodology for analysis of prior-data conflict proceeds as follows.

First, run the numerical model N times based on N random realizations of model hydraulic properties sampled from the prior. This produces N realizations of model outputs \mathbf{o} . Next, add realizations of measurement noise to these realizations of \mathbf{o} to generate realizations of \mathbf{h} . Collect these realizations of \mathbf{h} into an $n \times N$ matrix \mathbf{H} so that the vectors \mathbf{h} comprise the columns of \mathbf{H} .

Next a set of “leave-one-out” analyses is repeated N times. During each of these procedures, one member of the \mathbf{h} ensemble (i.e. one column of \mathbf{H}) is selected for individual attention. This becomes the “left-out” realization of \mathbf{h} . Denote this realization using the index i . The squared Mahalanobis distance D_{mhi} between this realization and the mean of the remaining realizations is then calculated. Then the squared Mahalanobis distance D_{mhmi} between the set of field measurements \mathbf{h}_m and this same mean is calculated.

Specifically:

1. Remove column i from \mathbf{H} .
2. Calculate $\underline{\mathbf{h}}$ and $\mathbf{C}^e(\mathbf{h})$ using the remaining $N-1$ realizations.
3. Using equation (14.5.1), calculate the squared Mahalanobis distance between the omitted \mathbf{h} (i.e. \mathbf{h}_i) and the empirical mean of the included realizations. Denote this as D_{mhi} .
4. Using equation (14.5.1) calculate the squared Mahalanobis distance between the measurement dataset \mathbf{h}_m and the empirical mean of the included realizations. Denote this as D_{mhmi} .

Calculation of $\mathbf{C}^{e-1}(\mathbf{h})$ (i.e. the approximate inverse of $\mathbf{C}^e(\mathbf{h})$) for use in equation (14.5.1) is described by Alfonzo and Oliver. The PEST utility PDC1 uses the same procedure, but with a slight variation. Before inversion, the matrix $\mathbf{C}^e(\mathbf{h})$ is replaced by a matrix Σ that is defined as follows.

$$\Sigma = \delta \mathbf{T} + (1 - \delta) \mathbf{C}^e(\mathbf{h}) \quad (14.5.2)$$

δ is a “shrinkage parameter” for which Alfonzo and Oliver employ

$$\delta = \frac{2}{N+2} \quad (14.5.3)$$

The matrix \mathbf{T} in equation (14.5.2) is formed from the identity matrix by dividing the diagonal elements of \mathbf{I} (which are all equal to unity) by the average variance of observations pertaining to the observation group to which respective elements of \mathbf{I} belong. (Alfonzo and Oliver use a single variance that is averaged over all observations.) Each element of \mathbf{T} is thus the reciprocal of a variance v that is specific to each observation group. For each observation, a variance is calculated over all realizations of \mathbf{h} comprising \mathbf{H} (except for the omitted \mathbf{h}_i), and then averaged over all members of the observation group to which the observation belongs. This allows \mathbf{T} to reflect the fact that observations that comprise a history-matching dataset may be of very different types.

Meanwhile, during each cycle of the above procedure, $C^e(\mathbf{h})$ is evaluated using the following equation (with \mathbf{h}_i omitted from \mathbf{H}):

$$C^e(\mathbf{h}) = \frac{(\mathbf{H} - \underline{\mathbf{H}})(\mathbf{H} - \underline{\mathbf{H}})^t}{N-2} \quad (14.5.4)$$

In equation (14.5.4) the columns of $\underline{\mathbf{H}}$ are all the same. Each is a vector $\underline{\mathbf{h}}$ that is the mean of all realizations of \mathbf{h} except for the omitted \mathbf{h}_i .

The outcomes of this sequential, leave-one-out procedure are a set of N values of D_{mhi} and corresponding values of D_{mhmi} . Ideally, values of D_{mhi} obtained in this way provide a basis for deciding whether values for D_{mhmi} obtained in this way are different enough to warrant the measurement dataset \mathbf{h}_m being declared as statistically different from model-generated datasets \mathbf{h} contained in \mathbf{H} . Different statistical characterization of D_{mhi} and D_{mhmi} implies prior-data conflict.

Following this succession of leave-one-out cycles, the median of the N values of D_{mhi} is calculated. This is compared with the median of the N values of D_{mhmi} . (Note that diversity of D_{mhmi} values calculated over the N leave-one-out iterations is generally much smaller than that of D_{mhi} values. This is because there is only one measurement dataset \mathbf{h}_m and many incidences of \mathbf{h}_i ; the latter may exhibit considerable variability.) A robust approximation to the standard deviation of these $N D_{mhi}$ values is then evaluated; this is calculated as 1.4825 times the mean absolute deviation (MAD) from the median of these values (a factor which is appropriate for normally distributed random variables).

Next a z -score is awarded to the median of all D_{mhmi} values using the mean and standard deviation calculated above. A z -score for a general random variable r is calculated from its mean μ and standard deviation σ as:

$$z = \frac{r - \mu}{\sigma} \quad (14.5.5)$$

A z -score characterizes the number of standard deviations by which a Gaussian random variable is removed from its mean. Obviously, high absolute values are unlikely values. For a Gaussian distribution, the interval that is spanned by 2 standard deviations either side of the mean is the 95% confidence interval. A z -score whose absolute value is greater than 2 therefore indicates an unlikely value for a random variable; a z -score whose absolute value is greater than 3 indicates a very unlikely value for a random variable. For this reason, a z -score of 3 or higher for median D_{mhmi} , calculated in the above manner, is indicative of prior-data conflict.

Alfonzo and Oliver suggest another metric for appraisal of prior-data conflict. They count how many D_{mhmi} values fall within the empirical 95% confidence interval of D_{mhi} values; both of these D_m values are accumulated over the N cycles of the leave-one-out procedure described above. This count is divided by N to yield a fractional count. The closer is this fractional count to 1.0, the less likely is prior-data conflict. Plotting of the empirical D_{mhi} and D_{mhmi} cumulative

probability distributions can also provide visual clues to a modeller who is trying to establish the existence, or otherwise, of prior-data conflict. See Alfonzo and Oliver (2019) for more details.

15. Model Defects

15.1 Introduction

No model is a perfect simulator of environmental processes at any study site. While this does not invalidate the use of models in environmental decision-making, it does mean that they should be used with caution. It also means that modellers should be aware of the repercussions of model defects so that, when called upon to make the many subjective decisions that modelling and model history-matching entails, these decisions can be as informed as possible.

Model defects arise from many sources. These include approximations used in the model algorithm, failure to provide enough parameters to represent system property heterogeneity, erroneous definition of temporal and spatial boundary conditions, the need for spatial and temporal discretization that supports numerical representation of partial differential equations, and from improper definition of system stresses and source terms. Defects may become visible during history-matching where it is discovered that model outputs cannot replicate measurements of system state. Alternatively, model defects may remain invisible during history-matching, especially in highly parameterized contexts where parameters have ample opportunity to compensate for model defects, so that high levels of model-to-measurement fit are attained for the wrong reason. As will be shown below, for some predictions this may be a good thing as model defects are thereby “calibrated out”. For other predictions made by the same model it will be a bad thing as history-matching induces a potential for error that was previously absent.

Once model defects are taken into account, assessment of the potential for error in history-matched model parameters, and in predictions made by a history-matched model becomes complicated. It also becomes subjective, because the nature and extent of model defects is normally unknown. Ideally, quantified predictive uncertainty intervals should include possible errors that arise from model structural or parametric defects. Or, better still, predictive errors arising from model defects should be small compared to the uncertainties of predictions of management interest. However this cannot always be guaranteed. Nevertheless, as will be discussed, steps can be taken to reduce the deleterious effects of model defects on post-history-matching model predictions.

15.2 An Example

Figure 15.1 shows a simple one-dimensional groundwater model. Water flows into the model domain at a known rate through its right boundary. The groundwater head is fixed at the left side of the model domain. Transmissivity within the model domain is uniform. Four observation wells are situated within the domain of this simple model. Heads are measured in these wells; these head measurements are accompanied by measurement noise.

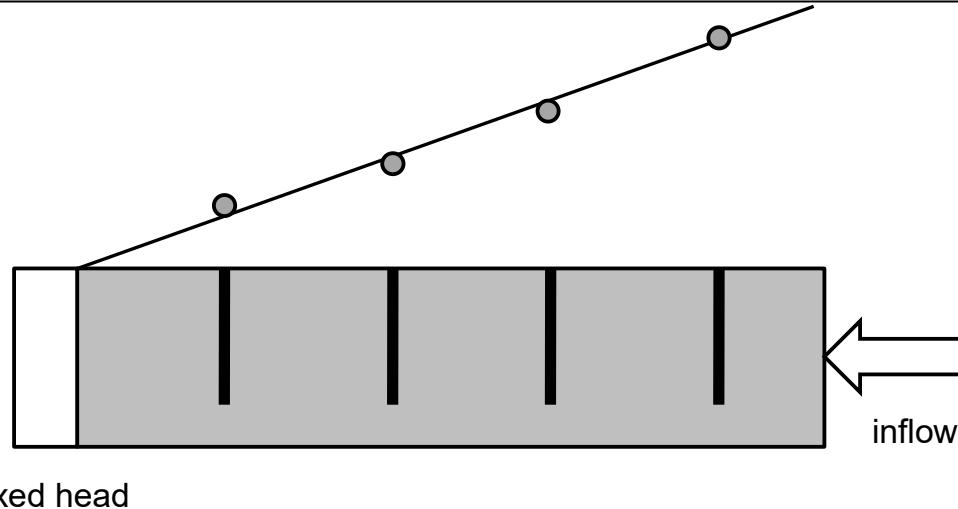


Figure 15.1. A simple groundwater model. The four circles show water levels recorded in four wells. The solid line shows heads throughout the model domain calculated using a uniform, calibrated transmissivity value. This and subsequent figures are modified from those presented in Doherty and Welter (2010).

Suppose that the model of figure 15.1 is calibrated against the four borehole water level measurements depicted in that figure in order to estimate the single parameter that characterizes domain-wide transmissivity. The solid line in figure 15.1 shows heads calculated by the model using the single estimated transmissivity value. The fit between model-calculated and measured heads is good. Measured heads are scattered randomly about this line of best fit; the amount of this scatter is commensurate with measurement noise.

A defect is introduced to the model in figure 15.2. The head at the left side of the model domain is assigned an incorrect value. The model is then re-calibrated to estimate the uniform domain transmissivity. The straight line of model-calculated heads depicted in figure 15.2 provides the best fit between model-calculated and measured heads in the least squares sense. However in this case the scatter of measured heads about the line of model-calculated heads appears to show some “structure”; uphill heads are over-estimated and downhill heads are underestimated by the history-matched model.

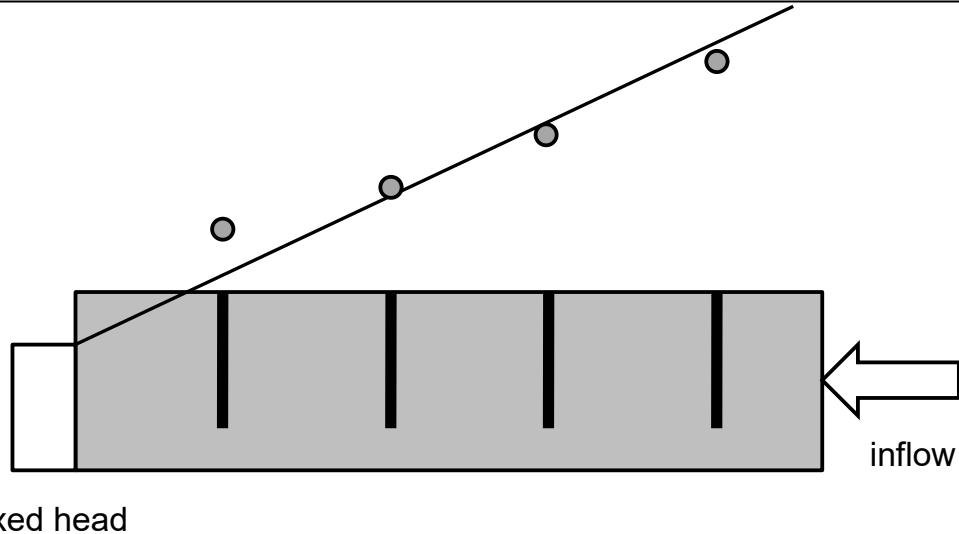


Figure 15.2. The same model domain as that depicted in figure 15.1, but with an incorrect boundary condition. The solid line shows heads throughout the one-dimensional model domain calculated by the calibrated model.

In this simple case, a modeller may recognize that history-matching has exposed a model defect. He/she may even be able to identify the defect. In more complex modelling contexts, recognition of a misfit-inducing defect may not be possible; or perhaps a defect may not be fixable. So the modeller may decide that he/she can live with the defect, but will ensure that it is not ignored when quantifying predictive uncertainties. In recognition that a good fit at one location or time must be traded against a good fit at another place and time, a modeller may decide to formalize this trade-off using a Pareto method; see, for example, Vrugt et al (2003) and Wöhling and Vrugt (2008).

Figure 15.3 implements this strategy. Two lines are shown in this figure. These lines pass through the extremes of the observation dataset. One of these lines has the maximum slope that it is possible for a line to possess and still respect any of the borehole head measurements. The other line is its minimum slope complement. It is apparent that estimation of the true transmissivity value requires a smaller slope than either of these lines. The true value of transmissivity is therefore not encompassed by parameter uncertainty limits that are assessed in this way.

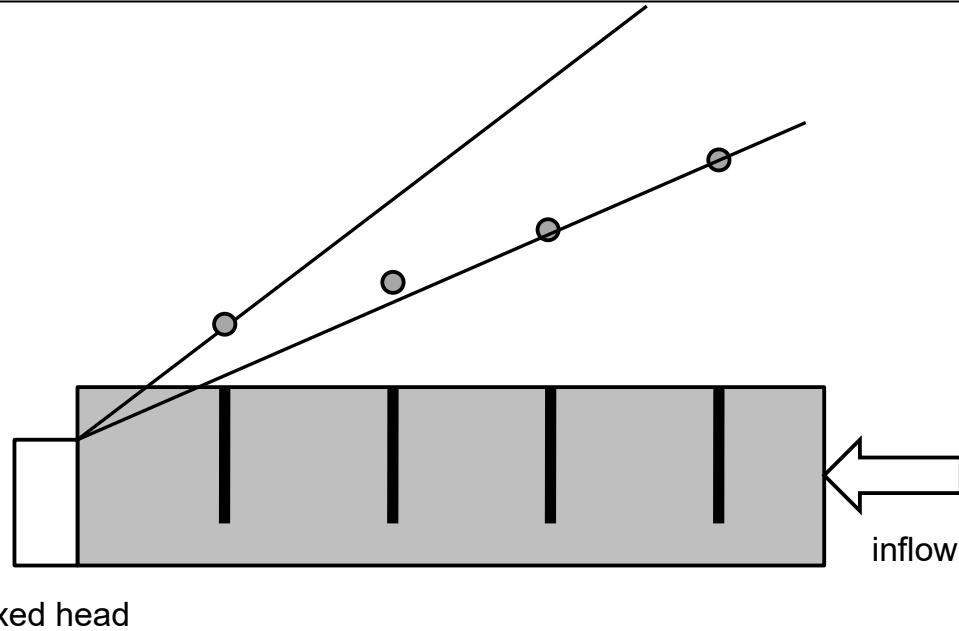


Figure 15.3. An attempt is made to bracket the true transmissivity by fitting data extremes.

The “best fit” though the four head measurements comprising the calibration dataset is that depicted in figure 15.4. The line of model-generated heads passes through none of the measured heads. However the slope of this line is correct; hence estimated transmissivity is correct.

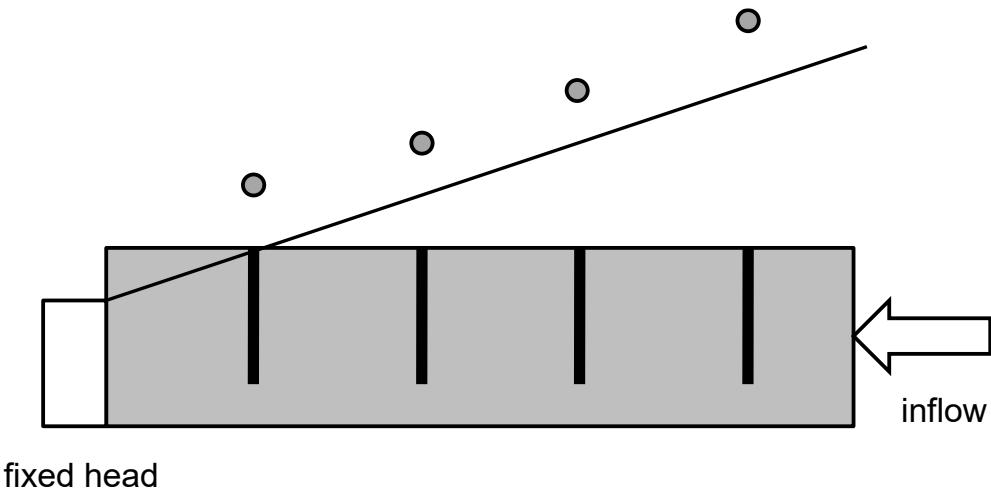


Figure 15.4. Line of “best fit” through the calibration dataset.

So would a modeller pursue a fit like that depicted in figure 15.4 when calibrating his/her model? Perhaps so. If the history-matching process is reformulated so that the model is asked to fit three measured lateral head differences rather than four measured heads, then exactly this fit is obtained and the value estimated for transmissivity is correct. There is a reason for this. The particular defect that afflicts this model does not affect computed lateral head differences; it only affects computed heads. We will return to this topic later. Meanwhile, history-matching using lateral head differences is schematized in figure 15.5.

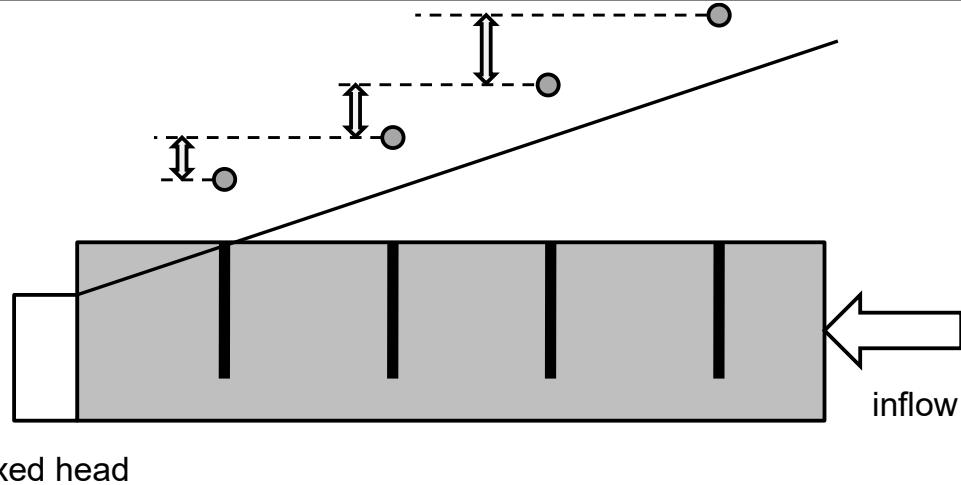


Figure 15.5. History-matching is achieved through matching lateral differences in measured heads rather than heads directly.

Rather than sensing a model defect, a modeller may attribute failure to connect all borehole heads to the left model boundary with a single straight line as evidence of the presence of transmissivity heterogeneity. He/she may then introduce parameters that can represent that heterogeneity. Figure 15.6 shows history-matching outcomes where two transmissivity parameters are employed. A good fit is obtained between model outputs and field measurements while residuals are random and independent. The T_2 parameter has “absorbed” structural noise and allowed correct estimation of the T_1 parameter.

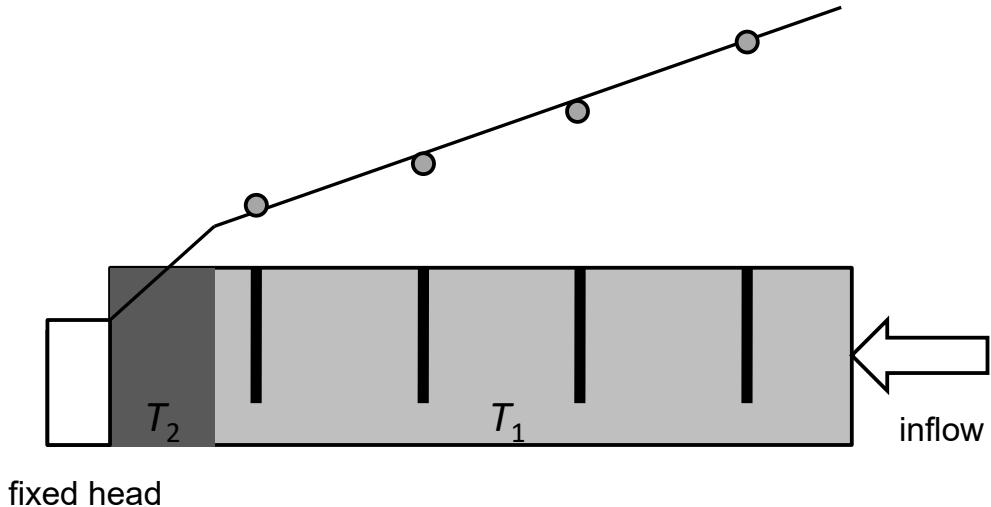


Figure 15.6. Estimation of an extra transmissivity parameter allows attainment of a good fit with borehole head data and estimation of a correct value for T_1 .

So how “valid” is the model of figure 15.6? It obviously has some validity as most of the model domain has been awarded the correct value of transmissivity. If the model is asked to predict drawdowns in and near a pumping well that is installed to the right of the T_2 boundary, its predictions will be correct. However if the model is asked to predict heads throughout the model domain following an increase in inflow from the right, then its predictions will be wrong; see figure 15.7. In fact its predictions will be more wrong than those made by the calibrated model of figure 15.2 over a substantial part of the model domain.

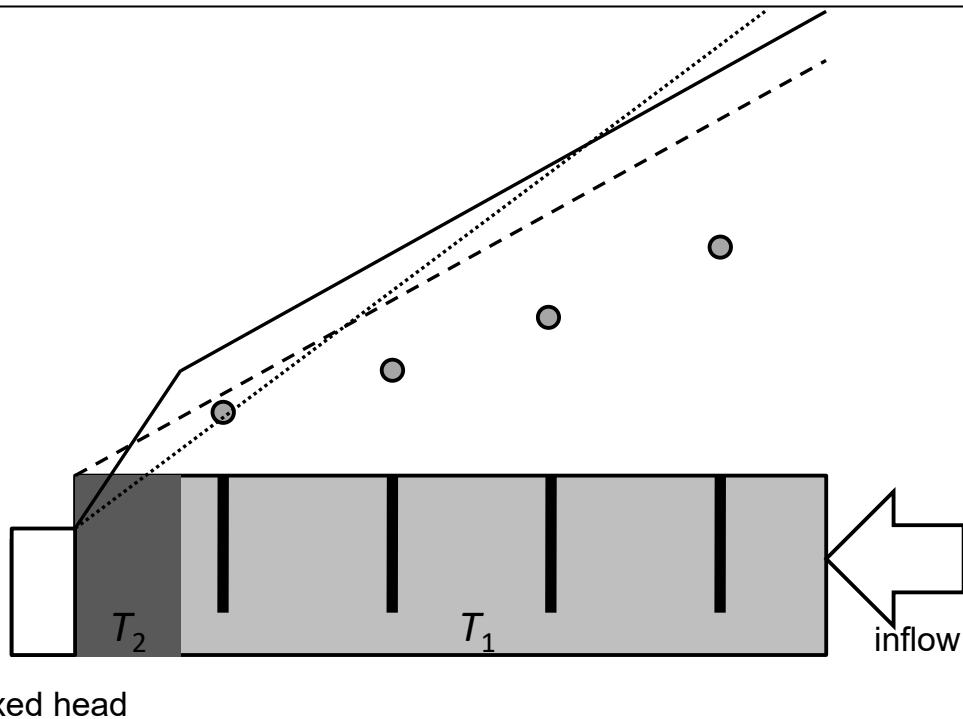


Figure 15.7. The full line shows heads predicted by the calibrated defective model of figure 15.6 following increased uphill inflow. The dashed line shows the correct heads under these conditions. The dotted line shows predictions calculated by the calibrated model depicted in figure 15.2 which employs no compensatory parameters.

During history-matching of the defective model of figure 15.6, the T_2 parameter plays a role that compensates for its defect. Its playing of this same compensatory role ensures the integrity of some model predictions, for example predictions of drawdown following pumping introduced to the right of the model domain. For another prediction however (namely the response of the system to increased lateral inflow) this compensatory role induces error.

15.3 Mathematical Formulation of the Problem

15.3.1 A Defective Model

Up until this point, the linearized action of a model when undergoing history-matching has been repeatedly described using the equation

$$\mathbf{h} = \mathbf{Zk} + \boldsymbol{\varepsilon} \quad (15.3.1)$$

Implied in all previous analyses that have been based on this equation is the assumption that the numerical simulator that is represented by \mathbf{Z} provides perfect emulation of system behaviour. Things are now about to change.

To simplify the equations that follow we assume the following.

$$\mathbf{C}(\boldsymbol{\varepsilon}) = \sigma^2_{\varepsilon} \mathbf{I} \quad (15.3.2)$$

$$\mathbf{C}(\mathbf{k}) = \sigma^2_{\mathbf{k}} \mathbf{I} \quad (15.3.3)$$

The former is effectively achieved through use of an appropriate weight matrix. Meanwhile, the latter assumption does not detract from the generality of conclusions to follow as it can be achieved through appropriate parameter transformation (e.g. Kahanen-Loève transformation).

Defects will now be introduced to the model so that equation (15.3.1) becomes

$$\mathbf{h} = \mathbf{Z}_m \mathbf{k}_m + \mathbf{Z}_d \mathbf{k}_d + \boldsymbol{\epsilon} \quad (15.3.4)$$

\mathbf{Z}_m and \mathbf{k}_m of equation (15.3.4) pertain to the model as it is actually “seen” by the modeller. Hence equation (15.3.3) is re-written as

$$\mathbf{C}(\mathbf{k}_m) = \sigma^2_{km} \mathbf{I} \quad (15.3.5)$$

The \mathbf{k}_d vector of equation (15.3.5) is comprised of “correction parameters”, while the \mathbf{Z}_d matrix is comprised of “correction processes” that act on these parameters. Of course representation of model defects in this way is linear; however, as in previous sections of this document, the linearity assumption allows us to draw some useful conclusions. These conclusions are readily extrapolated to nonlinear models.

The $\mathbf{Z}_d \mathbf{k}_d$ term of equation (15.3.4) embodies a (normally unknown) “correction term” for model outputs. The (normally unknown) \mathbf{k}_d parameters can be conceived of as expressing the reasons for which corrections are necessary. Take, for example, the model of figure 15.2. While not explicitly represented in the model’s parameter set, the single parameter comprising the \mathbf{k}_d vector for that model is the difference between the elevation of the left fixed-head boundary as represented in the model and its true elevation. Whether a modeller is aware of it or not, this parameter is present in the model’s parameter set (conceptually at least) as long as he/she has the capacity to misassign this elevation. The fact that the modeller is unaware of the possibility of this error, or does not wish to consider the boundary elevation as an adjustable parameter, relegates this parameter to the \mathbf{k}_d vector rather than to the \mathbf{k}_m vector. Inasmuch as model outputs used in the calibration process are sensitive to this parameter, a corresponding \mathbf{Z}_d matrix exists.

As for most other linear analyses documented herein, we choose parameters to express perturbations from prior expected values of the system properties that they represent. If a parameter’s value is zero, this indicates that the parameter is awarded its prior expected value. The same applies to correction parameters. If a particular element of \mathbf{k}_d has a value of zero, then the defect which it represents does not exist. This is the case in figure 15.1 where the elevation of the left fixed-head boundary is correct.

Appropriate “defect parameters” exist in other modelling contexts. For example, if a steady-state groundwater model is subjected to history-matching under transient conditions, \mathbf{k}_d parameters comprise recharge rate perturbations from their assumed steady-state values, as well as storage properties that do not feature in steady state simulation; meanwhile the pertinent \mathbf{Z}_d matrix represents the time-varying component of the real-world system over the time interval for which steady-state conditions are assumed to prevail. Where a groundwater model is assumed to possess a single layer instead of the many layers which may actually characterize the subsurface, \mathbf{k}_d parameters are comprised of the differences between individual layer properties and the properties of the single layer which replaces them in the model; meanwhile the \mathbf{Z}_d matrix computes model output correction terms that emerge from these layer property differences and from vertical head gradients.

Through comparison of equations (15.3.4) and (15.3.1) it is obvious that it is to $\mathbf{Z}_d \mathbf{k}_d$ that the term “structural noise” should be applied. In practice, it is often ascribed to its visible expression, these being residuals that are incurred by its presence as a model is history-matched. As will be discussed below, however, it is its invisible expression that inflicts damage on inferences of parameter values.

15.3.2 Parameters and Residuals

In analyses that follow, we assume that history-matching is effected through model calibration. Singular value decomposition is employed to examine the effects of model defects on model-to-measurement misfit, and on estimated parameters. However all conclusions that are drawn below are just as applicable to ensemble-based history-matching. Chapters 11 and 12 make it clear that the mathematics of reducing model-to-measurement misfit through parameter adjustment is the same whether this is done in order to subdue or express parameter diversity.

Through calibration, a model's "adjustable" parameters are estimated. These are the \mathbf{k}_m parameters of equation (15.3.4). Suppose that, when subjected to singular value decomposition, \mathbf{Z}_m can be written as

$$\mathbf{Z}_m = \mathbf{U}_m \mathbf{S}_m \mathbf{V}^t_m \quad (15.3.6)$$

Following determination of a suitable singular value truncation point based on considerations presented in section (8.2.5), \mathbf{k}_m is estimated as

$$\underline{\mathbf{k}}_m = \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \mathbf{h} \quad (15.3.7)$$

Use of equation (15.3.7) to estimate $\underline{\mathbf{k}}_m$ implies that the model \mathbf{Z}_m has no defects. However in recognition of the fact that the model is indeed defective, we now substitute (15.3.4) into (15.3.7) to obtain (after a little simplification)

$$\underline{\mathbf{k}}_m = \mathbf{V}_{m1} \mathbf{V}_{m1}^t \mathbf{k}_m + \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \mathbf{Z}_d \mathbf{k}_d + \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \boldsymbol{\epsilon} \quad (15.3.8)$$

Suppose that \mathbf{Z}_d can be decomposed as

$$\mathbf{Z}_d = \mathbf{U}_d \mathbf{S}_d \mathbf{V}^t_d \quad (15.3.9)$$

Then (15.3.8) can also be written as

$$\underline{\mathbf{k}}_m = \mathbf{V}_{m1} \mathbf{V}_{m1}^t \mathbf{k}_m + \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \mathbf{U}_d \mathbf{S}_d \mathbf{V}^t_d \mathbf{k}_d + \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \boldsymbol{\epsilon} \quad (15.3.10)$$

Residuals associated with the calibrated model are

$$\mathbf{r} = \mathbf{h} - \mathbf{Z}_m \underline{\mathbf{k}}_m \quad (15.3.11)$$

With substitution of (15.3.8) into (15.3.11) this becomes

$$\mathbf{r} = \mathbf{h} - \mathbf{Z}_m \mathbf{V}_{m1} \mathbf{V}_{m1}^t \mathbf{k}_m - \mathbf{Z}_m \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \mathbf{Z}_d \mathbf{k}_d - \mathbf{Z}_m \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \boldsymbol{\epsilon} \quad (15.3.12)$$

Substituting (15.3.6) for \mathbf{Z}_m , this becomes

$$\mathbf{r} = \mathbf{h} - \mathbf{U}_{m1} \mathbf{S}_{m1} \mathbf{V}_{m1}^t \mathbf{k}_m - \mathbf{U}_{m1} \mathbf{U}_{m1}^t \mathbf{Z}_d \mathbf{k}_d - \mathbf{U}_{m1} \mathbf{U}_{m1}^t \boldsymbol{\epsilon} \quad (15.3.13)$$

If (15.3.4) is substituted for \mathbf{h} and use is made of the relationship

$$\mathbf{U}_{m1} \mathbf{U}_{m1}^t + \mathbf{U}_{m2} \mathbf{U}_{m2}^t = \mathbf{I} \quad (15.3.14)$$

and the relationship

$$\mathbf{U}_m \mathbf{S}_m \mathbf{V}^t_m = \mathbf{U}_{m1} \mathbf{S}_{m1} \mathbf{V}_{m1}^t + \mathbf{U}_{m2} \mathbf{S}_{m2} \mathbf{V}_{m2}^t \quad (15.3.15)$$

we finally obtain for residuals

$$\mathbf{r} = \mathbf{U}_{m2} \mathbf{S}_{m2} \mathbf{V}_{m2}^t \mathbf{k}_m + \mathbf{U}_{m2} \mathbf{U}_{m2}^t \mathbf{Z}_d \mathbf{k}_d + \mathbf{U}_{m2} \mathbf{U}_{m2}^t \boldsymbol{\epsilon} \quad (15.3.16a)$$

which, using (15.3.9) can also be written as

$$\mathbf{r} = \mathbf{U}_{m2} \mathbf{S}_{m2} \mathbf{V}_{m2}^t \mathbf{k}_m + \mathbf{U}_{m2} \mathbf{U}_{m2}^t \mathbf{U}_d \mathbf{S}_d \mathbf{V}^t_d \mathbf{k}_d + \mathbf{U}_{m2} \mathbf{U}_{m2}^t \boldsymbol{\epsilon} \quad (15.3.16b)$$

The first term on the right of equations (15.3.16) is nonzero if singular values are truncated before they become zero. There is a good reason for doing this, as section 8.2.5 explains.

However if a problem is well-posed (whereby, as explained in previous sections, singular value decomposition becomes equivalent to the Gauss-Newton method), this term disappears. The second and third terms on the right of (15.3.16) constitute vector projections of structural noise ($\mathbf{Z}_d \mathbf{k}_d$) and measurement noise $\boldsymbol{\varepsilon}$ respectively onto the orthogonal complement of the range space of \mathbf{Z}_m (i.e. onto the subspace that is spanned by \mathbf{U}_{m2}). Any vector that occupies this subspace cannot be expressed as a linear combination of the columns of \mathbf{Z}_m ; hence no parameters can be given to the model which will allow it to fit data that is described by this vector. All terms on the right of (15.3.16) lie in this subspace; being therefore “unfitable”, they comprise residuals.

Equations (15.3.8) and (15.3.10) make it obvious that values estimated for \mathbf{k}_m parameters are likely to be influenced by the fact that a model is defective. History-matching can therefore induce parameter bias. Quantification of how much they are thus influenced, the extent to which predictions of future system behaviour inherit bias because of this, and whether any defences can be mounted against contamination of estimated parameters by model defects, are issues which will be addressed shortly.

Equations (15.3.8) and (15.3.10) state that if a model has no defects and hence \mathbf{k}_d is zero, and if no measurement noise is associated with a measurement dataset, then estimated values for all \mathbf{k}_m parameters are “correct” in the sense that \mathbf{k}_m is the projection of \mathbf{k}_m onto the solution space of \mathbf{Z}_m . With the introduction of model defects, however, the situation becomes more complicated. To the extent that the range space of the \mathbf{Z}_d matrix is aligned with that of the \mathbf{U}_{m1} matrix, expressions of model defects are “captured” by the history-matching process and thereby affect values estimated for the solution space projection of \mathbf{k}_m parameters. This happens where \mathbf{u}_{1mi} vectors comprising the columns of \mathbf{U}_{1m} have non-zero projections onto \mathbf{u}_{dj} vectors comprising the columns of \mathbf{U}_d for which corresponding elements of the diagonal \mathbf{S}_d matrix are non-zero. (Diagonal elements of the \mathbf{S}_{1m} matrix are all non-zero by definition). In contrast, equation (15.3.16b) shows that, to the extent that any \mathbf{u}_{dj} vector whose corresponding \mathbf{S}_d singular value is seriously non-zero is orthogonal to all \mathbf{u}_{1mi} vectors (and hence has a non-zero scalar product with one or more \mathbf{u}_{2mi} vectors), the real world system which the model attempts to emulate produces “signals” which the model simply cannot capture, for there is no place within the simplified model to store this information. This contributes to the values of residuals, and hence to model-to-measurement misfit. It becomes the “visible” component of structural noise. However, the component of structural noise that can be (erroneously) absorbed by model parameters (i.e. the part that projects onto \mathbf{U}_{m1}) is the invisible (and far more insidious) component of structural noise.

The above considerations can be restated as follows.

A model provides a limited number of receptacles for the information content of a measurement dataset. These receptacles are the elements of \mathbf{k}_m . Meanwhile, the elements of \mathbf{k}_d are fixed in the defective model that is being used to simulate processes that are operative in an environmental system of interest. Because they are fixed, they do not provide receptacles for information contained within the measurement dataset. Furthermore they may be fixed at values that are different from those of the system under study. Information pertaining to these parameters may be available through the measurement dataset. Two things can happen to this information, neither of which constitutes its correct fate, namely to provide a correct estimate of \mathbf{k}_d . Some of this information may flow wrongly into \mathbf{k}_m receptacles, as the defective model’s \mathbf{k}_d receptacles are closed to entry of this information. Elements of \mathbf{k}_m are thus awarded erroneous values which compensate for model defects; in doing so they reduce the visible expression of structural noise by absorbing some of it. Other aspects of this information simply

have nowhere to go. The latter “orphaned” information thus express itself through the only option available to it, namely as residuals that contribute to model-to-measurement misfit, or so-called “structural noise”.

There is a fundamental difference, however, between “structural noise” and measurement noise. This difference calls into question the common practice of lumping the two together and treating them as one when subjecting a model to history-matching.

Suppose that an inverse problem is well-posed. Suppose also that, in accordance with recommended practice, weights are chosen such that equation (15.3.2) is obeyed after weight-transformation of the Z_m matrix. Then the visible component of measurement noise that manifests itself as residuals after minimization of the calibration objective function can be used to estimate the invisible component of this same noise; see equation (7.2.23). It is the invisible component of measurement noise which introduces errors to estimated parameters, as this noise is “absorbed” by the parameter estimation process. In contrast, the visible component of measurement noise appears as the last term in equations (15.3.16). As the same ϵ that appears in (15.3.16) also appears in equations (15.3.8) and (15.3.10) (and in similar equations derived in previous chapters which omit model defects), noise-induced parameter error can therefore be estimated.

Unfortunately, there is no such relationship between the visible and invisible components of structural noise. Without explicit knowledge of Z_d , its visible component says nothing about its invisible component. In fact its visible component may be zero while its invisible component (the component that inflicts errors on parameters) may be substantial; see figure 15.6.

Despite this, some attempts have been made to assign a covariance matrix to structural noise so that it can be treated in a similar way to measurement noise; see for example Oliver and Alfonzo (2018), Cooley (2004) and Cooley and Christensen (2006). The latter authors show that where the nature of model defects can be characterized stochastically, an empirically-determined covariance matrix of structural noise can sometimes be used to support minimum error variance estimation of averaged parameters over a heterogeneous model domain. If the measurement noise contribution to model-to-measurement misfit is small, however, problems can arise in implementing this procedure. These arise from the fact that the covariance matrix of structural noise may be singular or near-singular. Hence it cannot be inverted and then used as a weighting (i.e. Q) matrix.

Let structural noise be denoted using the vector η . Hence from (15.3.4)

$$\eta = Z_d k_d \quad (15.3.17)$$

Its covariance matrix can then be written as

$$C(\eta) = Z_d C(k_d) Z_d^t \quad (15.3.18)$$

$C(\eta)$ is an $n \times n$ matrix, where n is the number of observations comprising a history-matching dataset. The rank of this matrix cannot be greater than the rank of $C(k_d)$ from which it is derived. If $C(\eta)$ is not of full rank (i.e. n), it cannot be inverted. It can only be of full rank if the number of elements of k_d equals or exceeds n . This does not occur for the model depicted in figure 15.2. k_d has a dimension of 1 in this example. This is because a single, erroneous fixed-head boundary elevation parameter is responsible for head misfit in all 4 observation wells. Hence the 4×4 covariance matrix of structural noise that is associated with these head observations is singular, and therefore noninvertible. That is why the line of best fit does not pass through any of the measured heads. It is also why, as in this example, structural noise does not “cancel out”; hence more measurements do not reduce the uncertainties of estimated

parameters.

In many environmental modelling contexts history-matching datasets are large. Nevertheless structural noise may be dominated by only a few model defects, far fewer defects than there are members of the history-matching dataset. (This, after all, can be construed as the mark of a “good model”.) $C(\eta)$ is therefore likely to approach singularity in many real-world modelling contexts. Treating it as if it were measurement noise then becomes impossible. Furthermore, even if it does not approach singularly, neither Z_d nor k_d is generally known. Hence inference of the invisible component of structural noise (the component that does the damage) from its visible component (the component that contributes to residuals) is impossible. Its effect on parameter (and predictive) bias is therefore impossible to calculate.

15.4 Defence against Structural Noise

15.4.1 Theory

Equations (15.3.8) and (15.3.10) show that history-matching of a defective model may induce errors in the estimation of some parameters. We now investigate whether it is possible to history-match a defective model in a way that affords parameters some protection from these errors.

Let us suppose that we can find an orthonormal matrix W which has the following important property

$$WW^t Z_d \approx \mathbf{0} \quad (15.4.1)$$

WW^t is thus an orthogonal projection operator. Decomposing Z_d through (15.3.9) we obtain

$$WW^t U_d S_d V_d^t \approx \mathbf{0} \quad (15.4.2)$$

From (15.4.1), W is chosen such that all of its columns are orthogonal to all of the columns of Z_d (or at least approximately so). If the Z_d matrix does not have many columns (because the number of k_d parameters is small), then it should not be too hard to find such a matrix. For reasons that will become obvious in a moment, the larger is the range space of W (that is, the larger the number of linearly independent columns that W contains) the more useful is this matrix. A small dimensionality of k_d is also desirable.

Unfortunately, as any environmental model is likely to be replete with defects, the number of elements comprising k_d is likely to be large. However, as stated above, in many circumstances a few obstinate defects are likely to dominate the others, for if a model is defective in too many ways it cannot be used. The number of seriously non-zero singular values featured in the diagonal S_d matrix is thus likely to be small. The columns of W need only be orthogonal to those columns of U_d for which corresponding S_d singular values are significantly non-zero.

If (15.3.4) is pre-multiplied by WW^t , we obtain

$$WW^t h = WW^t Z_m k_m + WW^t Z_d k_d + WW^t \epsilon \approx WW^t Z_m k_m + WW^t \epsilon \quad (15.4.3)$$

Equation (15.4.3) defines an inverse problem through which solution for k_m can be sought. Because of (15.4.1) and (15.4.2), W must have fewer columns than rows if all of these columns are linearly independent, for if it is of full rank then W^t cannot possess a null space. The more linearly independent columns that W possesses, the more of the information content of h does WW^t preserve.

Equation (15.4.3) transforms a history-matching problem in which there are n observations into one in which there are effectively fewer than n observations. Under favourable conditions this

may not reduce the dimensionality of the calibration solution space; the fewer the number of parameters that are being estimated (i.e. the fewer is the number of elements in \mathbf{k}_m) the more likely is this to be the case. Under unfavourable conditions, however, the dimensionality of the solution space may be smaller than without pre-multiplication of (15.3.4) by $\mathbf{W}\mathbf{W}^t$. However parameter combinations that thereby become inestimable are those whose estimates would have been contaminated by structural noise anyway, for this is why we choose to premultiply the terms of (15.3.4) by $\mathbf{W}\mathbf{W}^t$.

15.4.2 Practice

In everyday modelling practice, it is generally not too difficult to find ways of processing field measurements and corresponding model outputs in ways that “orthogonalize out” (i.e. filter out) at least the worst manifestations of model-generated structural noise.

For example, when history-matching a groundwater model, vertical differences in heads between model layers can be used explicitly in formation of the objective function that is reduced through history-matching rather than (or in addition to) heads in individual layers. Intuitively this makes sense. Vertical head differences may be small; however they are likely to be rich in information on the vertical conductance of material that separates different model layers from each other. For many models, the integrity of calculated head differences is likely to be maintained even if the integrity of heads from which these differences are calculated is degraded by errors in representation of recharge and boundary conditions. When using PEST, these head differences can be placed into a different observation group from that which contains head measurements. Head differences can then be weighted in a way that guarantees visibility in the overall objective function.

The benefits of matching lateral head differences calculated by a model to measured lateral head differences is discussed above; see figure 15.5.

A similar logic applies to temporal head differences calculated by a groundwater model, especially where this model is used to support water use sustainability. The information content of temporal head differences with respect to storage and recharge-related parameters employed by a model is likely to be high, notwithstanding the fact that these differences may be small in relation to the heads themselves, and that the integrity of model-calculated heads from which these differences are calculated may be degraded by defective boundary conditions or erroneous representation of permeability heterogeneity. By extracting these temporal differences from head measurements and complementary model outputs before matching them, and by then assigning weights to temporal head differences that guarantee their visibility in the overall objective function, the history-matching process is likely to endow a model with parameters that allow it to replicate these temporal head differences, even if it cannot replicate the individual heads themselves. Integrity of estimation of parameters to which these differences are sensitive is thereby enhanced.

In surface water model calibration a modeller may, in similar fashion, formulate a multi-component objective function in which a streamflow time series is processed in a number of different ways in order to expose to the history-matching process different aspects of the streamflow time series that are rich in information pertaining to different sets of watershed properties. For example, peak flows are informative of surface parameters such as roughness; post-peak recessions are informative of soil moisture store parameters; baseflows are informative of groundwater store parameters; monthly-averaged flows are informative of parameters controlling seasonality of evapotranspiration, etc. For each type of processing to which the flow time series is subjected, a different objective function component can be

formulated. Processed outcomes of the observed streamflow time series are thereby compared with similarly processed outcomes of the modelled streamflow time series. The overall history-matching dataset then includes all of these information-rich outcomes of streamflow time series processing, with each weighted for visibility.

Because procedures such as these use the same measurement dataset in different ways, they do not explicitly implement equation (15.4.3). Moreover, they actually increase the number of elements in the history-matching dataset beyond that of the original \mathbf{h} . At the same time, the weighting scheme is unashamedly pragmatic.

In theory, use of equation (15.4.3) requires adoption of a transformed measurement noise vector $\boldsymbol{\tau}$ calculated as

$$\boldsymbol{\tau} = \mathbf{W}\mathbf{W}^t \boldsymbol{\varepsilon} \quad (15.4.4)$$

From (15.3.2), the covariance matrix associated with this transformed noise becomes

$$\mathbf{C}(\boldsymbol{\tau}) = \sigma^2_{\varepsilon} \mathbf{W}\mathbf{W}^t \quad (15.4.5)$$

$\mathbf{C}(\boldsymbol{\tau})$ is rank-deficient and hence non-invertible. Hence an “ideal” \mathbf{Q} matrix for use in the history-matching process described by equation (15.4.3) cannot be obtained through inversion of $\mathbf{C}(\boldsymbol{\tau})$.

In spite of this, deployment of heuristic strategies such as those described above can be very effective in protecting the history-matching process from the adverse effects of structural noise. The nature and extent of model structural defects is never exactly known. Hence determination of a suitable \mathbf{W} for use in equation (15.4.3) can only ever be approximate. In fact, a modeller may not even be sure whether, in a particular modelling context, such protection is required at all. Meanwhile, he/she may feel uncomfortable about reducing the effective size of a history-matching dataset, and possibly the size of the calibration solution space, by implementing $\mathbf{W}\mathbf{W}^t$ filtering. Hence, he/she may decide that the best strategy for handling potential model defects is to “place a bet each way”. A modeller can do this by using the measurement dataset in its original form, as well as in a processed form, or in multiply-processed forms. In balancing the relative weighting of different objective function components, it is probably best for the modeller to adopt a pragmatic strategy based on the philosophy that if an objective function component is worth formulating, then it should be visible in the overall calibration objective function. It is thereby protected against dominance of another objective function component, at the same time as it is not allowed to dominate the objective function itself.

Before closing this subsection, it is worth noting that prevention of bias in estimates of \mathbf{k}_m parameter values in ways that are described above may or may not be a good thing. This depends on predictions that are required of a model. For some predictions, the ability of \mathbf{k}_m parameters to respond to the information contained in a history-matching dataset by adopting biased values that compensate for model defects may allow these defects to be effectively “calibrated out” as far as the making of those predictions is concerned. For other predictions, bias introduced to \mathbf{k}_m parameters through history-matching may be transferred directly to the predictions themselves; for these predictions, history-matching may do more harm than good. The situation is not clear-cut. We will return to this point shortly.

15.5 Model Defects: a Pictorial Representation

Figure 15.8 depicts three-dimensional parameter space with orthogonal axes which coincide with the three parameters that define this space. Two of the parameters specified by these axes are “m” parameters; that is, they are adjustable parameters in a defective model of an

environmental system. The third parameter is a “d” parameter. Hence its value is hardwired into the defective model’s construction. Let us assume that, collectively, the three parameters span the entirety of parameter space. Hence replication of past and future system behaviour is possible using these three parameters.

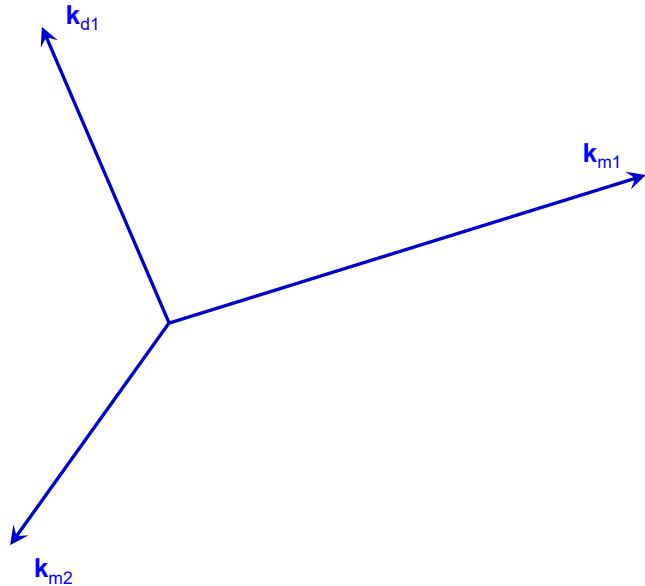


Figure 15.8. Two “m” parameters and one “d” parameter; collectively they span the entirety of parameter space. This and following figures are modified from Doherty and Christensen (2011).

Let the matrix \mathbf{Z} denote the “reality” model of the system. (The modeller does not have access to this model; he/she has access only to the \mathbf{Z}_m model). Thus

$$\mathbf{h} = \mathbf{Z} \begin{bmatrix} \mathbf{k}_m \\ \mathbf{k}_d \end{bmatrix} + \boldsymbol{\varepsilon} = [\mathbf{Z}_m \quad \mathbf{Z}_d] \begin{bmatrix} \mathbf{k}_m \\ \mathbf{k}_d \end{bmatrix} + \boldsymbol{\varepsilon} = \mathbf{Z}_m \mathbf{k}_m + \mathbf{Z}_d \mathbf{k}_d + \boldsymbol{\varepsilon} \quad (15.5.1)$$

We now subject the real-world model matrix \mathbf{Z} to singular value decomposition, so that

$$\mathbf{Z} = \mathbf{U} \mathbf{S} \mathbf{V}^t \quad (15.5.2)$$

The three \mathbf{v}_i vectors that comprise the columns of \mathbf{V} are added to the three native model parameter vectors in figure 15.9. \mathbf{v}_1 and \mathbf{v}_2 span the solution space of \mathbf{Z} while \mathbf{v}_3 spans its null space.

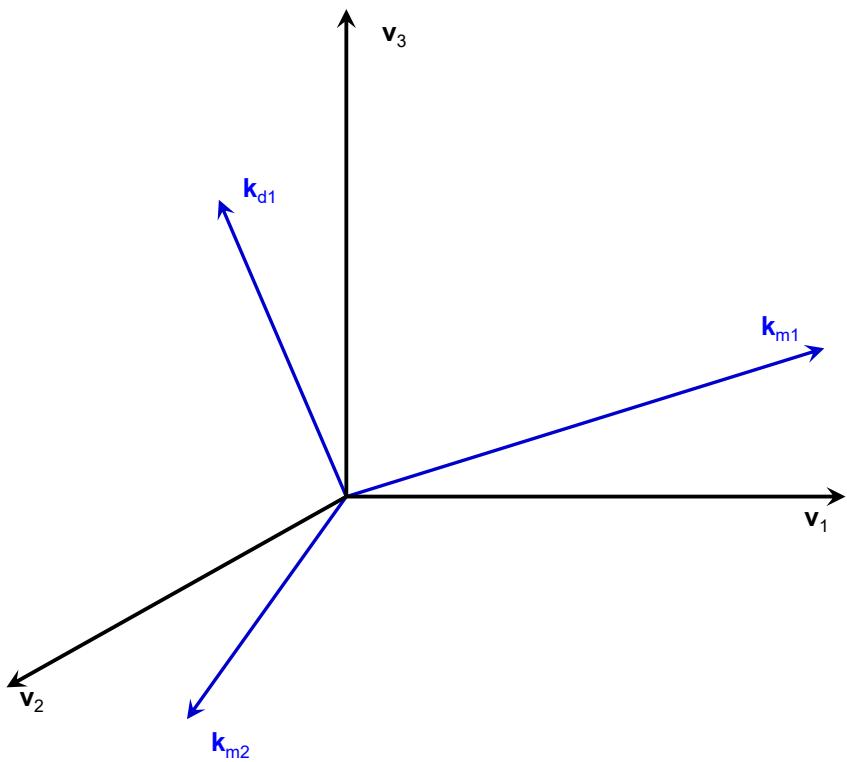


Figure 9.9. Model parameters together with the eigencomponents (shown in black) of the real-world model matrix Z .

Let the vector \mathbf{k} represent the three parameters of the real-world model \mathbf{Z} . Let us imagine, for a moment, that it is possible to actually build and then calibrate this model. The calibration process of the \mathbf{Z} model would yield the vector $\underline{\mathbf{k}}$ shown in figure 15.10. This is the projection of \mathbf{k} onto the real-world solution space. This estimate of real-world parameters is not, of course, correct. But it is of minimum error variance (and hence without bias) because the parameter set $\underline{\mathbf{k}}$ allows model outputs to fit the measurement dataset to a level that is commensurate with measurement noise while possessing no null space components; the latter are, by definition, unsupported by the measurement dataset.

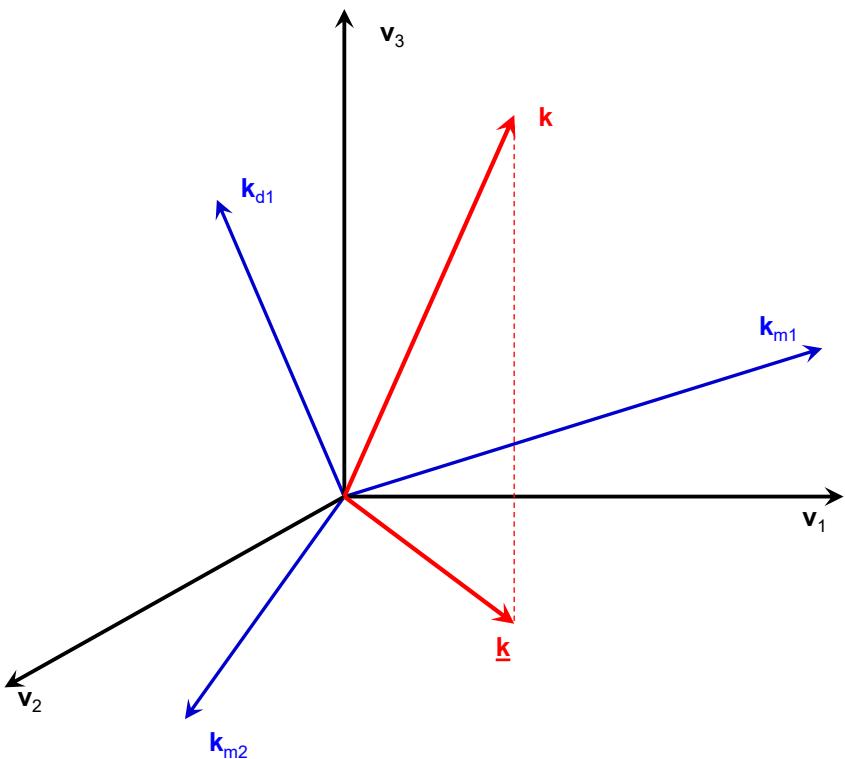


Figure 15.10. The true model parameter set \mathbf{k} and the parameter set $\underline{\mathbf{k}}$ (shown in red) that would be estimated through an “ideal” calibration process undertaken using the real-world model \mathbf{Z} .

Unfortunately, the real-world model \mathbf{Z} is not available to a modeller; he/she can calibrate only the simplified model \mathbf{Z}_m . It is through estimation of “m” parameters that a good fit is thereby sought with the history-matching dataset. In the present example there are just two of these, namely k_{m1} and k_{m2} , these comprising the \mathbf{k}_m parameter set. Meanwhile the third parameter k_d of the real-world model is fixed at a certain value, this value being implied in construction of the simplified, defective model.

As the dimensionality of the solution space of \mathbf{Z} is two, the two parameters of the simplified \mathbf{Z}_m model are enough to support a good fit between its outputs and the history-matching dataset (provided that neither of its parameters lies entirely within the null space of \mathbf{Z}). Let $\underline{\mathbf{k}}_m$ denote the parameter set achieved through calibration of the \mathbf{Z}_m model. The vector corresponding to this parameter set must lie in the k_{m1}/k_{m2} plane. At the same time, its projection onto the solution space of the real-world model (i.e. the space spanned by v_1 and v_2) must be $\underline{\mathbf{k}}$; if this is not the case then $\underline{\mathbf{k}}_m$ would not allow the simplified model to fit the calibration dataset. This is depicted in figure 15.11.

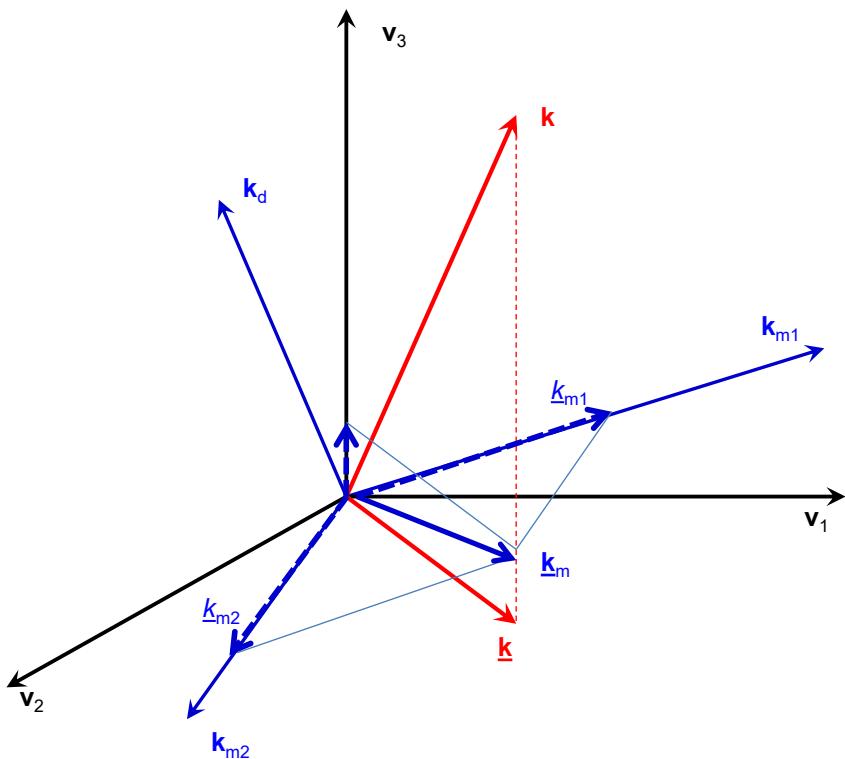


Figure 15.11. Calibration of the Z_m model through adjustment of only k_{m1} and k_{m2} with k_d fixed (at zero in this figure) leads to the vector \underline{k}_m (shown in blue) which projects onto the solution space of the real-world model Z as \underline{k} .

Immediately apparent from figure 15.11 is the fact that the projection of \underline{k}_m onto the v_3 axis is non-zero. Hence the solution of what is in reality a three-parameter, ill-posed inverse problem is non-optimal. This occurs because solution of that problem was obtained in a two-dimensional subspace that is not aligned with that which is obtained through singular value decomposition of the (unattainable) real-world model Z . Because this subspace is non-orthogonal to the solution space of Z it is possible to find a \underline{k}_m vector which fits the measurement dataset well. However because the two-dimensional Z_m subspace has a nonzero projection onto the null space of reality (i.e. Z), the \underline{k}_m vector has a non-zero projection onto that space. Optimality of solution to the inverse problem is therefore foregone.

k_d is fixed at zero in figure 15.11; according to the protocol that we have adopted in deriving many of the equations that are presented in this book, this is its minimum error variance estimate from an expert knowledge point of view. It is also the value that a modeller would choose in designing a simplified model. In reality, of course, k_d is probably nonzero, for it has variability that is described by its variance. If the acknowledged variability of all three real world model parameters is taken into account, together with the information content of the history-matching dataset, which is somewhat informative of all three model parameters (because none of them has zero projection onto the null space of Z) but not completely informative of any of them, the best “calibrated” value for k_d to take is given by the projection of \underline{k} (and not \underline{k}_m) onto the k_d parameter axis. Collectively \underline{k}_{m1} , \underline{k}_{m2} and \underline{k}_d will then have minimum error variance status from a post-calibration point of view; at the same time, the resulting vector \underline{k} will have zero projection onto the null space of Z . This cannot be accomplished using the simplified model. The result is parameter (and possibly predictive) bias.

Suppose that the real-world value of k_d actually differs from zero. Then a vector of length k_d in the direction of the k_d axis (we call this vector \mathbf{k}_d) has a non-zero projection onto the $\mathbf{v}_1/\mathbf{v}_2$ solution space of the real-world model \mathbf{Z} . In calibrating the \mathbf{Z}_m model, the projection of \mathbf{k}_d onto the $\mathbf{v}_1/\mathbf{v}_2$ subspace must be added to the projections of the k_{m1} and k_{m2} components of \mathbf{k}_m onto that subspace such that the total projection of \mathbf{k}_m plus \mathbf{k}_d has the correct projection onto the $\mathbf{v}_1/\mathbf{v}_2$ subspace. The correct projection of $\mathbf{k}_m + \mathbf{k}_d$ onto the $\mathbf{v}_1/\mathbf{v}_2$ subspace is, once again, \mathbf{k} . Given that the $\mathbf{v}_1/\mathbf{v}_2$ subspace is two dimensional, and that unit vectors in the k_{m1} and k_{m2} directions are orthogonal to each other and to \mathbf{k}_d , it will always be possible to achieve the correct projection onto \mathbf{k} regardless of the value of k_d . However the values that the individual k_{m1} and k_{m2} parameters must take to achieve this correct projection will depend on the value of k_d . Hence the values that k_{m1} and k_{m2} take through the history-matching process (namely k_{m1} and k_{m2}) are dependent on the nature and magnitude of the model's defects. At the same time, the total $\mathbf{k}_d + \mathbf{k}_m$ vector is likely to have a non-zero projection onto the null space of the real-world model \mathbf{Z} . This is illustrated in Figure 15.12.

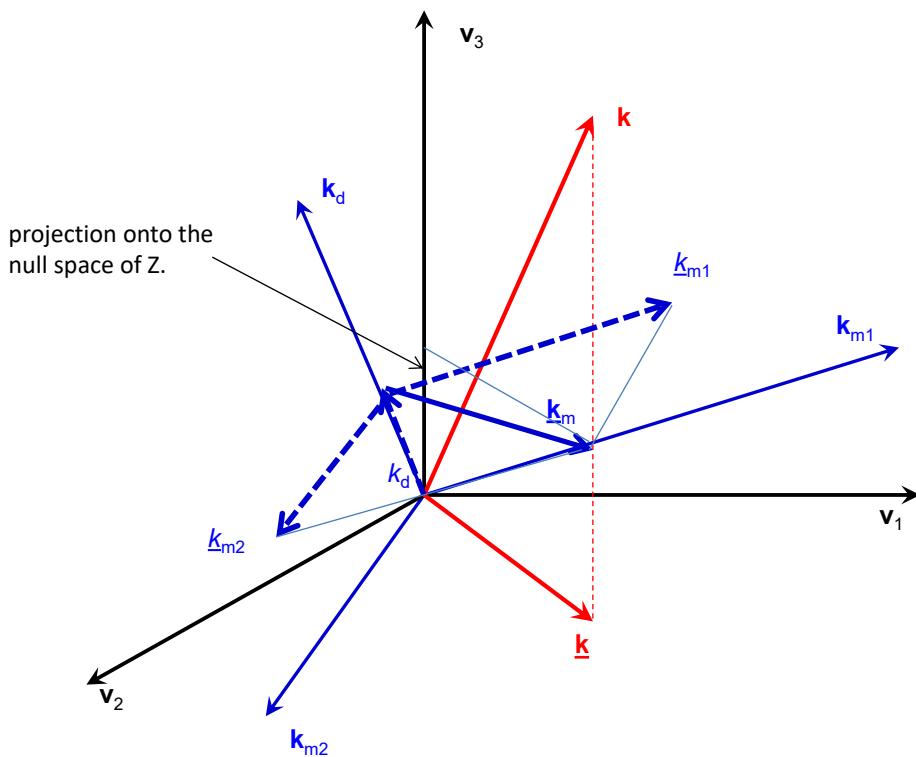


Figure 15.12. Same as figure 15.11, but with k_d having a non-zero value.

It is the measurement dataset \mathbf{h} , together with its information content, that define \mathbf{Z} , the singular value decomposition of \mathbf{Z} , and the \mathbf{V} matrix that emerges from singular value decomposition of \mathbf{Z} . For the model illustrated in the above figures, the measurement dataset has information which can resolve two *combinations* of parameters, but cannot resolve any *individual* parameter. Each of the three model parameters k_{m1} , k_{m2} and k_d project partly onto the solution space of \mathbf{Z} and partly onto its null space. Because the information content of the measurement dataset is shared between them, and because none of them can be resolved individually, the fixing of a value for any one of these parameters creates a well-posed inverse problem through which values can be estimated for the other two. However these values are dependent on the value ascribed to the fixed parameter. Hence any error in the value ascribed to the fixed parameter will be immediately transferred to the adjustable parameters. Unfortunately, the

value at which the fixed parameter is actually fixed is unlikely to be optimal from the point of view of a calibration process that takes into account the true unknown status of all three parameters. That value is determined through projection of $\underline{\mathbf{k}}$ (the ideal solution to the inverse problem based on the unknown real-world model matrix \mathbf{Z}) onto each of the individual parameter axes. $\underline{\mathbf{k}}$, of course, has a projection of zero onto the null space of \mathbf{Z} .

Any model is a simplification of reality. Construction of any model is thus likely to constitute imperfect regularization. It follows that, regardless of whether model history-matching takes place using ensembles or through regularized inversion, regularization has already taken place. Furthermore, this regularization is probably suboptimal.

In previous chapters of this book, optimal regularization has been defined. Optimal model *simplification* can now be defined. In fact the two definitions are aligned. If construction of a model \mathbf{Z}_m is such that simplifications of the real-world \mathbf{Z} model that are implied in its construction coincide with those which would have been obtained through singular value decomposition of \mathbf{Z} , then \mathbf{Z}_m model construction is optimal.

So if \mathbf{Z}_m model construction is such that \mathbf{k}_d parameters that represent \mathbf{Z}_m model defects lie entirely within the null space of \mathbf{Z} , this is a reasonable model construction outcome. \mathbf{k}_d parameters then have no effect on \mathbf{Z} model outputs under history-matching conditions. Hence they cannot induce errors in \mathbf{k}_m parameters that are adjusted as the \mathbf{Z}_m model is history-matched. They do have the capacity, of course, to introduce errors to \mathbf{Z}_m model predictions that are sensitive to \mathbf{k}_d defects.

Strategies for filtering out structural noise which are described in the preceding section of this chapter can be seen as attempts to re-align \mathbf{Z}_m model construction with singular value decomposition of \mathbf{Z} . For the three-parameter example described above, such a process would attempt to remove all information from the measurement dataset that is sensitive to the k_d parameter. k_d is therefore relegated to the null space of the revised (as far as history-matching is concerned) \mathbf{Z}_m model. This precludes the possibility of k_{m1} and k_{m2} parameters adopting values that compensate for an erroneous k_d value.

Yet another important lesson about defective model history-matching can be learned from the above three parameter example. Suppose that no steps are taken to mitigate the history-matching repercussions of assigning an erroneous value to k_d . As has been shown, this does not mean that the simplified, defective \mathbf{Z}_m model cannot be calibrated to achieve a good fit with a history-matching dataset. Regardless of the value taken by k_d , values for k_{m1} and k_{m2} can be found such that $\underline{\mathbf{k}}_m + \mathbf{k}_d$ projects exactly onto $\underline{\mathbf{k}}$. As has been discussed, the latter vector is the optimal solution to the ill-posed inverse problem of history-matching the real-world model \mathbf{Z} . This is because it is the solution that would have been obtained (if this were possible) through singular value decomposition of \mathbf{Z} while acknowledging adjustability of both \mathbf{k}_d and \mathbf{k}_m .

A direct consequence of this is that any prediction that is sensitive solely to solution space components of the \mathbf{Z} matrix (as distinct from solution space components of the \mathbf{Z}_m matrix), will be unaffected by the compensatory roles that k_{m1} and k_{m2} may need to play in history-matching the defective model; nor will they be sensitive to k_d defects that have been “calibrated out” because of k_{m1} and k_{m2} compensation. These types of predictions therefore deserve special consideration; they are not like other types of model predictions. In general, predictions which tend to be sensitive solely to solution space components of the real-world tend to be those that are similar in nature and location to measurements comprising a measurement dataset. Discussion of this issue will continue in the next section.

15.6 Defect-Induced and Defect-Immune Predictive Error

15.6.1 Linear Analysis

15.6.1.1 Formulation of Equations

From equation (15.3.8) the error in estimated \mathbf{k}_m parameters is given by

$$\underline{\mathbf{k}}_m - \mathbf{k}_m = -\mathbf{V}_{m2}\mathbf{V}_{m2}^t\mathbf{k}_m + \mathbf{V}_{m1}\mathbf{S}_{-1m1}\mathbf{U}_{m1}^t\mathbf{Z}_d\mathbf{k}_d + \mathbf{V}_{m1}\mathbf{S}_{-1m1}\mathbf{U}_{m1}^t\boldsymbol{\epsilon} \quad (15.6.1)$$

As usual, let s be a prediction of management interest. Its true value is given by

$$s = \mathbf{y}_{m1}^t\mathbf{k}_m + \mathbf{y}_{d1}^t\mathbf{k}_d \quad (15.6.2)$$

where the vector \mathbf{y}_m expresses sensitivities of s to \mathbf{k}_m model parameters while the vector \mathbf{y}_d expresses sensitivities of s to \mathbf{k}_d parameters. However the value of the prediction made by the defective, history-matched model is given by

$$\underline{s} = \mathbf{y}_{m1}^t\underline{\mathbf{k}}_m \quad (15.6.3)$$

Equation (15.6.3) assumes that \mathbf{k}_d is fixed at zero in the defective model. This is its prior expected value. It is reasonable to assume that any simplification strategy would attempt to achieve this conceptual outcome.

The error in the prediction made by the simplified model is calculated by subtracting (15.6.2) from (15.6.3) to obtain

$$\underline{s} - s = \mathbf{y}_{m1}^t(\underline{\mathbf{k}}_m - \mathbf{k}_m) - \mathbf{y}_{d1}^t\mathbf{k}_d \quad (15.6.4)$$

After substitution of (15.6.1) into (15.6.4) followed by some simplification, we obtain

$$\underline{s} - s = -\mathbf{y}_{m1}^t\mathbf{V}_{m2}\mathbf{V}_{m2}^t\mathbf{k}_m + \mathbf{y}_{m1}^t\mathbf{V}_{m1}\mathbf{S}_{-1m1}\mathbf{U}_{m1}^t\boldsymbol{\epsilon} + (\mathbf{y}_{m1}^t\mathbf{V}_{m1}\mathbf{S}_{-1m1}\mathbf{U}_{m1}^t\mathbf{Z}_d - \mathbf{y}_{d1}^t)\mathbf{k}_d \quad (15.6.5)$$

Under the assumption that \mathbf{k}_m and \mathbf{k}_d are statistically independent, predictive error variance can then be calculated as follows

$$\begin{aligned} \sigma_{\underline{s}-s}^2 = & \mathbf{y}_{m1}^t\mathbf{V}_{m2}\mathbf{V}_{m2}^t\mathbf{C}(\mathbf{k}_m)\mathbf{V}_{m2}\mathbf{V}_{m2}^t\mathbf{y}_m + \mathbf{y}_{m1}^t\mathbf{V}_{m1}\mathbf{S}_{-1m1}\mathbf{U}_{m1}^t\mathbf{C}(\boldsymbol{\epsilon})\mathbf{U}_{m1}\mathbf{S}_{-1m1}\mathbf{V}_{m1}^t\mathbf{y}_m + \\ & (\mathbf{y}_{m1}^t\mathbf{V}_{m1}\mathbf{S}_{-1m1}\mathbf{U}_{m1}^t\mathbf{Z}_d - \mathbf{y}_{d1}^t)\mathbf{C}(\mathbf{k}_d)(\mathbf{y}_{m1}^t\mathbf{V}_{m1}\mathbf{S}_{-1m1}\mathbf{U}_{m1}^t\mathbf{Z}_d - \mathbf{y}_{d1}^t)^t \end{aligned} \quad (15.6.6)$$

This equation can be simplified by making use of the assumptions embodied in equations (15.3.2) and (15.3.5). Let us also make the following additional assumption.

$$\mathbf{C}(\mathbf{k}_d) = \sigma_{kd}^2 \mathbf{I} \quad (15.6.7)$$

Equation (15.6.6) then becomes

$$\begin{aligned} \sigma_{\underline{s}-s}^2 = & \sigma_{km}^2 \mathbf{y}_{m1}^t\mathbf{V}_{m2}\mathbf{V}_{m2}^t\mathbf{y}_m + \sigma_{\epsilon}^2 \mathbf{y}_{m1}^t\mathbf{V}_{m1}\mathbf{S}_{-2m1}\mathbf{V}_{m1}^t\mathbf{y}_m + \\ & \sigma_{kd}^2 (\mathbf{y}_{m1}^t\mathbf{V}_{m1}\mathbf{S}_{-1m1}\mathbf{U}_{m1}^t\mathbf{Z}_d - \mathbf{y}_{d1}^t) (\mathbf{y}_{m1}^t\mathbf{V}_{m1}\mathbf{S}_{-1m1}\mathbf{U}_{m1}^t\mathbf{Z}_d - \mathbf{y}_{d1}^t)^t \end{aligned} \quad (15.6.8)$$

The first two terms on the right side of equation (15.6.8) are identical to those on the right of equation (8.2.34). As is discussed in section (8.2.5) and illustrated in figure 8.4, these terms can be used to compute the point at which singular values are best truncated during SVD-based calibration. If truncation takes place at too few singular values, the history-matching process has effectively denied passage of information from a measurement dataset to model parameters. This information could have reduced the error variances of those parameters, and of predictions that are sensitive to them. However if truncation takes place too late, then the information content of a measurement dataset becomes excessively contaminated by noise associated with that dataset as it passes from observations to parameters. Estimates of parameters (and of

predictions that are sensitive to them) therefore lose their credibility.

Contributions that simplifications required for model construction make to predictive error variance are described by the third term on the right side of (15.6.8). Unfortunately a modeller cannot calculate this term as he/she does not have access to the Z_d matrix that simulates differences between his/her model and the real world, nor to the parameters k_d on which Z_d operates. Unfortunately, however, this term can have a significant impact on predictive error variance, and on optimization of predictive error variance.

Lack of direct access to Z_d and k_d does not necessarily mean that equation (15.6.8) is unusable. In everyday modelling practice, calculations based on equation (15.6.8) can proceed if a modeller introduces extra parameters to his/her model to represent at least some of the environmental system complexity that his/her model ignores. These will often be parameters that a modeller would not actually subject to history-match adjustment, but for which calibration and predictive sensitivities can nevertheless be calculated for inclusion in the Z_d matrix and the y_d vector used in equation (15.6.8). In some cases these parameters may be viewed as surrogates for other types of complexity that differentiate a model from the real-world system that the model is meant to simulate.

By way of example, for the purposes of linear analysis using (15.6.8), a groundwater modeller may parameterize some of the boundary conditions of his/her model that were fixed during history-matching. For example, some boundary conditions may have been assigned a “no-flow” status during history-matching, but are endowed with adjustable inflow/outflow values for the purpose of sensitivity calculation. A groundwater modeller may also parameterize historical pumping rates that were endowed with assumed values during history-matching. A surface water modeller may add temporally and spatially varying historical rainfall factors to the parameter set employed by his/her model in order to explore the repercussions of history-matching that model using rain that fell in an unrepresentative gauge. In all of these examples, implementation of equation (15.6.8) requires only that sensitivities with respect to the added parameters be calculated. It does not require that these parameters be subjected to history-match adjustment.

Equation (15.6.8) shows that computation of predictive error variance where the existence of model defects is acknowledged requires the summation of three terms, rather than the two that are required if a model’s simulation of real-world processes is perfect. These terms sum to the total predictive error variance. Figure 15.13 modifies figure 8.4 to illustrate this.

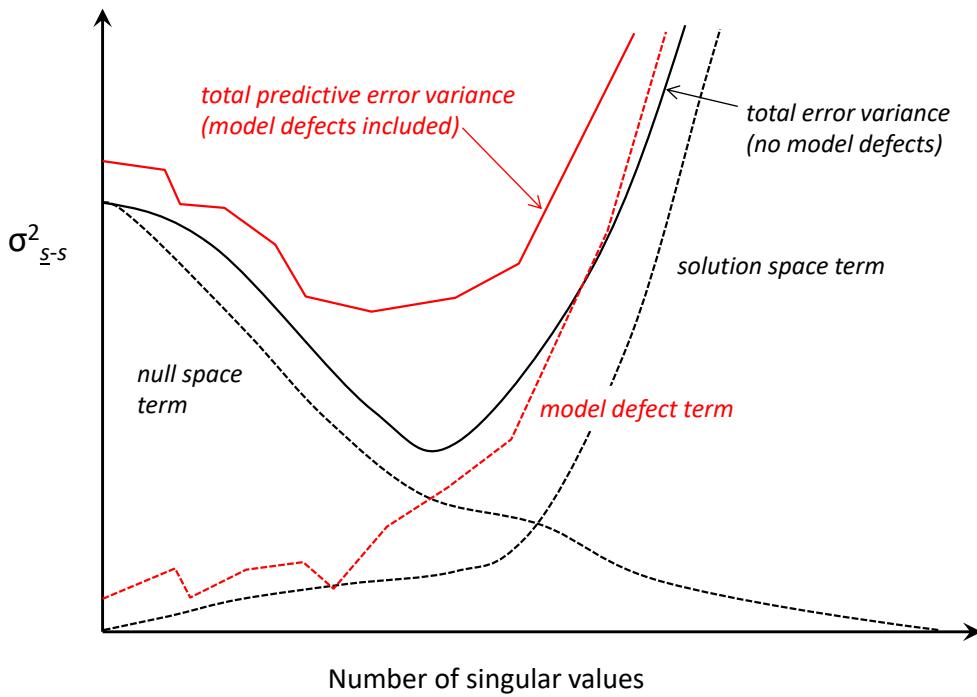


Figure 15.13. The three terms of equation (15.6.8).

15.6.1.2 Insights Gained from Linear Model Defect Analysis

Equation (15.6.8) provides many useful insights into the effects that model defects have on parameters that are estimated through history-matching, and on predictions that are made by a history-matched, defective model. These insights can have a profound effect on the way that environmental modelling is used to support environmental management. Furthermore, as stated earlier in this chapter, these insights are applicable whether a model is subjected to regularized inversion of ensemble-based Bayesian analysis.

The last term of equation (15.6.8) can only be zero or positive. However, unlike the first and second terms of this equation, there can be no guarantee that it increases monotonically with increasing number of pre-truncation singular values. Eventually, however, it must rise with singular value index because it features two incidences of the S^{-1}_{m1} matrix. When singular values of Z_m become zero, it rises to infinity. (Recall that singular values decrease with increasing singular value index.)

Where singular value index is low, the rate at which this term increases as history-matching is effected with an increasing number of singular values can be slow or rapid, this depending on the prediction. For some predictions, this term will be non-zero at the very left of the graph where the number of pre-truncation singular values is zero, and therefore the model is not subjected to history-matching at all. This occurs where model defects that are encapsulated in Z_d contribute to the error of the prediction purely because of sensitivity of the prediction to defect parameters k_d . With inclusion of zero singular values in the inversion process (and hence no history-matching at all), equation (15.6.8) becomes

$$\sigma^2_{ss} = \sigma^2_{km} y_m^t y_m + \sigma^2_{kd} y_d^t y_d \quad (15.6.9)$$

Whether the total predictive error variance falls below its initial value as the number of singular values employed in the history-matching process rises above zero also depends on the prediction. Experience in using equation (15.6.8) in real-world modelling contexts

demonstrates that it is far from impossible for the total error variance to rise and never fall as the number of singular values increases. This type of behaviour is normally associated with predictions that do not benefit much from model history-matching because they are not sensitive to hydraulic properties that are well-informed by a measurement dataset. It follows that history matching can do more harm than good for these predictions!

This has profound consequences for the way in which a model is built and deployed if the support that modelling provides to decision-making is based on these predictions. A modeller may decide to eschew the “benefits” of parameter-based history-matching in recognition of the fact that history-matching may actually do more harm than good. Instead, he/she may build a model that is designed to express expert knowledge with some degree of sophistication, taking into account the fundamentally stochastic nature of expert knowledge. This may involve the use of random, categorical realizations of hydraulic properties based on complex geostatistical characterization of the study site. Uncertainty quantification and reduction is then best achieved using data space inversion.

For other predictions, the overall predictive error variance may fall with increasing number of singular values. At some number of singular values the predictive error variance is minimized. However because the defect term of equation (15.6.8) is not necessarily monotonic, the total error variance curve may have multiple minima. At some point however, total error variance rises significantly. Meanwhile its minimum will have been shifted to the left (possibly by a large number of singular values) when compared to that which would have been calculated if only the first two terms of (15.6.8) are used. (These are generally the only terms that are available to a modeller.)

This shift to the left that is undergone by the total predictive error variance minimum indicates that in history-matching a defective model, a modeller should seek a reduced level of fit from that which he/she is entitled to seek when history-matching a perfect model. This, of course, makes sense as it recognizes the existence of structural noise in addition to measurement noise. Unfortunately, however, in most real-world modelling circumstances, a modeller cannot know the extent to which his/her goodness-of-fit ambitions should be reduced.

A significant outcome of the existence of model defects is that the singular value location of the minimum of the predictive error variance curve may be significantly different for some predictions than it is for others. This is in contrast to the case of no model defects where, although error variance reduction may differ from prediction to prediction, the number of singular values corresponding to the minimum of predictive error variance is unlikely to differ by too much from prediction to prediction. Sadly, where the existence of model defects is taken into account, the minimum of a predictive error variance curve may be indistinct, possibly serrated, and highly prediction-specific. Hence while, for some predictions, model history-matching may enable a considerable reduction in uncertainty, for others (as already stated) even a little history-matching can do more harm than good.

The popular belief that a history-matched groundwater model is good for the making of predictions of many and varied types is therefore called into question. See the epilogue of this book for more on this subject.

15.6.1.3 Prediction Immunity

In the previous section of this chapter, it was suggested that some predictions may carry virtual immunity from the deleterious effects of model defects. This will now be demonstrated using equation (15.6.5). For these predictions, a good fit should be sought with a history-matching

dataset regardless of the surrogate roles that \mathbf{k}_m parameters may adopt to achieve this fit, and regardless of the fact that the values estimated for some parameters may transgress the limits of system property credibility. It was suggested above that these predictions are those that are sensitive only to parameters which comprise the solution space of the real-world model \mathbf{Z} (as distinct from the solution space of the defective model \mathbf{Z}_d). These tend to be predictions which bear similarities to measurements that comprise a history-matching dataset.

The third term on the right side of (15.6.5) is zero if

$$\mathbf{y}_d^t \mathbf{k}_d = \mathbf{y}_m^t \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \mathbf{Z}_d \mathbf{k}_d \quad (15.6.10a)$$

This can also be written as

$$[\mathbf{y}_m \quad \mathbf{y}_d]^t \begin{bmatrix} \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \mathbf{Z}_d \\ -\mathbf{I} \end{bmatrix} \mathbf{k}_d = \mathbf{0} \quad (15.6.10b)$$

Suppose that measurement noise is zero. Suppose also that model imperfections can be “calibrated out” because \mathbf{k}_m parameters can be endowed with values that compensate for non-zero \mathbf{k}_d parameter values. It follows that for any alteration $\delta \mathbf{k}_d$ to the defect parameter set \mathbf{k}_d which alters model outputs from \mathbf{h} to $\delta \mathbf{h}$ so that

$$\delta \mathbf{h} = \mathbf{Z}_d \delta \mathbf{k}_d \quad (15.6.11)$$

an altered \mathbf{k}_m parameter set can be found such that

$$\delta \mathbf{h} = \mathbf{Z}_d \delta \mathbf{k}_d = \mathbf{Z}_m \delta \mathbf{k}_m = \mathbf{U}_m \mathbf{S}_m \mathbf{V}_m^t \delta \mathbf{k}_m = \mathbf{U}_{m1} \mathbf{S}_{m1} \mathbf{V}_{m1}^t \delta \mathbf{k}_m \quad (15.6.12)$$

The last two equalities in equation (15.6.12) follow from (15.3.6). From (15.6.12)

$$\delta \mathbf{k}_m = \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \mathbf{Z}_d \delta \mathbf{k}_d \quad (15.6.13)$$

From (15.6.13) and (15.6.12) it then follows that

$$[\mathbf{Z}_m \quad \mathbf{Z}_d] \begin{bmatrix} \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \mathbf{Z}_d \\ -\mathbf{I} \end{bmatrix} \delta \mathbf{k}_d = \mathbf{Z} \begin{bmatrix} \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \mathbf{Z}_d \\ -\mathbf{I} \end{bmatrix} \delta \mathbf{k}_d = \mathbf{0} \quad (15.6.14)$$

where \mathbf{Z} is the real-world model matrix. From this it can be concluded that the vector $[\mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \mathbf{Z}_d \\ -\mathbf{I}] \delta \mathbf{k}_d$ is in the null space of the real world model \mathbf{Z} . Equation (15.6.10.b) therefore states that a prediction will be immune to model structural defects if that prediction is insensitive to null space components of \mathbf{Z} and is therefore sensitive only to solution space components of \mathbf{Z} , as was stated above. For the making of such a prediction, the model can be virtually a “black box”, as long as it is a well-calibrated black box. The principle requirement for the making of data-driven predictions such as these is therefore machine learning, and not expert knowledge.

This conclusion makes the author wonder how many environmental models have been decreed as “good” simply because they have demonstrated alacrity in making one particular type of prediction. The above theory shows that when a model has defects, history-match enabled aptitude for the making of some types of predictions may guarantee suboptimal performance in the making of other types of predictions. Assessment of a model cannot be separated from predictions that the model is required to make.

White et al (2014) demonstrate use of equation (15.6.8) in exploring errors associated with predictions made by a defective groundwater model. They also demonstrate how this equation can be used to explore the efficacy of different objective function formulations in reducing calibration-induced predictive error.

15.6.2 Nonlinear Analysis

A methodology that has much in common with data space inversion is now presented. However, in contrast to data space inversion where a complex numerical model is replaced by a surrogate statistical model in order to make history-match-constrained predictions and quantify their uncertainties, a complex numerical model is replaced by a machine-learning-trained simple numerical model.

With the slightest of re-arrangements, equation (15.6.5) can be re-written as

$$s = \underline{s} + \mathbf{y}_m^t \mathbf{V}_{m2} \mathbf{V}_{m2}^t \mathbf{k}_m - \mathbf{y}_m^t \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \boldsymbol{\varepsilon} - \mathbf{y}_m^t \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \mathbf{Z}_d \mathbf{k}_d + \mathbf{y}_d^t \mathbf{k}_d \quad (15.6.15)$$

This formula was used by Doherty and Christensen (2011) who developed a methodology for joint use of two models - a complex model and a complementary simple model – at the same study site. The simple model is calibrated against a measurement dataset, and is then used to make predictions of management interest and quantify their uncertainties. However before it is deployed in this manner, corrections for calibration-induced predictive bias (to which a simple model is prone for reasons outlined in previous sections of this chapter) are established through repeatedly calibrating the simple model against synthetic calibration datasets computed by the complex model as the latter is endowed with a suite of stochastic hydraulic property fields that express expert knowledge of conditions at a study site. This methodology can be used where complex model run times are long, and/or where site conditions and processes are complex. Furthermore, the complex model can be endowed with geostatistically-based, categorical parameter fields that are difficult or impossible to adjust through history-matching. This frees the complex model to express expert knowledge, and the simple model (which is supposedly endowed with parameters that can be readily adjusted through inversion) to respond to information that is available in a history-matching dataset.

As usual, let \mathbf{h} denote a measurement dataset that is available at a study site. The complex model is used to simulate conditions that prevailed at the site over the history-matching period. This is done many times; on each occasion, it is provided with a different stochastic hydraulic property field. We denote each set of complex model outputs that correspond to the measurement dataset \mathbf{h} as \mathbf{h}_{ci} ; in this notation, the “c” subscript symbolizes “complex-model-generated” while “i” symbolizes the i ’th hydraulic property realization. (Note that each set of complex model outputs should be sprinkled with a realization of measurement noise.)

Suppose that this is done N times. Suppose also that, for each of these N hydraulic property fields, predictions of management interest are also made using the complex model. We will focus on just one of these predictions which we denote as s ; the prediction made using the i ’th hydraulic property field is denoted as s_i . If N is large enough, the prior probability distribution of s is thereby explored.

Now let the simple counterpart to the complex model be calibrated against each of the N realizations of \mathbf{h}_{ci} . Let it be supposed that the simple model is complex enough for it to reproduce these \mathbf{h}_{ci} reasonably well. Thus, for each \mathbf{h}_{ci} , a simple model parameter field is calculated corresponding to the complex model hydraulic property field. Each of these calibrated parameter fields endows the simple model with the capacity to reproduce the respective \mathbf{h}_{ci} .

Suppose also that, once the i ’th simple model has been calibrated against the i ’th complex model to obtain the i ’th simple model parameter field, the simple model is then run into the future to make the prediction s . We denote the i ’th prediction made by the thus-calibrated simple model as \underline{s}_i .

After the above calibration and prediction processes have been repeated N times, s_i can be plotted against \underline{s}_i for all i . Figure 15.14 schematizes such a plot. The prior uncertainty of the prediction is shown in the plot; it comprises the total span of complex model predictions.

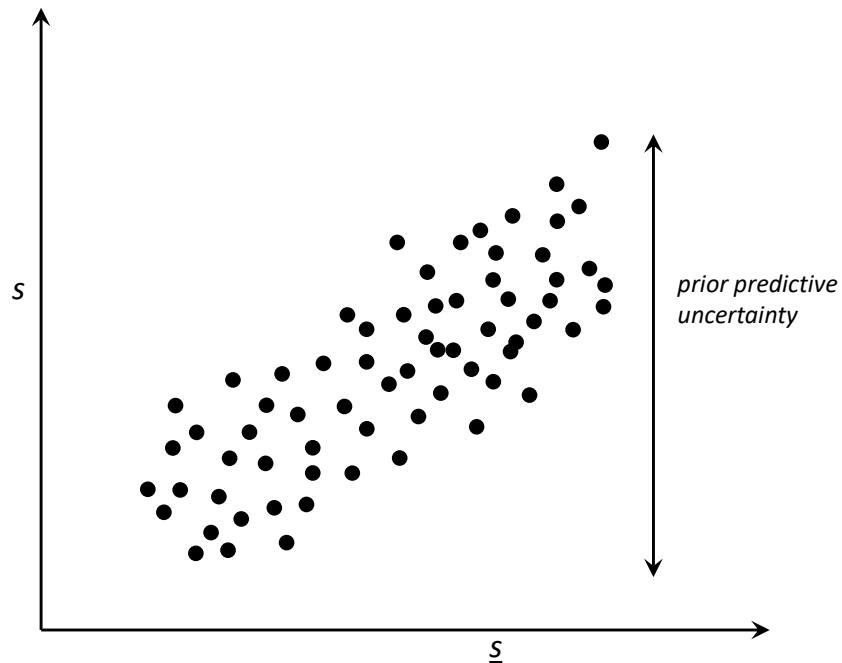


Figure 15.14. Plot of the value of a prediction s made using different complex model hydraulic property field realizations against predictions \underline{s} made by a partnered simple model. Each simple model is calibrated against a complex-model-generated measurement dataset whose contents mimic a real-world measurement dataset.

The simple model is now calibrated just one more time. However this time it is calibrated against the real-world measurement dataset \mathbf{h} . The calibrated simple model is then used to make the prediction of interest. Suppose that the value that it calculates for this prediction is \underline{s}_r ; the subscript “r” stands for “real”. The post-calibration potential for error of this prediction (which we use as a surrogate for the posterior uncertainty of the prediction) can then be obtained in the manner suggested by figure 15.15. The predictive interval depicted in this figure records the range of complex model predictions that is compatible with the prediction made by the simple model when the latter is calibrated against the real-world measurement dataset \mathbf{h} .

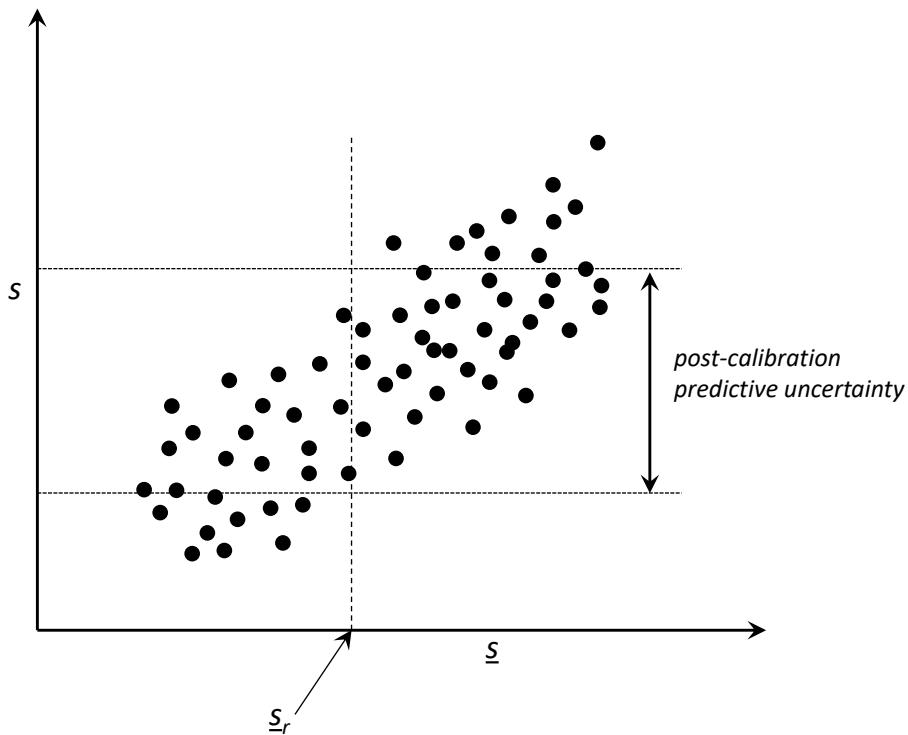


Figure 15.15. A simple model calibrated against the real-world calibration dataset makes the prediction s_r . The post-calibration uncertainty of that prediction is established as the range of complex model predictions through which a vertical line through s_r passes.

Various aspects of figures 15.14 and 15.15 can be explained by equation (15.6.15).

The second term on the right of (15.6.15) is $\mathbf{y}_m^t \mathbf{V}_{m2} \mathbf{V}_{m2}^t \mathbf{k}_m$. This represents predictive error that prevails if there is no noise in the history-matching dataset and no defects in the model. It is simply the null space contribution to post-calibration predictive error, and hence to posterior predictive uncertainty. Because this affects s rather than s_r in equation (15.6.15), it promotes vertical scatter of the points depicted in figures 15.14 and 15.15. If this is the only contribution to posterior model predictive uncertainty (because measurement noise is zero and simple model defects are either non-existent or do not affect the prediction s) then the line of best fit through the s vs. s_r scatterplot has a slope of 1.

The third term on the right of equation (15.6.15) (i.e. $\mathbf{y}_m^t \mathbf{V}_{m1} \mathbf{S}_{m1}^{-1} \mathbf{U}_{m1}^t \boldsymbol{\varepsilon}$) describes the effect of measurement noise on predictive error. As this affects s_r and not s , it promotes horizontal scatter of points that appear in figures 15.14 and 15.15. In many groundwater modelling contexts, this contribution to posterior predictive uncertainty is much smaller than the vertical, null space contribution to posterior predictive uncertainty. This is the case where lack of information in a history-matching dataset, rather than measurement noise associated with this dataset, is responsible for the greater part of posterior predictive uncertainty.

The contribution made to predictive error by calibration-induced parameter compensation is described by the fourth term on the right of equation (15.6.15). As this affects s_r but not s , it introduces horizontal scatter to the s vs. s_r scatterplot. Experience shows that horizontal scatter of this type can be considerable for some predictions. Where it occurs, the line of best fit through the s vs. s_r scatterplot may have a slope that is considerably less than unity. In fact for some predictions, the horizontal width of an s vs. s_r scatterplot can be much greater than its vertical width, this providing clear evidence that, for these predictions, the simple model

calibration process has the potential to induce considerable predictive error as simple model parameters compensate for simple model structural/parametric defects. Nevertheless, by using an s vs. \underline{s} scatterplot of this type, calibration-induced predictive error can be corrected. Meanwhile, prediction-pertinent information that is held by the history-matching dataset is harvested.

The last term on the right of equation (15.6.15) expresses the direct effect of model defects on predictive error. It arises purely from sensitivity of the prediction to these defects. It promotes vertical scatter within the s vs. \underline{s} scatterplot as it affects s and not \underline{s} .

Equation (15.6.15) assumes linear model behaviour. However implementation of the paired model methodology described by Doherty and Christensen (2011) requires that neither the complex nor simple model be linear. Bias correction of a prediction made by a simple calibrated model, and association of an uncertainty interval with this prediction that embraces the range of predictions which a corresponding stochastic, complex model would make if it were subjected to posterior predictive uncertainty analysis, is a concept that requires no linearity assumption.

For nonlinear models, an s vs. \underline{s} scatterplot may show curvature; furthermore, scatter may be wider in some parts of the scatterplot than in others. If a scatterplot possesses nonlinear features such as these, then inference of scatter along a vertical line projected upwards from the \underline{s} axis as in figure 15.15 may be a little more difficult than for linear models wherein scatter is more uniform. This is because scatter assessment involves implicit projection of scatter about other parts of the line of best fit along the direction of this line to the vertical \underline{s} line. Where scatter about this line of best fit is variable, projection is more difficult than where scatter is uniform. These errors can be accommodated if a modeller is conservative in his/her assessment of scatter about parts of the s vs. \underline{s} best fit line that are used in this projection.

Despite first appearances, construction of an s vs. \underline{s} scatterplot is not necessarily a numerically intensive exercise. Recall that the complex model must be run a total of N times, where N is the number of points comprising the scatterplot. The simple model must, of course, be run many more times than this, for it must be calibrated N times. However if it runs reasonably quickly then the numerical burden of conducting N calibration exercises may not be great.

16 Other Issues

16.1 Introduction

This final chapter tidies up a few loose ends.

However this is not quite the end of the book. See the epilogue. The epilogue examines, from a Bayesian perspective, the “journey of information” through an environmental model. It begins with the same simple example of an ill-posed history-matching problem that is used to illustrate singular value decomposition in chapter 8. Using some simple equations, it reinforces conclusions about decision-support model construction and design that have been drawn throughout this text. It finishes with some suggestions about when a model should undergo conventional, parameter-based history-matching, and when this should be eschewed in favour of data space inversion, or even use of a complex geostatistical prior or worst case scenario analysis as a surrogate for posterior uncertainty analysis.

16.2 Objective Function Formulation

16.2.1 General

The previous chapter of this book devotes some discussion to formulation of a sum-of-squares objective function that is minimized as a model undergoes history-matching. The Gauss-Marquardt-Levenberg method is formulated to explicitly minimize this function. Its reduction is implicit in solution of inverse problems using singular value decomposition. Its reduction is also central to formulation of equations that are used by Tikhonov regularization, and in implementation of Bayesian sampling of the posterior parameter probability distribution using an iterative ensemble smoother, data space inversion or the randomized maximum likelihood method.

The lower is the objective function, the better is the fit between model outputs and corresponding field measurements. The greater, therefore, has information that is resident in field measurements of system behaviour been absorbed by model parameters.

In the previous chapter it was recommended that consideration be given to processing field measurements and complementary model outputs in one or a number of different ways before subtracting one set from the other to form residuals; weighted squared residuals are then summed to form the overall objective function. Processed field measurements and complementary model outputs may comprise the entirety of a calibration dataset. Alternatively, unprocessed field measurements and model outputs that complement them may be retained to form one component of the total objective function. As is discussed in the previous chapter, if measurement and model output processing can be designed in such a way as to “orthogonalize out” expressions of structural noise in model outputs, this strategy can inoculate the inversion process, to some extent at least, against the bias-inducing effects of model defects on estimated model parameters, and on predictions that are sensitive to them.

Some other possibilities for creative formulation of an objective function are now discussed. We note, however, that a formulaic approach to this important matter is not possible as every data assimilation context is different. Subjectivity is therefore unavoidable. Nevertheless, a universal underlying principle should be respected. This is to ensure that as much information as possible is transferred, in as “safe” a way as possible, to predictions of management interest that are made by an environmental model.

16.2.2 Some Issues

16.2.2.1 Large Observation Range

Measurements comprising some history-matching datasets span many orders of magnitude. Flows within ephemeral streams are an example. Groundwater contaminant concentrations in the vicinity of a contaminant spill are another example. Both of these measurement datasets may contain values that range from zero to very large.

Where measurements span a large range of values, it may be inappropriate to assign equal weights to all of them. A strategy of equal weighting ensures that measurements whose values are large dominate the objective function while measurements whose values are small are “lost”. In the case of streamflow, this strategy may ensure that history-matching can match flood peaks well, but may not be able to match flood recessions, nor diminished streamflows during times of drought. In the case of groundwater contamination, a strategy of equal weighting may ensure that history-matching is able to match concentrations close to a contaminant source, but is unable to reproduce concentrations near the periphery of a contaminant plume. In both of these cases important information may be lost unless these low-valued measurements are rendered more “visible” in the overall objective function. Furthermore, predictions of management interest required of a model may pertain to times and/or locations within the model domain which are similar to those at which low-valued measurements were recorded; the information that is implied in a low value of streamflow or contaminant concentration may be very relevant to these predictions.

This situation is easily rectified. Measurements and model outputs can be transformed (for example log-transformed) before being matched to each other so that small measured values are as visible as large measured values. Alternatively, weights may be calculated from measurements in such a way as to ensure that higher weights are assigned to lower-valued measurements.

In the author’s opinion, the best guide to formulation of an appropriate weighting strategy is common sense. Common sense can be expressed in many ways.

For example, if high measured values are rich in information pertaining to one aspect of a system, and low measured values are rich in information pertaining to another aspect of a system, a modeller may decide to formulate two distinct components of a multi-component objective function. Each component employs the same measurement dataset. In the first of these components, weights may be uniform. In the second of these components, measurements whose values are low or intermediate in value may be rendered more visible in the overall objective function through logarithmic (or other) transformation. A group-specific weight multiplier can then be applied to each of these two groups of measurements in order to ensure that the contribution made to the overall objective function by each of them is roughly equal at the commencement of the history-matching process.

16.2.2.2 Non-Detects

Suppose that measurements of contaminant concentration comprise at least part of a history-matching dataset. Suppose also that some measured concentration values are recorded as “non-detect”. The non-detect threshold will normally have varied over time – a fact that needs to be taken into account where measurement datasets span many years. What must also be taken into account is that “non-detect” does not necessarily mean zero.

Measurements of low or non-detectable concentration are often an important component of a

history-matching dataset. In a groundwater remediation context these define the periphery of a contaminant plume. Hence high weights may need to be assigned to these measurements.

History-matching requires that a numerical value be assigned to “non-detect”. However caution must be exercised. If a modeller associates a value of zero with “non-detect”, then the history-matching process may alter model parameters so that model-calculated concentrations which are below the detection threshold, but are not exactly zero, are moved downwards towards zero. This may not reflect reality. Hence, declaring non-detects as zero provides misinformation to the inversion process.

A better option is to declare observed non-detects as equal to the detection limit. Meanwhile a model post-processor (or the model itself) should raise model-calculated concentrations up to the detection limit if they are below it. A zero-valued residual is thus achieved if both modelled and observed concentrations are at or below the detection limit. A non-zero residual is recorded if either is at or below this limit and the other is above it. Meanwhile continuity of model outputs with respect to parameters is preserved as modelled concentrations cross the detection limit.

16.2.2.3 Asymmetric Sensitivities

Another problem that is often associated with history-matching a groundwater model against contaminant concentrations is their nonlinearity. Suppose that a history-matching dataset includes observations of contaminant concentrations at a certain well. When endowed with a certain set of parameters, a model may calculate a contaminant concentration of zero at that well. The sensitivity of the model-calculated well concentration with respect to every model parameter is also therefore zero. An inversion process is therefore given no opportunity to adjust parameters in ways that move the plume to the well so that modelled concentrations can match observed concentrations.

There are a number of ways in which this problem can be addressed. It is the author’s experience that when undertaking history-matching, particles may be a more productive way of specifying model-calculated plume locations than concentrations. Part of a measurement dataset may therefore be comprised of “distance to plume” observations. This is easily calculated from particle trajectories using utility software supplied with the PEST suite. Where non-zero contaminant concentrations are recorded for a well, the measured value of “distance to plume” is zero. Meanwhile where, for a given set of parameters, a model-calculated plume does not intersect a well, the modelled “distance to plume” is nonzero. This distance is sensitive to local hydraulic property parameters. Adjustment of these parameters therefore allows the plume to intersect the well.

PEST utility software is able to associate an “in the plume” halo with every particle. This has a value of one along particle trajectories and diminishes continuously (using a Gaussian curve) to zero at a user-specified distance from each particle trajectory. Meanwhile, a modeller can assign each observation well an “in-the-plume” status or an “out-of-the-plume” status at times at which the well was sampled for contaminants. This status is either zero or one. By reducing discrepancies between observed values of “in-or-out-of-the-plume” at measurement wells with model-calculated “in the plume” particle halo values, a history-matching process is able to continuously adjust parameters until contaminated wells lie within a model-calculated, particle-based contaminant plume.

Hugman et al (2022) extend the use of particles to include “particle inconsistency status” in history-matching models that are built to explore contaminant source locations. A particle that

passes near a well in which contaminant has been measured cannot also pass near a well in which no contaminant has been detected. Using ancillary linear analysis, Hugman et al (2022) show that inclusion in their history-matching dataset of the “observation” that the inconsistency status of all particles is zero considerably reduces the uncertainties of history-match-inferred, near-well hydraulic conductivities.

Sensitivity asymmetry can also occur where rate of springflow comprises part of a measurement dataset through which a groundwater model is history-matched. If the head calculated by a groundwater model is above the level of a spring, then model-calculated springflow is non-zero; as such, it is sensitive to at least some model parameters. In contrast, if the model-calculated groundwater head is below the level of the spring, then the spring does not flow. Because springflow is then insensitive to all model parameters, an inversion process may experience difficulties in adjusting those model parameters which can revive flow of the spring. In this case, sensitivity asymmetry can be rectified by introducing near-spring groundwater head to the history-matching dataset. The observed value of near-spring head may be the spring elevation. If the model-calculated head at the spring is below the elevation of the spring (so that the spring does not flow), a non-zero, parameter-sensitive residual is associated with the head observation. An inversion process therefore gains the incentive to adjust parameters so that the head is raised, and the spring can flow.

16.2.2.4 Penalty Observations

There are occasions where one-sided observations, and one-sided sensitivities, may be purposefully deployed in order to allow a modeller to ensure that the outputs of a history-matching process are in accordance with qualitative understandings of a study site.

Suppose, for example, that ensemble-based history-matching is undertaken in order to sample the posterior probability distribution of parameters that belong to a groundwater model. Members of the posterior ensemble may all perpetrate good fits between measured borehole heads and complementary model-calculated heads. However some parameter fields may incur high groundwater levels over parts of the model domain where a modeller does not know exact groundwater depths, but nevertheless knows that they are not as shallow as those calculated by the model. He/she may decide to prevent this numerical occurrence by building into the objective function a set of one-sided observations that heads in all model cells that occupy the offending part of the model domain are, for example, 3 metres deep. At the same time, the modeller decrees that these observations do not contribute to the objective function if model-calculated heads are deeper than 3 metres. In contrast, the objective function increases where model-calculated heads are shallower than this. Weighting is therefore purposefully asymmetric.

Of course, weighting asymmetry leads to sensitivity asymmetry. Problems associated with sensitivity asymmetry are discussed above. However these problems are less likely to occur where penalty functions are used to enforce model output sensibility because the occurrence of zero sensitivities coincides with the occurrence of parameter fields that do not offend a modeller.

16.3 Differences

Environmental models are generally better at calculating differences than absolutes – for example horizontal, vertical and temporal differences of system states and fluxes. These differences tend to be more immune from model defects than absolute quantities calculated by models. This has repercussions for how a model should be history-matched, and how a model

should be deployed to support decision-making.

One of these repercussions is discussed in the previous chapter. This is the inclusion of observed and model-calculated differences in the objective function that is minimized as a model is history-matched. This strategy can defend parameters from the need to adopt values that compensate for a model's inherent inability to reproduce certain absolute aspects of a system's behaviour.

A second repercussion pertains to model predictions. In the author's experience, model-evaluated posterior uncertainties of predictive differences (temporal or spatial) are invariably smaller than model-evaluated posterior uncertainties of the raw model outputs that are used to form these differences. The same applies to comparative predictions, for example the extent to which system states calculated for one management plan are different from those calculated for another. It follows that, where possible, where modelling is used to support environmental management, then decision-making should be based on differences rather than absolutes.

While predictive uncertainty is quantifiable, the same does not apply to predictive bias. However repeated evidence that model-calculated differences are more credible than model-calculated absolutes suggests that model-calculated temporal, spatial and comparative differences are less susceptible to prior and posterior bias than model-calculated absolutes. Consistent predictive bias often cancels as differences are taken. This adds weight to the notion that, insofar as it is possible, model-based environmental management should rely on differences rather than absolutes.

16.4 Derivatives

Methodologies for highly parameterized inversion (whether achieved using conventional methods or through dimensional reduction) that are described in this book depend on an ability to successively linearize the action of a model on its parameters. The same applies to linear parameter/predictive uncertainty analysis and to certain types of nonlinear uncertainty analysis that are discussed herein. Implementation of all of these methods relies on the filling of a Jacobian matrix. This replaces the **Z** or **X** matrices in equations that are used extensively in this book.

Some models can calculate derivatives of their outputs with respect to their parameters through an extension of their numerical solution procedure. See, for example, the Mf6Adj adjoint solver for MODFLOW 6. Unfortunately, the need for flexibility in formulation of creative, multi-component objective functions (see above) makes model-calculated derivatives difficult to use in everyday inversion and uncertainty analysis settings. Hence derivatives with respect to parameters of model-generated counterparts to members of a history-matching dataset are generally calculated using finite parameter differences.

PEST provides many options for calculation of finite-difference derivatives. It also provides mechanisms for detecting problematic derivatives (often a consequence of poor model algorithmic design or difficulties experienced by model numerical solvers), and for taking remedial action during the inversion process when poor numerical derivatives are encountered. See PEST documentation for details. However there is a limit to the measures that PEST can take to ameliorate the effects that bad finite-difference derivatives can have on a history-matching process. Hence a modeller must take whatever steps that he/she can to enhance the integrity of finite-difference derivatives.

Calculation of finite-difference derivatives with respect to a particular parameter requires that the value of the parameter be varied incrementally, and that differences in model outputs

incurred by the incremental parameter difference be computed. Because the parameter's value is varied only incrementally, alterations to model outputs are small. The first few significant figures of many of these outputs are therefore the same for both the original and incremented parameter value; hence numerical precision of the last few significant figures is important. (Unfortunately, opportunities for these last few significant figures to suffer degradation can be increased if, following advice provided above, horizontal, vertical and/or temporal differences of system states are included in a calibration dataset; these quantities must be differenced twice in formulation of finite-difference derivatives.)

The following is a short checklist of measures that can maintain or enhance the integrity of finite-difference derivatives calculation. All of these suggestions attempt to ensure that model outputs reflect as accurately as possible the values of its parameters.

- When using a model-independent parameter estimator such as PEST which reads numbers from a model's output files, these numbers should be recorded by the model with maximum numerical precision, using scientific notation. Note that between 6 and 7 significant figures are required for representation of a single precision real number that is stored using 4 bytes. Around 16 significant figures are required for representation of a double precision real number that is stored using 8 bytes.
- If, as is often the case, the "model" with which PEST is interacting is a batch or script file comprised of multiple executable programs (for example a model and a number of pre- and postprocessors), these executable programs should transfer numbers between each other without loss of numerical precision.
- Convergence settings for model numerical solvers should be set tight.
- If possible, strategies such as adaptive time-stepping that can hasten model execution speed, but which may result in differences in numerical solution trajectories for different sets of parameters, should be avoided. (It is recognized, however, that for many models the cost of eschewing powerful solution stabilizers like adaptive time-stepping may result in unduly long execution times and/or model solver non-convergence. Under these circumstances, there is no other option but for PEST to simply accommodate errors that adaptive time stepping may introduce to finite-difference derivatives. Its tolerance for adaptive time-stepping is normally high; however this is model-specific.)

Despite the dangers that sometimes accompany finite-difference derivatives calculation, experience has shown that highly parameterised inversion and highly parameterized, calibration-constrained uncertainty analysis, when undertaken using software such as PEST that implements methodologies that are documented herein, is remarkably robust. However, experience has also shown that where the numerical performance of these methodologies does not live up to expectations, problematic calculation of finite-difference derivatives is nearly always the reason.

It should also be pointed out that ensemble methods such as IES are somewhat more immune to model output numerical granularity than methods that require the filling of a Jacobian matrix. This is because they do not require derivatives of model outputs with respect to parameters; instead, they rely on covariances (see chapter 12). However, the random parameter fields that they employ may upset delicate numerical solvers in other ways. Some parameter realizations that comprise an ensemble must sometimes be abandoned as they precipitate instability of a model's solver. The loss of too many realizations may prevent IES from attaining a good fit between model outputs and field measurements.

16.5 Conclusions

This book attempts to provide a working knowledge of the theory on which history-matching and uncertainty analysis of environmental models in general, and groundwater models in particular, is based. It is hoped that this theory is presented in a way that allows it to be properly understood and strategically implemented. Just as importantly, it is hoped that it is presented in a way that allows both modellers and decision-stakeholders to recognize the significant repercussions of this theory for the way in which decision-support modelling is practiced.

Some of these repercussions are profound. They pertain not just to how simulator history-matching and uncertainty analysis should be implemented. They pertain to the metrics by which decision-support modelling should be judged.

Decision-making requires information. It is the task of decision-support modelling to harvest and deliver information to decision-makers. Importantly, some of this information pertains to the repercussions of information insufficiency. Decision-making can therefore be risk-aware.

In many management contexts, numerical simulation is pivotal to informed decision-making. This is not because simulation can provide an accurate replication of complex subsurface processes, together with detailed representation of system properties on which these processes depend. What it can provide, however, are optimal receptacles for information that can illuminate the repercussions of these processes and properties for attainment of management goals, whatever these goals may be. In doing so, it can identify deficiencies in critical information, and suggest strategies by which these deficiencies can be rectified.

Numerical simulation should not be considered as an end in itself. Nor should decision-support simulation be judged according to metrics that imply that numerical simulation, if sufficiently complex, can provide anything that approaches digital replication of an unseen reality. This is not possible; furthermore, pursuit of this goal may erode the support that simulation can actually provide to decision-making.

A model plays host to parameters. In days gone by, the design of a parameterization scheme for a numerical model received very little consideration other than it be “geologically reasonable” and parsimonious. Few recognized that these two attributes are often mutually exclusive. Assignment of values to parameters through model calibration was seen as the last step in building a “fit for purpose” simulator that could be used to make a wide variety of decision-pertinent predictions. Uncertainty analysis was rarely undertaken, as it was inconceivable that a calibrated model could make predictions that are too far from the truth.

The view of decision support modelling that is presented in this book focusses on the centrality of information to decision-making. In doing so, it reverses the role of simulators and parameters in the decision-support modelling process. Parameters are not accessories. They are the very reason for building a numerical model, for they provide receptacles for the information that modelling must harvest. The design of a parameterization scheme, and the assignment of prior probability distributions to model parameters, is as important as the setting of layer, grid and boundary condition specifications for the simulator that hosts these parameters. In addition to this, all aspects of decision-support model construction, parameterization and history-matching must be attuned to the particular set of predictions that a model is required to make.

The primary design criterion for decision-support modelling is that it provide an unobstructed path through which information can flow from field measurements to predictions, and ultimately to decision-makers who rely on these predictions. As much of this information as possible must arrive at its destination. As little of this information as possible must be morphed

into disinformation that expresses itself as invisible predictive bias, or as uncertainty intervals that do not properly reflect information insufficiency with respect to predictions that matter.

These tasks go well beyond simulation. And they require much more than a simulation package. Use of packages such as PEST and PEST++, assisted by utility software that can give full stochastic expression to decision-salient expert knowledge, and that can formulate and implement model-to-measurement fit criteria that protect model-based decision-making from simulator and parameterization inadequacies, are fundamental to achieving their successful implementation.

It is hoped that this book provides modellers with an understanding of the theoretical basis behind PEST, PEST++ and their support packages. More than this, it is hoped that this book also provides modellers with an understanding of what numerical simulation can, and cannot, contribute to the making of critical decisions that accompany everyday environmental management.

17. References

Aanonsen, S.L., Naevdal, G., Oliver, D.S., Reynolds, A.C. and Vallès, B., 2009. The ensemble Kalman filter in reservoir engineering – a review. *Soc Petrol Eng J*, 14(3), 393-412, doi:10.2118/117274-PA.

Albert, A., 1972. Regression and the Moore-Penrose pseudoinverse. Academic Press, New York.

Alfonzo, M. and Oliver, D.S., 2019. Evaluating prior predictions of production and seismic data. *Comput. Geosci.*, 23:1331-1347.

Belsley, D.A., Kuh, E. and Welsch, R.E., 1980. Regression Diagnostics: Identifying Influential Data and Source of Collinearity. John Wiley, New York.

Beven, K., 2005. On the concept of model structural error. *Water Sci. Technol*, 52(6), 167–175.

Beven, K.J., Smith, P.J. and Freer, J., 2008. So just why would a modeller choose to be incoherent? *J Hydrol*, 345 (1), 15-32.

Blatter, D., Morzfeld, M., Ley, L. and Constable, S., 2022. Uncertainty quantification for regularized inversion of electromagnetic geophysical data. Part 1: motivation and theory. *Geophysical Journal International*, 231(2). doi:10.1093/gji/ggac241.

Bredehoeft, J., 2005. The conceptualization model problem – surprise. *Hydrogeol. J.* 13, 17-46.

Carle, S., 1999. T-PROGS: Transition Probability Geostatistical Software. University of California, Davis.

Certes, C. and deMarsily, G., 1991. Application of the pilot point method to the identification of aquifer transmissivities. *Adv. Water Resour.*, 145, 284–300, 1991.

Chada, N.K., Iglesias, M.A., Roininen, L. and Stuart, A.M., 2018. Parameterizations for ensemble Kalman inversion. *Inverse problems*, 34, 055009 (31pp).

Chen, Y. and Oliver, D.S., 2013. Levenberg-Marquardt forms of the iterative ensemble smoother for efficient history-matching and uncertainty quantification. *Comput. Geosci.* 17:689-703.

Chen, Y. and Oliver, D.S., 2017. Localization and regularization for iterative ensemble smoothers. *Comput. Geosci.* 21:13-30.

Christensen, S. and Cooley, R.L., 1999. Evaluation of prediction intervals for expressing uncertainties in groundwater flow model predictions. *Water Resour Res*, 35 (9), 2627-2639.

Collenteur, R.A., Bakker, M., Caljé, Klop, S.A. and Schaars, F., 2019. Pastas: open source software for the analysis of groundwater time series. *Groundwater*. 57 (6): 877-885.

Cook, R.D. and Weisberg, S., 1982. Residuals and Influence in Regression. *Monogr. Stat. Appl. Probability*, vol 18, Chapman and Hall, New York, 1982.

Cooley, R.L., 1983. Some new procedures for numerical solution of variably saturated flow problems. *Water Resour. Res*, 19 (5), 1271-1285.

Cooley, R.L. and Christensen, S., 2006. Bias and uncertainty in regression-calibrated models of groundwater flow in heterogeneous media. *Adv Water Resour*, 29, 639-656.

Cooley, R.L., 2004. A theory for modeling ground-water flow in heterogeneous media. *U.S. Geological Survey Professional Paper 1679*, 2004. 220 pp.

Cui, T., Fox, C. and O'Sullivan, M.J., 2011. Bayesian calibration of a large-scale geothermal reservoir model by a new adaptive delayed acceptance Metropolis Hastings algorithm. *Water Resour. Res.*, 47, W10521, doi:10.1029/2010WR010352.

Dausman, A.M., Doherty, J., Langevin, C.D., and Sukop, M.C., 2010. Quantifying data worth toward reducing predictive uncertainty. *Groundwater*, 48 (5), 729-740.

De Groot-Hedlin, C. and Constable, S., 1990. Occam's inversion to generate smooth, two-dimensional models from magnetotelluric data. *Geophysics*, 55 (12), 1613-1624.

Deutsch, C. and Journel, A.G., 1997. *GSLIB: Geostatistical Software Library and User's Guide*. Oxford University Press, New York. Second edition, 369pp.

Delottier, H., Doherty, J. and Brunner, P., 2023. Data space inversion for efficient uncertainty quantification using an integrated surface and subsurface hydrologic model. *Geoscientific Model Development*, 16(14): 4213-4231.

Doherty, J., 2003. Groundwater model calibration using pilot points and regularization. *Groundwater*, 41 (2), 170-177.

Doherty, J., 2022. The Decision-Support Modelling Workflow. GMDSI Monograph. Downloadable from <http://www.gmdsi.org>

Doherty, J., 2025a. PLPROC: A Parameter List Processor. Watermark Numerical Computing, Queensland, Australia. Downloadable from <http://www.pesthomepage.org>

Doherty, J., 2025b. PEST Groundwater Utilities. Watermark Numerical Computing, Queensland, Australia. Downloadable from <http://www.pesthomepage.org>

Doherty, J. and Christensen, S., 2011. Use of paired simple and complex models in reducing predictive bias and quantifying uncertainty. *Water Resour Res*, doi:10.1029/2011WR010763.

Doherty, J. and Hunt, R.J., 2009. Two statistics for evaluating parameter identifiability and error reduction. *J Hydrol*, 366, 119-127.

Doherty, J. and Lotti, F., 2025. Simulation and Data Assimilation for an Open Loop Shallow Geothermal System using MODFLOW 6 GWE, Non-Stationary Geostatistics and Data Space Inversion. *GMDSI Tutorial*. Downloadable from <http://www.gmdsi.org>

Doherty, J. and Moore, C., 2021. Decision-Support Modelling Viewed through the Lens of Model Complexity. GMDSI Monograph. Downloadable from <http://www.gmdsi.org>

Doherty, J., and Moore, C., 2023. Problem Decomposition in Decision-Support Groundwater Modelling. GMDSI Monograph. Downloadable from <http://www.gmdsi.org>

Doherty, J. and Simmons, C.T., 2013. Groundwater modelling in decision support: reflections on a unified conceptual framework. *Hydrogeol J*, 21, 1531–1537.

Doherty, J. and Welter, D., 2010. A short exploration of structural noise, *Water Resour. Res.*, 46, W05525, doi:10.1029/2009WR008377.

Emerick, A.A. and Reynolds, A.C., 2013. Ensemble smoother with multiple data assimilation. *Computers & Geosciences*. 55:3-15.

Evensen, G., Vossepoel, F.C. and van Leeuwen, P.J., 2022. *Data Assimilation Fundamentals. A Unified Formulation of the State and Parameter Estimation Problem*. Springer Textbooks in

Earth Sciences, Geography and Environment. <https://doi.org/10.1007/978-3-030-96709-3>

Freeze R.A., Massmann J., Smith L., Sperling T. and James B., 1990. Hydrogeological decision analysis: 1 a framework. *Groundwater* 28 (5), 738–766.

Fuentes, M., 2002. Spectral methods for nonstationary spatial processes. *Biometrika*, 89: 197-210.

Gallagher, M. R., and Doherty, J., 2007. Parameter interdependence and uncertainty induced by lumping in a hydrologic model. *Water Resour Res.*, 43, W05421, doi:10.1029/2006WR005347.

Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B., Vehtari, A. and Rubin, D.B., 2025. Bayesian Data Analysis. Third Edition. Downloadable from <https://sites.stat.columbia.edu/gelman/book/BDA3.pdf>

Hadi, A.S., 1992. A new measure of overall potential influence in linear regression. *Comput Stat Data An*, 14, 1-27.

Halko, N., Martinsson, P.-G. and Tropp, J. A., 2011. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2), 217–288.

He, J., Sarma, P., Bhark, E., Tanaka, S., Chen, B., Wen, X-H and Kamath, L., 2018. Quantifying expected uncertainty reduction and value of information using ensemble-variance analysis. *SPE Journal*. 23. 10.2118/182609-PA.

Higdon, D., Swall, J. and Kern J., 1999. Non-stationary spatial modeling, in: *J.M. Bernardo et al. (Eds.), Bayesian Statistics 6*, Oxford University Press, Oxford, pp. 761–768.

Hill, M. C., 1992. A Computer Program (MODFLOWP) for Estimating Parameters of a Transient, Three-Dimensional, Ground-Water Flow Model using Nonlinear Regression. *U. S. Geological Survey Open-File Report 91-484*.

Hill, M.C. and Tiedeman, C.R., 2007. Effective groundwater model calibration: With analysis of data, sensitivities, predictions, and uncertainty. Wiley and Sons, New York, New York, 455p.

Hugman, R., Lotti, F. and Doherty, J., 2022. Probabilistic contaminant source assessment – getting the most out of field measurements. *Groundwater* 61 (3): 363-374.

Iglesias, M.A., Law, K.J.H. and Stuart, A.M., 2013. Ensemble methods for inverse problems. *Inverse problems*, 39, 045001 (20pp).

Jiang, S., Sun, W. and Durlofsky, L., 2020. A data space inversion procedure for well control optimization and closed-loop reservoir management. *Comput. Geosci.* 24:361-379.

Kavetski, D., Fenicia, F., Reichert, P., and Albert, C., 2018. Signature-domain calibration of hydrological models using approximate Bayesian computation: theory and comparison to existing applications. *Water Resour Res.* 54, 4059–4083. <https://doi.org/10.1002/2017WR020528>

Kay, J. and King, M., 2020. Radical Uncertainty – Decision-Making Beyond Numbers. Bridge Street Press.

Kitlasten, W., Moore, C. and Doherty, J., 2025. Exploring the use of new data assimilation technologies to map groundwater quality vulnerability in a large alluvial aquifer. *Front. Earth. Sci. Sec. Hydrosphere*. vol 13 – 2025, doi: 10.3389/feart.2025.1609778

Koch, K-R. 1999. Parameter Estimation and Hypothesis-Testing in Linear Models. Springer, Berlin, Heidelberg.

Larsen, R.M., 1998. Lanczos bidiagonalization with partial reorthogonalization: Aarhus, Denmark, Aarhus University, Computer Science Department, 90 p. *PROPACK* is downloadable from <http://sun.stanford.edu/~rmunk/PROPACK/>

Levenberg, K., 1944. A method for the solution of certain non-linear problems in least squares. *Q Appl Math*, 2, 164-168.

Lima, M. M., Emerick, A. A. and Ortiz, C. E. P., 2020. Data-space inversion with ensemble smoother. *Comput. Geosci.*, 24:1179–1200, <https://doi.org/10.1007/s10596-020-09933-w>, 2020.

Liu, D., Rao, X., Zhao, H., Xu, Y-F and Gong, R-X., 2021. An improved data space inversion method to predict reservoir state fields via observed production data. *Petroleum Science* 18:1127-1142.

Luo, X., Bhakta, T. and Naevdal, G., 2018. Correlation-based adaptive localization with applications to ensemble-based 4d seismic history-matching. *SPE Journal*. 23(2) 396-427.

Manewell, N., Doherty, J. and Hayes, P., 2023a. Spatial averaging implied in aquifer test interpretation: the meaning of estimated hydraulic properties. *Front. Earth. Sci. Hydrosphere*: doi.org/10.3389/feart.2022.1079287.

Manewell, N., Doherty, J. and Hayes, P., 2023b. Translating pumping test data into groundwater model parameters: a workflow to reveal aquifer heterogeneities and implications in regional model parameterisation. *Frontiers in Water. Water and Hydrocomplexity*. doi: 10.3389/frwa.2023.1334022.

Mariethoz, G. and Caers, J., 2015. Multiple-Point Geostatistics. Stochastic Modeling with Training Images. John Wiley and Sons, UK.

Marquardt, D. W., 1963. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society of Industrial and Applied Mathematics*, 11 (2), p431-441.

Mikhail, E.M., 1976. Observations and Least Squares. IEP Series in Civil Engineering. Dun-Donnelley, New York.

Moore, C. and Doherty, J., 2005. The role of the calibration process in reducing model predictive error. *Water Resour Res*, 41, W05020, doi:10/1029/2004WR003501.

Moore, C. and Doherty, J., 2006. The cost of uniqueness in groundwater model calibration. *Adv in Water Resour*, 29, (4), 605 – 623.

Moore, C., Wöhling, T., and Doherty, J., 2010. Efficient regularization and uncertainty analysis using a global optimization methodology. *Water Resour Res*, 46, W08527, doi:10.1029/2009WR008627.

Oliver, D.S., 1995. Moving averages for Gaussian simulation in two and three dimensions. *Math. Geol.* 27 (8), 939-960.

Oliver, D.S. and Chen, Y., 2010. Recent progress on reservoir history matching: a review. *Computat Geosci*, 15 (1), 185-221.

Oliver, D.S. and Alfonzo, M., 2018. Calibration of imperfect models to biased observations. *Comput. Geosci*. 22:145:161.

Oliver, D.S., 2022. Hybrid iterative ensemble smoother for history-matching of hierarchical

models. *Math. Geosci.* 54:1289-1313.

Opazo, T., 2025. The Prior and the Likelihood: Accomodating Uncertainty Priors and Model Defects in Groundwater Modelling Decision Support. *PhD Thesis. College and Science and Engineering, Flinders University.*

Paige, C.C. and Saunders, M.A., 1982a. LSQR: An algorithm for sparse linear equations and sparse least squares. *ACM T Math Software*, 8 (1), 43–71.

Paige, C.C., and Saunders, M.A., 1982b. Algorithm 583 LSQR: Sparse linear equations and least squares problems. *ACM T Math Software*, 8 (2), 195–209.

Petropoulos, F. and many others, 2022. Forecasting: theory and practice. *International Journal of Forecasting*. 38:705:871.

Poeter, E.P. and Hill, M.C., 2007. MMA, A Computer Code for Multi-Model Analysis. *U.S. Geological Survey Techniques and Methods*, 6-E3, 113p.

RamaRao, B.S., LaVenue, A.M., De Marsily, G. and Marietta, M.G., 1995. Pilot point methodology for automated calibration of an ensemble of conditionally simulated transmissivity fields: 1. Theory and computational experiments. *Water Resour Res*, 31(3), 475-493.

Renard, P. and Allard, D., 2013. Connectivity metrics for subsurface flow and transport. *Advances in Water Resources*, 51, 168-196.

Renard, P. and Ababou, R., 2022. Equivalent permeability tensor of heterogeneous media: upscaling methods and criteria (review and analyses). *Geosciences*, 12, 269.

Remy, N., Boucher, A. and Wu, J., 2011. Applied Geostatistics with SGeMS. Cambridge University Press. 264pp.

Suescun, L.C., Cook, P., Partington, D., Hugman, R. and Wallis, I., 2025. Influence of aquifer properties on optimal location and duration of groundwater level monitoring to predict aquifer recovery after mine closure. *Journal of Hydrology*. 661. <https://www.sciencedirect.com/science/article/pii/S0022169425010248?via%3Dihub>

Sun, W. and Durlofsky, L. J., 2017. A new data-space inversion procedure for efficient uncertainty quantification in subsurface flow problems. *Math. Geosci.*, 49, 679–715, <https://doi.org/10.1007/s11004-016-9672-8>, 2017.

Tahmasebi, P., 2018. Multiple Point Statistics: A Review. In: Daya Sagar, B., Cheng, Q., Agterberg, F. (eds) *Handbook of Mathematical Geosciences*. Springer, Cham. https://doi.org/10.1007/978-3-319-78999-6_30

Tonkin, M. and Doherty, J., 2005. A hybrid regularised inversion methodology for highly parameterised models. *Water Resour Res*, 41, W10412, doi:10.1029/2005WR003995, 2005.

Vecchia, A.V. and Cooley, R.L., 1987. Simultaneous confidence and prediction intervals for nonlinear regression models with application to a groundwater flow model. *Water Resour Res*, 23 (7), 1237-1250.

Vrugt, J.A., Gupta, H.V., Bastidas, L.A., Bouten, W and Sorooshian W., 2003. Effective and efficient algorithm for multiobjective optimization of hydrologic models. *Water Resour Res*, 39(8), 1214, doi:10.1029/2002WR001746.

Vrugt, J. A., 2011. DREAM: an adaptive markov chain Monte Carlo simulation algorithm to

solve discrete, noncontinuous, posterior parameter estimation problems. *Hydrol. Earth Syst. Sc.*, 8, 4025-4052, doi:10.5194/hessd-8-4025-2011.

Vrugt, J.A., 2016. Markov chain Monte Carlo simulation using the DREAM software package: theory, concepts and MATLAB implementation. *Environmental Modelling & Software*. 75, 273-316.

Vrugt, J.A. and Beven, K.J., 2018. Embracing equifinality with efficiency: limits of acceptability sampling using the DREAM(LOA) algorithm. *Journal of Hydrology*. 559, 954-971.

Watson, T.A., Doherty, J.E. and Christensen, S., 2013. Parameter and predictive outcomes of model simplification. *Water Resour Res.*, 49 (7), 3952-3977. DOI: 10.1002/wrcr.20145.

White, J., 2018. A model-independent iterative ensemble smoother for efficient history-matching and uncertainty quantification in very high dimensions. *Environmental Modelling and Software*. 109, 191-213. DOI: 10.1016/j.envsoft.2018.06.009

White, J.T., Doherty, J.E. and Hughes, J.D., 2014. Quantifying the predictive consequences of model error with linear subspace analysis. *Water Resour Res.*, 50(2), 1152-1173. doi: 10.1002/2013WR014767.

Wöhling, T. and Vrugt, J.A., 2008. Combining multiobjective optimization and Bayesian model averaging to calibrate forecast ensembles of soil hydraulic properties. *Water Resour Res.*, 44, W12432, doi:10.1029/2008WR007154.

Woodbury, A. and Ulrych, T., 2000. A full-Bayesian approach to the groundwater inverse problem for steady state flow. *Water Resour Res.*, 36 (8), 2081-2093, doi:10.1029/2000WR900086.

Yager, R.M., 1998. Detecting influential observations in nonlinear regression modelling. *Water Resourc Res.*, 34 (7), 1623-1633.

Zhang, J., Lin, G., Li, W., Wu, L., and Zeng, L., 2018. An iterative local updating ensemble smoother for estimation and uncertainty assessment of hydrologicl model parameters with multimodel distributions. *Water Resources Research*. 54: 1716-1733.
<https://agupubs.onlinelibrary.wiley.com/doi/full/10.1002/2017WR020906>

Epilogue: The Journey of Information

E1. Introduction

This epilogue is extracted from a GMDSI monograph that was written by the author of this book. Its perspective on decision-support modelling is the same as that which is presented herein. However this perspective is presented in a slightly different way.

“GMDSI” stands for “Groundwater Modelling Decision Support Initiative”. It is an industry-funded, industry-targeted group that attempts to provide perspectives, education and examples of decision-support groundwater modelling in order to promote discussions that can improve its everyday deployment. See <https://gmdsi.org/> for details.

This epilogue begins by considering a simple inverse problem. This is the same problem that is used to demonstrate singular value decomposition in section 8.2. However this time the history-matching problem is approached from a Bayesian perspective. Conclusions that are drawn from equations that emerge from Bayesian analysis of this problem are then expanded in order to draw broader conclusions about decision-support modelling in general.

As the title of this epilogue suggests (and as is discussed in the body of this book), there is didactic benefit in viewing decision-support groundwater modelling as construction of a path along which information can travel. This creates a “modelling narrative”. Though not mathematically exact, narratives are important. In fact, narratives are often viewed as central to human decision-making; see, for example, Kay and King (2020).

In the author’s opinion, the narrative that presently underpins the commissioning and construction of decision-support groundwater models does not serve groundwater management well. It is based on the premise that complex, three-dimensional subsurface processes that operate in extremely heterogeneous media whose properties are only vaguely known can be replicated on a computer. Furthermore, the narrative suggests that, with greater time and money devoted to construction of such a model, fidelity of simulation is both achievable and inevitable.

A number of GMDSI monographs address this same issue. See, for example, Doherty and Moore (2021), Doherty(2022) and Doherty and Moore (2023). These monographs suggest that decision-support groundwater modelling is more likely to achieve its decision support potential if this potential is defined more realistically, and if design specifications of the decision-support modelling process seek justification in this definition.

It is suggested herein that the decision-support modelling narrative that places “integrity of simulation” as a metric against which groundwater models should be judged should be replaced by a narrative that gives modelling a more decision-relevant purpose. This purpose is that of reducing the uncertainties associated with predictions of management interest, particularly those that are associated with unwanted management outcomes. Uncertainty is reduced by information. Hence the task of decision-support groundwater modelling is to harvest information, and to convey information to the “pointy end” of a model at which decision-pertinent predictions are made.

The simple example that is presented in this epilogue illustrates that the journey taken by information from its original repository in site data, through a model’s parameters, to a place where it can express and reduce the range of predictive possibilities is fraught with danger. If insufficient attention is paid to these dangers, the result can be predictive bias and/or

underestimation of risks posed to management. These dangers are identified, as are decision-support modelling strategies that can reduce them. It is suggested that in some decision-support circumstances, traditional regularized or ensemble-based history-matching may introduce obstacles to the free passage of information. In these circumstances, serious consideration should be given to adoption of alternative mechanisms for information harvesting and delivery such as data space inversion.

Before describing the simple model that is the focus of the first part of this epilogue, the aspirations of decision-support groundwater modelling are stated. They are as follows.

- Quantification of the uncertainties of decision-critical model predictions. Decision-makers are therefore made aware of the risks that they face in managing, or interfering with, a groundwater system.
- Reduction of these uncertainties to the extent that available information allows, or to a point where the hypothesis of an unwanted management outcome can be rejected.

Doherty and Simmons (2013) point out that a decision-support modelling exercise *fails* where the uncertainties of decision-critical model predictions are understated. On the other hand, decision-support modelling runs the risk of being *useless* if predictive uncertainties are carelessly overstated in order to circumvent decision-support modelling failure. Uselessness is, in fact, an outcome of failure to assimilate prediction-salient information that is available to the modelling process.

E2. A Model

E2.1 Model Equations

Figure E2.1 shows a simple model. Water flows into the right boundary of a one-dimensional confined aquifer. The flow rate is 1.0. Conditions are steady-state. A Dirichlet boundary occupies the left end of the model domain. The head in this boundary is set to 0.0.

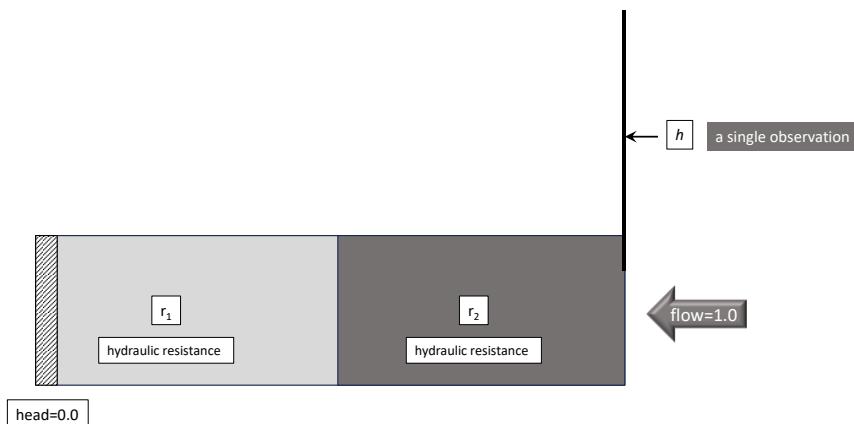


Figure E2.1 A simple model.

Groundwater flows through two materials on its way to the fixed head boundary that is situated at the left of the model domain. One of these materials presents a resistance of r_1 to groundwater flow, while the other presents a resistance of r_2 .

The model calibration dataset is comprised of a single head measurement taken at the right end of the system. We denote its value as h . Obviously, history-matching of this simple model poses an ill-posed inverse problem.

As is done in the body of this book, the following equation is used to characterise flow in the groundwater system.

$$\mathbf{h} = \mathbf{Zk} + \boldsymbol{\varepsilon} \quad (\text{E2.1.1})$$

In equation (E2.1.1), elements of the vector \mathbf{h} comprise a history-matching dataset while elements of the vector \mathbf{k} comprise model parameters. The $\boldsymbol{\varepsilon}$ vector hosts noise associated with measurements \mathbf{h} . Equation (E2.1.1) states that what is measured in the field is an outcome of the action of a model (linearised using the matrix \mathbf{Z}) on its parameters plus measurement noise.

The system that is depicted in figure E2.1 requires no linearisation as it is already linear. For this system, equation (E2.1.1) becomes equation (E2.1.2):

$$[h] = [1 \ 1] \begin{bmatrix} r_1 \\ r_2 \end{bmatrix} + [\varepsilon] = r_1 + r_2 + \varepsilon \quad (\text{E2.1.2})$$

E2.2 Model Predictions

The model is required to make two predictions. The action of a model in making a prediction is described using the following equation:

$$s = \mathbf{y}^t \mathbf{k} \quad (\text{E2.2.1})$$

In equation (E2.2.1), s is a scalar. Elements of the vector \mathbf{y} denote sensitivities of the prediction s to individual model parameters.

Figure E2.2 schematises the first prediction that is considered in this epilogue. This prediction is the head in a well that lies at the right end of material #1. We denote this prediction as s_1 . Using vector notation, its value can be calculated using the following linear equation. This equation states that the value of the prediction is the same as the value of r_1 .

$$[s_1] = [1 \ 0] \begin{bmatrix} r_1 \\ r_2 \end{bmatrix} = [r_1] \quad (\text{E2.2.2})$$

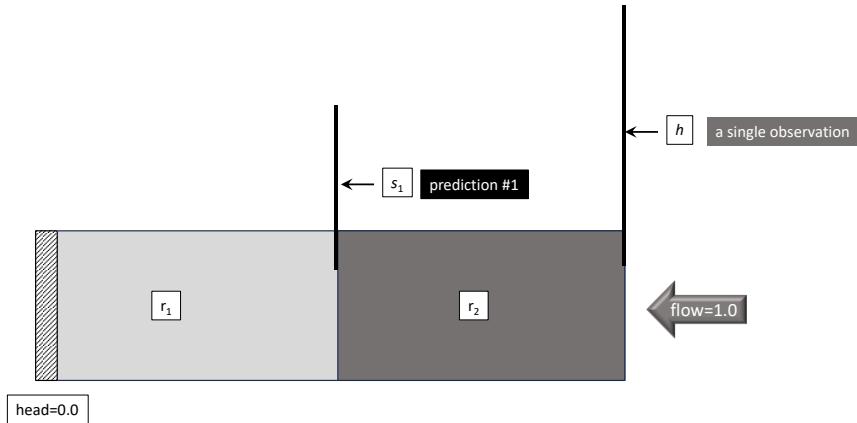


Figure E2.2. Model prediction #1.

A second model prediction is schematised in figure E2.3. This is the head at the extreme right of the system under a regime in which groundwater flow into the system is doubled. This prediction can be made using equation (E2.2.3).

$$[s_2] = [2 \ 2] \begin{bmatrix} r_1 \\ r_2 \end{bmatrix} = 2r_1 + 2r_2 \quad (\text{E2.2.3})$$

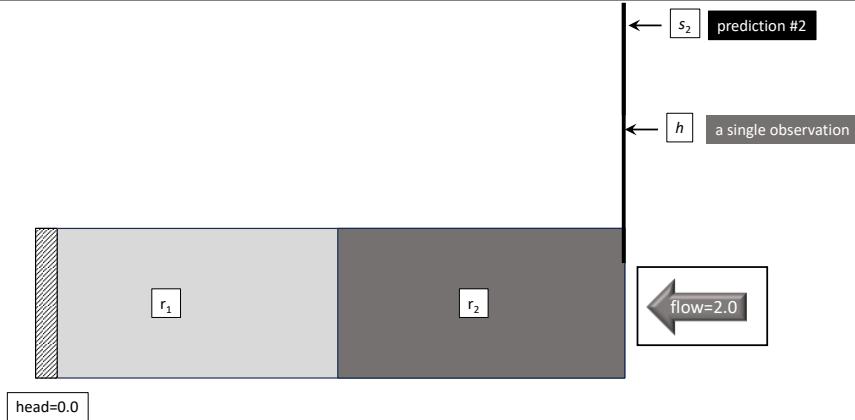


Figure E2.3. Model prediction #2.

Intuition informs us that predictions #1 and #2 will be accompanied by very different levels of uncertainty. Distinguishing between these fundamentally different types of prediction is fundamental to decision-support modelling design.

E2.3 Prior Parameter Probability Distribution

Whether history-matching is undertaken through regularised inversion (i.e. model calibration) or in a Bayesian framework using parameter ensembles, it must begin with definition of prior parameter means and uncertainties. The latter are often encapsulated in a covariance matrix. If undertaking calibration, the prior covariance matrix is used for regularization. If undertaking ensemble-based history-matching, the prior covariance matrix is used for generation of random parameter realisations.

We use the vector \underline{r} to host prior parameter means. The elements of this vector are the means of the individual resistance parameters. That is:

$$\underline{r} = \begin{bmatrix} r_1 \\ r_2 \end{bmatrix} \quad (\text{E2.3.1})$$

Meanwhile, $C(\underline{r})$ denotes the prior covariance matrix of the two resistance parameters. The elements of this matrix are depicted in equation (E2.3.2). The diagonal elements of $C(\underline{r})$ denote prior parameter variances while off-diagonal elements denote prior parameter covariances. Recall that variance is the square of standard deviation. Recall also that a covariance matrix is symmetrical; hence σ_{12} and σ_{21} have the same values.

$$C(\underline{r}) = \begin{bmatrix} \sigma_{11}^2 & \sigma_{12} \\ \sigma_{21} & \sigma_{22}^2 \end{bmatrix} \quad (\text{E2.3.2})$$

It must never be forgotten that the prior covariance of model parameters is something that is ascribed to those parameters by a modeller. Ideally, this has its roots in expert knowledge emerging from site characterisation. There is little doubt, however, that filling of a prior parameter covariance matrix requires a high degree of personal judgement. Furthermore, the importance of filling this matrix judiciously often goes unacknowledged. This is a matter to which we shall return shortly.

E3. History-Matching the Model

E3.1 Conditioning

Because the action of the model on its parameters under both calibration and predictive

conditions is linear, implementation of Bayes equation is simple. Recall that Bayes equation describes how the uncertainties of model parameters, and predictions that are sensitive to them, are reduced through history-matching. When history-matching follows the precepts of Bayes equation, it is sometimes referred to as “conditioning”. Use of this term makes it clear that uncertainty is reduced, but not eliminated, through history-matching.

In what follows, simple matrix equations are used to calculate posterior mean values of both parameters and predictions, as well as their posterior uncertainties. Calculation of posterior mean parameter values is equivalent to model calibration. In more complex modelling circumstances, calibration can be implemented using Tikhonov regularization. So-called “Bayesian Tikhonov” regularization ensures minimum departure of estimated parameter values from prior mean parameter values. The prior parameter covariance matrix governs the manner in which these departures are allowed to occur. See section 8.3 of this book.

For the present analysis, we use conditioning equations that are introduced in section 3.8. These apply where random vectors possess multiGaussian probability distributions. For convenience, they are repeated here.

Consider a general vector \mathbf{z} that is comprised of two sub-vectors \mathbf{x} and \mathbf{y} . That is

$$\mathbf{z} = \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \quad (\text{E3.1.1})$$

Let the vector $\underline{\mathbf{z}}$ denote the prior mean of \mathbf{z} . That is:

$$\underline{\mathbf{z}} = \begin{bmatrix} \underline{\mathbf{x}} \\ \underline{\mathbf{y}} \end{bmatrix} \quad (\text{E3.1.2})$$

We denote the covariance matrix of \mathbf{z} as $\mathbf{C}(\mathbf{z})$. It can be partitioned into submatrices that express the individual covariances of \mathbf{x} and \mathbf{y} on the one hand, and cross-covariances between them on the other hand.

$$\mathbf{C}(\mathbf{z}) = \begin{bmatrix} \mathbf{C}_{\mathbf{xx}} & \mathbf{C}_{\mathbf{xy}} \\ \mathbf{C}_{\mathbf{yx}} & \mathbf{C}_{\mathbf{yy}} \end{bmatrix} \quad (\text{E3.1.3})$$

Suppose that we acquire knowledge of all elements of the vector \mathbf{y} . Let us assign measured values of these elements to the vector \mathbf{y}_c (where the “c” subscript denotes “conditioning”). The conditional mean (i.e. the posterior mean) of \mathbf{x} can then be calculated using the following equation.

$$\underline{\mathbf{x}}' = \underline{\mathbf{x}} + \mathbf{C}_{\mathbf{xy}} \mathbf{C}_{\mathbf{yy}}^{-1} (\mathbf{y}_c - \underline{\mathbf{y}}) \quad (\text{E3.1.4})$$

Meanwhile, the posterior covariance matrix of \mathbf{x} is calculated as follows.

$$\mathbf{C}'_{\mathbf{xx}} = \mathbf{C}_{\mathbf{xx}} - \mathbf{C}_{\mathbf{xy}} \mathbf{C}_{\mathbf{yy}}^{-1} \mathbf{C}_{\mathbf{yx}} \quad (\text{E3.1.5})$$

Equation (E3.1.5) makes it clear that conditioning has the potential to reduce the uncertainties of elements of \mathbf{x} from their prior uncertainty values.

For convenience (and so that this epilogue is self-contained), another much-used formula that is presented in the body of this book is repeated; see section 3.9. Suppose that \mathbf{x} is a random vector, and that \mathbf{y} is calculated from \mathbf{x} using the following matrix equation. \mathbf{A} is a general matrix.

$$\mathbf{y} = \mathbf{Ax} \quad (\text{E3.1.6})$$

Because \mathbf{x} is a random vector, \mathbf{y} is a random vector. Its mean and covariance matrix are

calculated from \underline{x} (the mean of \mathbf{x}) and $C(\mathbf{x})$ (the covariance matrix of \mathbf{x}) using the following formulas.

$$\underline{\mathbf{y}} = \mathbf{Ax} \quad (\text{E3.1.7})$$

$$C(\mathbf{y}) = \mathbf{AC}(\mathbf{x})\mathbf{A}^t \quad (\text{E3.1.8})$$

E3.2 Prediction #1

E3.2.1 Model Calibration

The equations for steady-state flow of groundwater through the system that is described in the previous section must be formulated in a manner that affords application of the conditioning equations presented above. From (E2.1.2) and the obvious fact that

$$\mathbf{r} = \mathbf{Ir} \quad (\text{E3.2.1})$$

where \mathbf{I} is the identity matrix, we obtain

$$\begin{bmatrix} r_1 \\ r_2 \\ h \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \\ \varepsilon \end{bmatrix} \quad (\text{E3.2.2})$$

Using (E3.1.8) we can now obtain the covariance matrix of $\begin{bmatrix} r_1 \\ r_2 \\ h \end{bmatrix}$ as

$$C\left(\begin{bmatrix} r_1 \\ r_2 \\ h \end{bmatrix}\right) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} C\left(\begin{bmatrix} r_1 \\ r_2 \\ \varepsilon \end{bmatrix}\right) \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix}^t \quad (\text{E3.2.3})$$

Let σ_ε^2 denote the variance of error associated with measurement of h . Then, from (E2.3.2) we obtain (after transposing the last matrix)

$$C\left(\begin{bmatrix} r_1 \\ r_2 \\ h \end{bmatrix}\right) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} \sigma_{11}^2 & \sigma_{12} & 0 \\ \sigma_{21} & \sigma_{22}^2 & 0 \\ 0 & 0 & \sigma_\varepsilon^2 \end{bmatrix} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \quad (\text{E3.2.4})$$

After expansion through matrix multiplication, equation (E3.2.4) becomes

$$C\left(\begin{bmatrix} r_1 \\ r_2 \\ h \end{bmatrix}\right) = \begin{bmatrix} \sigma_{11}^2 & \sigma_{12} & \sigma_{11}^2 + \sigma_{12} \\ \sigma_{21} & \sigma_{22}^2 & \sigma_{22}^2 + \sigma_{12} \\ \sigma_{11}^2 + \sigma_{12} & \sigma_{22}^2 + \sigma_{12} & \sigma_{11}^2 + 2\sigma_{12} + \sigma_{22}^2 + \sigma_\varepsilon^2 \end{bmatrix} \quad (\text{E3.2.5})$$

If desired, the equation for element (3,3) of the matrix on the right of equation (E3.2.5) can be slightly reformulated by noting that

$$\sigma_{r_1+r_2}^2 = \sigma_{11}^2 + 2\sigma_{12} + \sigma_{22}^2 \quad (\text{E3.2.6})$$

where $\sigma_{r_1+r_2}^2$ is the prior variance of $(r_1 + r_2)$. This is the total resistance between the head measurement point on the right of the system and the fixed head boundary on the left of the system.

Now that we have formulated a joint covariance matrix between r_1 , r_2 and h , we can use equation (E3.1.4) to obtain the posterior means of r_1 and r_2 as follows.

$$\begin{bmatrix} \underline{r}'_1 \\ \underline{r}'_2 \end{bmatrix} = \begin{bmatrix} \underline{r}_1 \\ \underline{r}_2 \end{bmatrix} + \frac{\begin{bmatrix} \sigma_{11}^2 + \sigma_{12}^2 \\ \sigma_{22}^2 + \sigma_{12}^2 \end{bmatrix} (\underline{h} - \underline{h})}{\sigma_{11}^2 + 2\sigma_{12}^2 + \sigma_{22}^2 + \sigma_{\varepsilon}^2} \quad (\text{E3.2.7})$$

To make the following equations a little less cumbersome, we define the perturbation of a parameter from its prior mean using the symbol p . Therefore:

$$p_1 = r_1 - \underline{r}_1 \quad (\text{E3.2.8a})$$

$$p_2 = r_2 - \underline{r}_2 \quad (\text{E3.2.8b})$$

Note also that from (E2.1.2):

$$h = r_1 + r_2 + \varepsilon \quad (\text{E3.2.9a})$$

while

$$\underline{h} = \underline{r}_1 + \underline{r}_2 \quad (\text{E3.2.9b})$$

so that (E3.2.7) can be written as the following two equations.

$$\underline{p}'_1 = \frac{\sigma_{11}^2 + \sigma_{12}^2}{\sigma_{11}^2 + 2\sigma_{12}^2 + \sigma_{22}^2 + \sigma_{\varepsilon}^2} [p_1 + p_2 + \varepsilon] \quad (\text{E3.2.10a})$$

$$\underline{p}'_2 = \frac{\sigma_{22}^2 + \sigma_{12}^2}{\sigma_{11}^2 + 2\sigma_{12}^2 + \sigma_{22}^2 + \sigma_{\varepsilon}^2} [p_1 + p_2 + \varepsilon] \quad (\text{E3.2.10b})$$

The first of these equations is of most interest to us. Define q as the perturbation of a prediction from its prior mean, so that:

$$q_1 = s_1 - \underline{s}_1 \quad (\text{E3.2.11a})$$

$$q_2 = s_2 - \underline{s}_2 \quad (\text{E3.2.11b})$$

By construction:

$$\underline{q}'_1 = \underline{p}'_1 \quad (\text{E3.2.12})$$

so that, from (E3.2.10a):

$$\underline{q}'_1 = \underline{p}'_1 = \frac{\sigma_{11}^2 + \sigma_{12}^2}{\sigma_{11}^2 + 2\sigma_{12}^2 + \sigma_{22}^2 + \sigma_{\varepsilon}^2} [p_1 + p_2 + \varepsilon] \quad (\text{E3.2.13})$$

Equation (E3.2.13) provides some useful insights.

The first of these insights is actually rather obvious. If measurement noise is high (and hence σ_{ε}^2 is high), then the calibrated value of r_1 is equal to its prior mean (because \underline{p}'_1 approaches zero). The same applies to the posterior mean value of s_1 . In this case, the information content of the history-matching dataset is lost because of its accompaniment by a large amount of measurement noise.

A second insight is that the calibrated value of r_1 , and hence prediction s_1 made by the calibrated model, is influenced by the hydraulic properties of all media through which water flows. So while the prior mean value of s_1 depends only on the prior mean value ascribed to r_1 (the resistance of the material which separates the prediction from the left fixed head boundary), its posterior mean value depends on both r_1 and r_2 .

Another important insight from equation (E3.2.13) is that the calculated value of \underline{p}'_1 (and hence \underline{q}'_1) is also affected by the prior parameter probability distribution. Hence, not only is it affected by all media through which water flows. It is also affected by all of the parameters through which information flows. Recall that in the present example, information is harvested from the

single head measurement h which comprises the entirety of the history-matching dataset.

If the prior uncertainty ascribed to r_2 is large, then the denominator of the right side of equation (E3.2.13) is large. The posterior mean value of r_1 is therefore equal to its prior mean value. All of the head drop between the right and left of the model domain that is not explained by mean values of r_1 and r_2 is therefore attributed to perturbation of r_2 from its mean under these circumstances. In other words, r_2 absorbs all of the information that is contained in the single head measurement h .

If the prior uncertainty ascribed to r_2 is small, then both σ_{22}^2 and σ_{12} are small. If measurement noise is also small, then the calibrated perturbation of r_1 from its prior mean value is equal to the summed perturbations of both r_1 and r_2 from their prior mean values as the numerator and denominator on the right side of equation (E3.2.13) cancel. Neither r_1 nor r_2 is known, of course. However the calibration process assigns all potential resistance variability required for model calibration to r_1 instead of partitioning it between r_1 and r_2 . In other words, all of the information that resides in the single head measurement is absorbed by r_1 .

Suppose, for a moment, that material upstream of the s_1 predictive well is occupied in part or in whole by a model boundary condition, so that r_2 partly or wholly pertains to this boundary. Let us also suppose that this boundary condition conductance is in error, and that uncertainty in its parameterisation is ignored (as is often the case for boundary conditions); its uncertainty is therefore presumed to be small or zero. The above discussion implies that r_1 can partly or wholly compensate for this situation during model calibration, as its value can be adjusted in order to accommodate possibly erroneous upstream conditions that are associated with this boundary.

A final point that emerges from an inspection of equation (E3.2.13) is that the value assigned to r_1 (and hence s_1) through model calibration (which is also its posterior mean), is influenced by any prior covariance that is thought to exist between r_1 and r_2 ; that is, it is influenced by σ_{12} . This matter is further discussed below.

E3.2.2 Predictive Uncertainty

If equation (E3.1.5) is now applied to equation (E3.2.5), we obtain the following expression for the posterior covariance matrix of the two parameters r_1 and r_2 of our simple model.

$$C'(\mathbf{r}) = C(\mathbf{r}) - \frac{\begin{bmatrix} \sigma_{11}^2 + \sigma_{12}^2 & \sigma_{11}^2 + \sigma_{12}^2 & \sigma_{22}^2 + \sigma_{12}^2 \\ \sigma_{22}^2 + \sigma_{12}^2 & \sigma_{11}^2 + 2\sigma_{12}^2 + \sigma_{22}^2 + \sigma_{\epsilon}^2 \end{bmatrix}}{\sigma_{11}^2 + 2\sigma_{12}^2 + \sigma_{22}^2 + \sigma_{\epsilon}^2} \quad (\text{E3.2.14})$$

By multiplying out, we find that

$$C'(\mathbf{r}) = C(\mathbf{r}) - \frac{\begin{bmatrix} (\sigma_{11}^2 + \sigma_{12}^2)(\sigma_{11}^2 + \sigma_{12}^2) & (\sigma_{11}^2 + \sigma_{12}^2)(\sigma_{22}^2 + \sigma_{12}^2) \\ (\sigma_{11}^2 + \sigma_{12}^2)(\sigma_{22}^2 + \sigma_{12}^2) & (\sigma_{22}^2 + \sigma_{12}^2)(\sigma_{22}^2 + \sigma_{12}^2) \end{bmatrix}}{\sigma_{11}^2 + 2\sigma_{12}^2 + \sigma_{22}^2 + \sigma_{\epsilon}^2} \quad (\text{E3.2.15})$$

so that

$$\sigma'_{11}^2 = \sigma_{11}^2 - \frac{(\sigma_{11}^2 + \sigma_{12}^2)(\sigma_{11}^2 + \sigma_{12}^2)}{\sigma_{11}^2 + 2\sigma_{12}^2 + \sigma_{22}^2 + \sigma_{\epsilon}^2} \quad (\text{E3.2.16})$$

Equation (E3.2.16) demonstrates the potential for history-matching to reduce uncertainty. Recall that σ'_{11}^2 is the posterior uncertainty variance of r_1 , and hence of prediction s_1 , while σ_{11}^2 is the prior uncertainty variance of r_1 , and hence of s_1 .

The following points can be gleaned from equation (E3.2.16).

If measurement noise is large (i.e. if σ_ϵ^2 is large), then the denominator of the second term on the right of (E3.2.16) is large. The posterior uncertainty of prediction s_1 is therefore reduced very little from its prior uncertainty.

The same occurs if the prior uncertainty of r_2 is large. In this case, as was stated above, r_2 absorbs all of the information that is contained in the single head measurement.

Conversely, if measurement noise is small and the prior uncertainty of r_2 is small (i.e. if both σ_{22}^2 and σ_{12} are small), then the posterior uncertainty of r_1 (and hence of prediction s_1) is also small, regardless of its prior uncertainty. This is a matter of concern, especially if r_2 pertains to an erroneously assigned boundary condition (as was discussed above). It seems that uncertainty attracts information. Conversely, lack of uncertainty repels, or cancels, information.

A final point that can be gleaned from equation (E3.2.16) is that if r_1 and r_2 are thought to exhibit a high degree of prior spatial correlation, this reduces the posterior uncertainty of r_1 and hence of s_1 . In the extreme case that the correlation coefficient between r_1 and r_2 approaches unity and the prior uncertainties of r_1 and r_2 are the same, the second term on the right of equation (E3.2.16) approaches σ_{11}^2 (if measurement noise is small), so that the posterior uncertainty of r_1 (and hence of s_1) approaches zero.

E3.2.3 Journey of Information

Insights that are outlined above can be summarized in a way that enhances the “information as a journey” narrative that this epilogue is attempting to develop.

1. The ability of information to reach its target (in this case prediction s_1) is diminished if it encounters too much uncertainty along its journey to that target.
2. On the other hand, failure to recognise the existence of uncertainty causes information to inform the wrong variable. It thereby becomes misinformation.
3. Assignment by a modeler of spatial correlation between parameters comprises a form of information itself. This information can reduce the uncertainties of decision-critical model predictions. It must be noted however, that spatial hydraulic property correlation is a difficult stochastic specification to characterise, especially in heterogeneous materials that comprise the subsurface through which groundwater flows. These materials are often permeated by narrow zones of highly-connected permeability that afford no simple geostatistical description.

From the above it follows that information’s journey through a decision-support modelling process that is designed to harvest it, and deliver it to places where it must illuminate decisions, is fraught with danger. If information encounters unexpected uncertainty along its way, it may get lost. If it is denied access to uncertainty that actually exists, then it lands in the wrong place and misinforms the wrong parameter. On the other hand, if it is complemented by modeler-supplied information on likely subsurface hydraulic property continuity, then its utility may be enhanced.

It follows that the path along which information must travel must be carefully prepared by a decision-support groundwater modeler.

E3.3 Prediction #2

E3.3.1 Model Calibration

Equation (E3.2.10) provides posterior mean values of calibration-induced deviations of r_1 and

r_2 from their prior means. Let us define \underline{q}'_2 as the posterior deviation of prediction s_2 from its prior mean. A value can be assigned to \underline{q}'_2 by summing equations (E3.2.10a) and (E3.2.10b) and multiplying by 2. The factor of 2 is necessary because inflow into the system is doubled when making prediction s_2 . Hence

$$\underline{q}'_2 = 2 \left(\underline{p}'_1 + \underline{p}'_2 \right) = 2 \frac{\sigma_{11}^2 + 2\sigma_{12} + \sigma_{22}^2}{\sigma_{11}^2 + 2\sigma_{12} + \sigma_{22}^2 + \sigma_\varepsilon^2} [p_1 + p_2 + \varepsilon] \quad (\text{E3.3.1})$$

After dividing the numerator and denominator of the right side of the above equation by $\sigma_{11}^2 + 2\sigma_{12} + \sigma_{22}^2$ and recalling that the prior variance of total system resistance is given by:

$$\sigma_{r_1+r_2}^2 = \sigma_{11}^2 + 2\sigma_{12} + \sigma_{22}^2 \quad (\text{E3.3.2})$$

we obtain the following equation for \underline{q}'_2 :

$$\underline{q}'_2 = 2 \frac{\frac{p_1 + p_2 + \varepsilon}{\sigma_\varepsilon^2}}{1 + \frac{\sigma_\varepsilon^2}{\sigma_{r_1+r_2}^2}} \quad (\text{E3.3.3})$$

From equation (E3.3.3) it is apparent that if measurement noise is large (and hence the denominator of the right side of this equation is large), then the post-history-matching value of prediction s_2 is equal to its prior mean. This results from dilution of information contained in the single head measurement by noise associated with this measurement.

On the other hand if measurement noise is small, then the posterior mean of prediction s_2 is equal to the sum of system resistances times the inflow rate of water. This, of course, is its true value. Accuracy of this particular prediction is achieved despite the fact that it is a function of two resistances that are not uniquely resolved by history-matching!

The path along which information must travel to make this prediction is therefore partitioned into two separate routes. However, once this information coalesces again for the making of the prediction, its integrity is undiminished. (It is natural to ask whether the model is therefore unnecessarily complicated for the making of this prediction.)

E3.3.2 Predictive Uncertainty

In order to derive a formula for the posterior uncertainty of prediction s_2 , we first form a covariance matrix to which we can apply the same conditioning equations as those that were applied above.

Equations (E2.2.3) and (E2.1.2) can be combined into a single matrix equation.

$$\begin{bmatrix} s_2 \\ h \end{bmatrix} = \begin{bmatrix} 2 & 2 & 0 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \\ \varepsilon \end{bmatrix} \quad (\text{E3.3.4})$$

Using (E3.1.8) which describes propagation of variance, a joint covariance matrix between s_2 and h can be established as follows.

$$C \left(\begin{bmatrix} s_2 \\ h \end{bmatrix} \right) = \begin{bmatrix} 2 & 2 & 0 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} \sigma_{11}^2 & \sigma_{12} & 0 \\ \sigma_{21} & \sigma_{22}^2 & 0 \\ 0 & 0 & \sigma_\varepsilon^2 \end{bmatrix} \begin{bmatrix} 2 & 1 \\ 2 & 1 \\ 0 & 1 \end{bmatrix} \quad (\text{E3.3.5})$$

By multiplying out, collecting terms, and then using equation (E3.3.2) we obtain

$$C \left(\begin{bmatrix} s_2 \\ h \end{bmatrix} \right) = \begin{bmatrix} 4\sigma_{r_1+r_2}^2 & 2\sigma_{r_1+r_2}^2 \\ 2\sigma_{r_1+r_2}^2 & \sigma_{r_1+r_2}^2 + \sigma_\varepsilon^2 \end{bmatrix} \quad (\text{E3.3.6})$$

Application of (E3.1.5) then yields

$$\sigma_{s_2}^{\prime 2} = 4\sigma_{r_1+r_2}^2 - \frac{4\sigma_{r_1+r_2}^2\sigma_{r_1+r_2}^2}{\sigma_{r_1+r_2}^2 + \sigma_e^2} \quad (\text{E3.3.7})$$

Conclusions drawn from equation (E3.3.7) reinforce those drawn from equation (E3.3.3) about the journey taken by information from the single head measurement to the single prediction.

The posterior uncertainty variance of prediction s_2 is undiminished from its prior uncertainty variance of $4\sigma_{r_1+r_2}^2$ if measurement noise is large, for information is drowned in this noise. Conversely, if measurement noise is zero, the second term on the right of equation (E3.3.7) becomes equal to the first term. The posterior uncertainty of prediction s_2 is therefore zero. This pleasing outcome is obtained in spite of the “complexity” of the simulated system, and in spite of the fact that information reaches s_2 after being split between two parameters that may have relatively high posterior uncertainties before it is re-combined in order to render the s_2 prediction relatively uncertainty-free.

E3.3.3 Journey of Information

Conclusions that emerge from an examination of prediction s_2 that are pertinent to the development of the “journey of information” narrative that is the subject matter of this epilogue are now outlined.

It is apparent that some predictions of future system states and/or fluxes required of a decision-support groundwater model may be well informed by historical measurements of system states and fluxes. These predictions may be sensitive to parameters whose posterior uncertainties are high, or whose values have been corrupted during the history-matching process as they compensate for model inadequacies or for errors/inconsistencies in assignment of prior parameter probability distributions. The integrity of these predictions is not adversely affected by the degraded integrity of parameters to which they are sensitive. Nor are their posterior uncertainties increased by the high posterior uncertainties of these same parameters.

E4. Reflections on the Journey

E4.1 General

As was stated above, the purpose of this epilogue is to develop a narrative that assists modellers and modelling stakeholders to make the many choices that decision-support modelling requires of them. The reference point for these choices should be the decision support modelling goals that are outlined in section E1. These goals are not “integrity of simulation”. Instead, they are harvesting and delivery of information.

The above-documented solution of a very simple ill-posed inverse problem (in fact, as simple as an ill-posed problem can be while still being an ill-posed problem) provides useful insights into the journey of information that is enabled by more complex decision-support modelling. In doing so, it offers insights into decision support modelling appropriateness. We now reflect on these insights, with the help of some schematics.

E4.2 Information Splitting

Figures E4.1 and E4.2 schematize decision-support modelling when viewed from a “journey of information” perspective.

Figure E4.1 pertains to a prediction that is fully informed by a model’s history-matching dataset

(such as prediction s_2 of the model described in previous sections). Such a prediction is denoted as “solution-space-dependent” or “data-driven” in the body of this book. That is to say, it is entirely dependent on combinations of parameters that lie entirely within a subspace of parameter space that is spanned by a subset of orthogonal unit vectors obtained from singular value decomposition of the model Z matrix. These unit vectors lie to the left of the point at which singular values are set to zero because of lack of information, or because of increasing dilution of information by amplification of measurement noise. See section 8.2 of this book for details. In this epilogue, we refer to the prediction that is illustrated in figure E4.1 as a “type 2 prediction” because it is exemplified by prediction #2 in the above discussion.

In contrast, figure E4.2 pertains to a prediction that is partly solution-space-dependent and partly null-space-dependent (such as prediction s_1 of the previous example). Its uncertainty is reduced through history-matching. However, its partial sensitivity to combinations of parameters that span the null space of the history-matching inverse problem guarantees that uncertainty reduction is incomplete because of a deficit of information with respect to this prediction in the history-matching dataset. We refer to the prediction that is illustrated in figure E4.2 as a “type 1 prediction” in the discussion that follows.

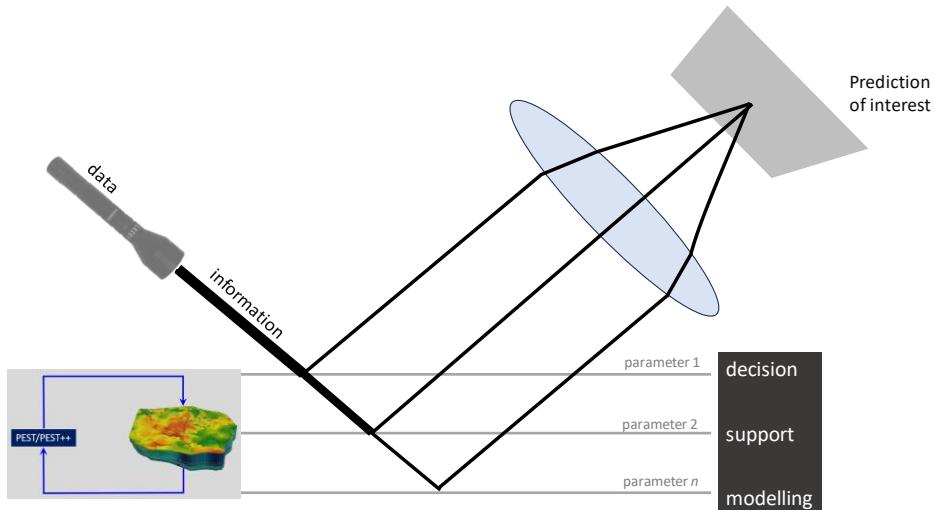


Figure E4.1. Journey of information to a solution-space-dependent prediction (i.e. a type 2 model prediction).

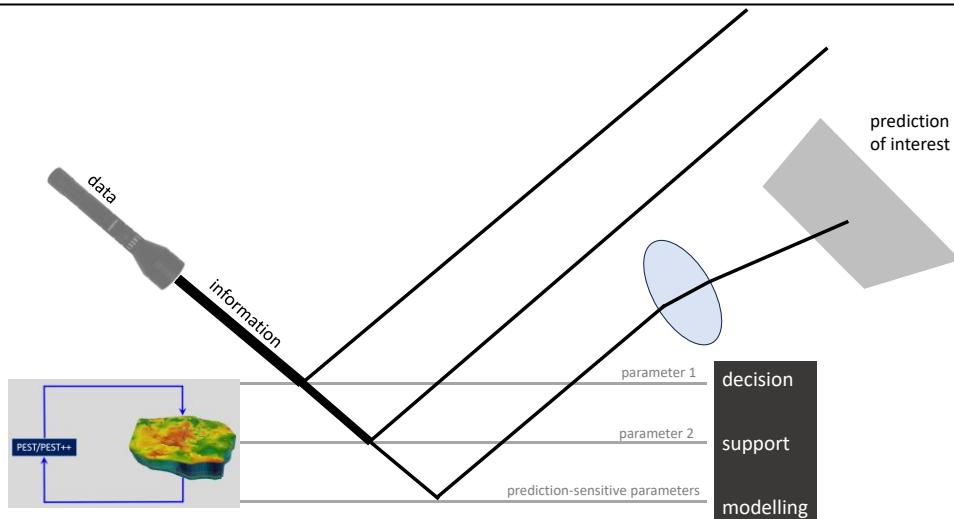


Figure E4.2. Journey of information to a prediction which is partially solution-space-dependent and partially null-space-dependent (i.e. a type 1 model prediction).

In both of figures E4.1 and E4.2, decision-support modelling is represented as a process of information harvesting and delivery. Information that is harvested through history-matching is stored in a model's parameters, for this is what parameters are designed to do. A parameter is (by definition) adjustable; it can therefore express both the presence and absence of information. For many parameters, post-history-matching adjustability is less than pre-history-matching adjustability as their prior adjustability is constrained by history-match-harvested information.

Most history-matching datasets host information that is stored by more than one model parameter after being harvested through history-matching. This is especially the case where a model is endowed with a multiplicity of parameters. As explained in the body of this book, highly parameterised history-matching is recommended decision-support groundwater modelling practice for reasons that include the following.

1. It avoids the possibility of information rejection incurred by failure to attain a high level of model-to-measurement fit.
2. It reduces (but does not eliminate) the possibility of predictive bias incurred by storage of model-harvested information in inappropriate parameter receptacles.
3. It avoids underestimation of predictive uncertainty through failure to represent parameters that are either uninformed by a measurement dataset, or are informed to only a limited extent.

All of these reasons support decision-support modelling goals that are outlined in section E1.

E4.3 Prior Parameter Probability Distribution

While the introduction of many parameters to a decision-support model is an essential ingredient of decision-support modelling practice in many groundwater management circumstances, it is accompanied by some important responsibilities. These include the assignment of a joint prior probability distribution to these parameters. The more complex is the system that is being simulated, and the more complex is the parameterization scheme with which a numerical model is endowed, the more complex must be the prior probability distribution that accompanies this parameterisation scheme.

As he/she assigns many parameters to a groundwater model for reasons outlined above, a modeller must also beware of omitting important parameters from a model. Parameter omission occurs where non-adjustable simulator structure or boundary conditions are used to represent those aspects of a system that are incompletely known, and hence are better depicted as adjustable. Uncertain “parameters” are thereby wrongly awarded zero prior uncertainty. Equations that are presented in section E3 demonstrate that model parameterisation and structural inadequacies of this nature can adversely affect values and uncertainties assigned to type 1 model predictions through history-matching. However, these parameterisation/structural inadequacies do not affect type 2 model predictions.

Unfortunately, many predictions that are required of a groundwater model are only partly informed by the historical behaviour of a simulated system, and hence are of type 1. Indeed, the making of predictions of this type is often the reason for construction of a groundwater model in the first place. History-matching transfers information to these predictions, but not enough information to constrain them completely. Their values must therefore be expressed probabilistically, using a posterior probability distribution that inherits many characteristics (including possible parameter correlations) from the prior parameter probability distribution.

E4.4 Modelling as Information Partitioning

The above discussion makes it clear that a fundamental task of decision-support groundwater modelling is that of harvesting information that is resident in historical measurements of system behaviour and then (just as importantly) partitioning it into two components. Importantly, this partitioning is prediction-specific. For a given prediction, one of these information components is prediction-relevant, while the other component is prediction-irrelevant. If a prediction is the recipient of information that should not be assigned to it then it incurs bias; at the same time its capacity to express uncertainty is reduced.

Factors that degrade the ability of decision-support groundwater modelling to perform this information separation task include the following.

1. Incorrect or inadequate assignment of a joint prior probability distribution to model parameters. The propensity for error in construction of a prior parameter probability distribution increases with the necessity for decision-support model parameters to be both upscaled and adjustable, and hence somewhat abstract in nature.
2. Failure to denote imperfectly-known components of a decision-support model (such as layering and boundary conditions) as adjustable, and therefore parameterizable. This applies not only to those aspects of a model to which a decision-pertinent prediction may be sensitive. It also applies to any aspect of a model which is illuminated by a history-matching dataset, and hence which should absorb information that is not intended to illuminate that prediction.

The occurrence of both of these problems is inevitable. The mathematics of section E3 suggests that they endow type 1 predictions with a high degree of vulnerability. As stated above, this is expressed as a propensity for predictive bias and for underestimation of predictive uncertainty. (Predictive bias can be viewed as unquantifiable predictive uncertainty.)

E4.5 Parameterisation Limits

The above considerations suggest that there may be an upper limit to the ability of model parameters to harvest and deliver information to predictions that are only partly informed by a history-matching dataset. Other methods of information harvesting and delivery must therefore

be sought when a history-matching dataset contains prediction-relevant information, but where parameter-based information harvesting and delivery cannot be relied upon to separate and deliver to a prediction only that information that is pertinent to it.

E4.6 Numerical Burden Reduction

For a type 2 model prediction, information that is resident in a history-matching dataset is separated and then reassembled. Errors that are incurred in information separation become irrelevant as this information is recombined to make the prediction. These types of prediction have reduced vulnerability to bias and uncertainty under-assessment. Their uncertainties are inherited from noise that accompanies measurements of system behaviour. Parameter nonuniqueness contributes little (if anything) to their uncertainties despite the fact that they may be sensitive to parameters whose values cannot be uniquely estimated.

Because of their relative immunity to errors induced by model imperfections and by shortcomings in designation of a prior parameter probability distribution, it is natural to ask whether it is possible to harvest information that is salient to these types of prediction with reduced numerical cost. Conceptually, this can be accomplished using a model that is structurally simple while retaining a moderate amount of parameterisation complexity; see Doherty and Moore (2021).

In real-world modelling circumstances where data are more abundant than for the simple model that is discussed in sections E2 and E3, parameterisation complexity should be sufficient to ensure that all prediction-salient information is harvested from a history-matching dataset. However values estimated for some model parameters may not lie within the expected range of values of hydraulic properties after which they are named as these parameters compensate for model imperfections. As has been discussed, this does not erode the integrity of type 2 model predictions that are sensitive to them. Meanwhile, the structural simplicity of the model that is designed specifically to make this type of prediction may expedite model run times and reduce the model's propensity for troublesome numerical behaviour, thereby removing impediments to information harvesting. In fact, for predictions of this type, it is permissible for the decision-support model, together with its parameters, to provide a type of "black box" that is optimally designed for prediction-specific machine learning.

E4.7 Ameliorating Predictive Vulnerabilities

E4.7.1 Pre-Model Data Processing

For type 1 model predictions, bias and uncertainty underestimation are dangers that attend parameter-based information separation. Vulnerabilities can be induced by one or more of the following:

- Misrepresentation in a model's structure of the disposition of subsurface hydrogeological units;
- Use of upscaled parameters in which hydraulic property heterogeneity is represented in a somewhat abstract fashion in order to achieve parameter adjustability;
- Endowment of upscaled parameters with prior probability distributions that do not reflect the complex spatial interrelationships of hydraulic properties that exist in the subsurface;

- Use of boundary conditions that simplify simulation of interactions between a local groundwater system and a regional groundwater system, or with local sources and sinks of groundwater;
- Endowment of simplistic and/or variance-diminished prior probability distributions to boundary condition parameters.

As an alternative to using parameters as agents for information separation, it may be possible to remove information that is either irrelevant, or of secondary importance, to a type 1 prediction before that information enters the model and runs the risk of being mishandled. This situation is schematised in figure E4.3.

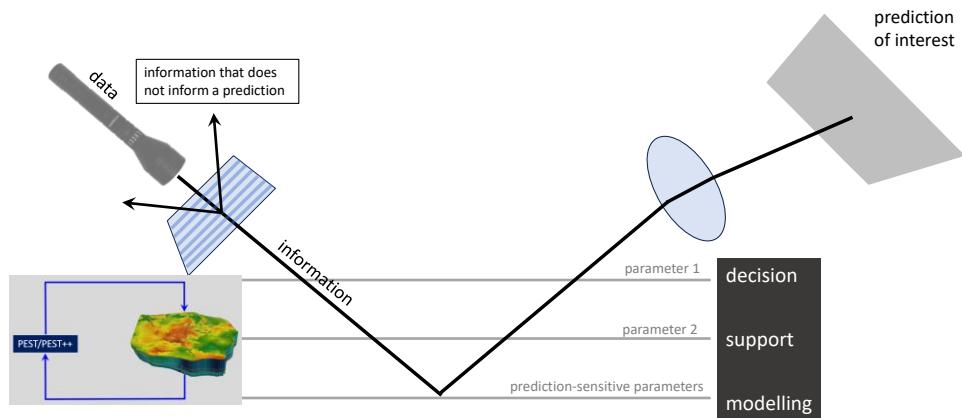


Figure E4.3. Journey of information, where some information has been removed from data prior to its entry into a model.

The extent to which the situation that is depicted in figure E4.3 can be achieved in practice depends heavily on the decision-support modelling context, and on predictions required of a model. Often, something approaching this situation can be achieved by focussing on observation differences (and corresponding model output differences) rather than on the raw values of model outputs. The use of lateral, vertical and temporal head differencing in model history-matching is discussed in sections 15.4 and 16.3 of this book. Some loss of information is incurred by taking these differences; this may result in slightly higher (but quantifiable) uncertainties for some predictions. However the remaining information may require adjustment of fewer parameters as it is assimilated by a model. Errors in information separation may thereby be reduced. This may reduce the propensity for bias of some decision-critical model predictions, particularly those that rely on differencing themselves, for example differences between heads or fluxes that result from different posited management practices, or differences between system states under management plans A and B; see section 16.3.

Other strategies for filtering out non-prediction-salient information before it enters a model may be easy to devise in other modelling circumstances. In most instances, information separation is achieved by strategic processing of field measurements and complementary model outputs before minimising differences between them through history-matching.

E4.7.2 Dispensing with Data Assimilation

If a modeller judges that historical measurements of system behaviour collectively inform a prediction of management interest to only a small extent, then he/she may decide to dispense with parameter-based data assimilation altogether. Justification for this decision may rest on the premise that, under these circumstances, history-matching may do more harm than good to

the decision-support modelling process for reasons that are outlined above. This decision frees a modeller to take either of two modelling paths. The more effective path depends on the decision-support modelling circumstances that the modeller faces.

One option may be to undertake worst-case scenario analysis. This option generally entails significant simplification of a model, thereby allowing a modeller to concentrate on representation of impact pathways that are most salient to a prediction of management interest. Ideally, use of a simplified model for this purpose allows a modeller to ensure that worst case is indeed worst case, as he/she is able to focus his/her attention on those aspects of a system that are of primary predictive importance. This, in turn, may allow a modeller to provide a strong guarantee that the full extent of the pessimistic end of a predictive probability distribution of management interest has been explored by the modelling process. This is especially important where the pessimistic end of a predictive probability distribution is multimodal because of the possible effects of structural or alluvial features that embody connected hydraulic conductivity, but whose existence and properties are difficult to characterise stochastically. See figure E4.4.

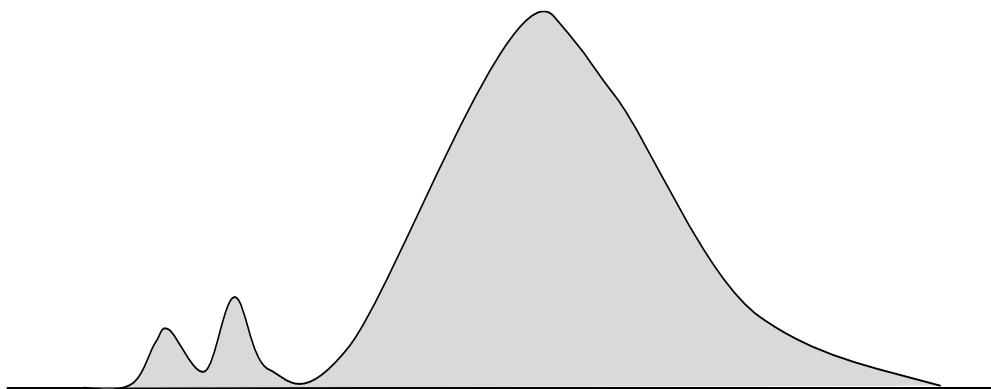


Figure E4.4. Schematic of the probability distribution of a prediction whose value may depend on the existence, or otherwise, of subsurface structural or alluvial features.

As an alternative to simplistic, targeted modelling, exoneration from history-matching may induce a modeller to build a detailed model that is then devoted to Monte-Carlo exploration of predictive possibilities. This may be more appropriate than worst case scenario analysis where predictions of management interest are sensitive to spatially distributed system details such as fracture density, width and orientation that are difficult (or impossible) to adjust during history-matching. Prediction of slope stability behind an expanding mining pit may require this approach.

E4.7.3 Dispensing with Parameters: Data Space Inversion

Data space inversion (DSI) is described in chapter 14 of this book.

DSI uses realisations of hydraulic properties and boundary conditions that are embodied in a complex model to build statistical relationships that directly link the measured past to the managed future. These linkages are then encapsulated in a statistical model whose parameterisation is based on standard normal variates. It is the latter parameters that are adjusted as information is harvested from data and delivered directly to a prediction. Meanwhile, because they do not have to undergo adjustment, system hydraulic properties encapsulated in a complex model can be represented as stochastic, categorical, non-adjustable geobodies of arbitrary complexity whose locations and properties can vary from realisation to realisation. Uncertainties in geostatistical descriptors of these geobodies can also vary between

realisations.

Use of DSI is schematised in figure E4.5.

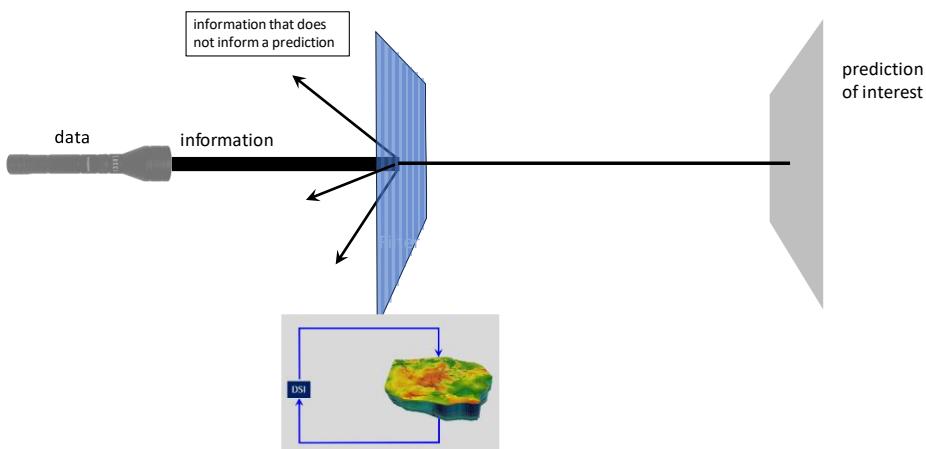


Figure E4.5. Schematic of data space inversion.

As well as avoiding predictive vulnerabilities incurred by parameter-based information harvesting, separation and delivery, DSI can generally be implemented with small numerical cost. It often requires only a few hundred runs of an arbitrarily complex numerical model over a time span that includes that over which measurements of system behaviour were made and that over which management-salient predictions are required.

DSI is a relatively new technology. At the time of writing, experience in its use in decision-support groundwater modelling is somewhat limited. Nevertheless, this experience gives cause for optimism. Furthermore, it is difficult to see how its use could introduce significant predictive bias, nor result in undue corruption of posterior predictive uncertainties, particularly in contexts where the information content of a calibration dataset with respect to a prediction of management interest is medium to low.

E5. Modelling Appropriateness

E5.1 General

This epilogue concludes with some reflections on decision-support modelling appropriateness drawn from insights gained in history-matching the simple model that is described above. These should not be interpreted as “recommendations”. This is because every decision-support modelling context is different, and because technology, and ideas that emerge from use of technology, are changing all the time. Furthermore, as for any other branch of science, a formulaic approach is inappropriate. Nevertheless, these reflections may strike a chord with some readers as they wonder what modelling strategies are best for their particular situation.

E5.2 The Variables

In the graph that is developed hereunder, the “independent variable” (i.e. the variable that appears on the x axis) is “prediction-pertinent information content of history-matching dataset”. This varies from zero to one. Where it is zero, historical measurements of system behaviour have nothing to say about a prediction of management interest. Where it is one, the prediction is data-driven (i.e. solution-space dependent).

Note the prediction-specificity of the independent variable. This implies that decision-support model design and construction should be prediction-based. The tuning of modelling to the

making of a specific prediction is a theme that is echoed throughout this book.

The dependent variable in the following graphs (i.e. the variable that appears on the y axis) is model complexity. This is a rather vague term. Furthermore, Doherty and Moore (2021) distinguish between model structural complexity and model parameterisation complexity. Nevertheless, the term “model complexity” is employed in the present instance in order to simplify the graph. It can be loosely equated to “the amount of detail that a model includes”. This detail certainly includes structural and process complexity. It may also include complexity of representation of hydraulic properties; however, as will be discussed, these may not always be represented as parameters.

Figure E5.1 shows the graph so far.

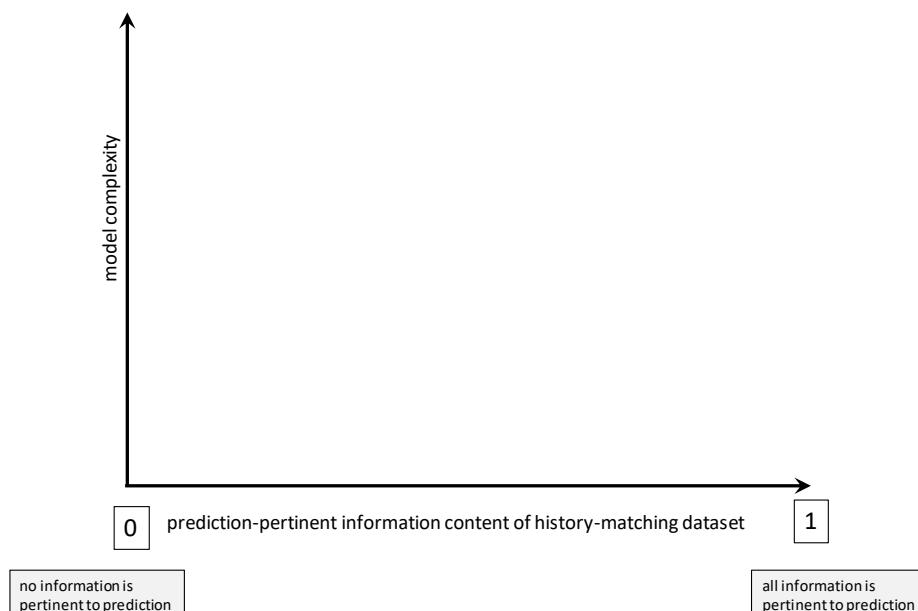


Figure E5.1. Independent and dependent variables of model-appropriateness graph.

E5.3 Parameter-Based History-Matching

We now consider the right-hand portion of the graph. This is where a history-matching dataset is rich in prediction-salient information. In decision-support modelling contexts where this is the case, traditional parameter-based history-matching (implemented through calibration and calibration-constrained uncertainty analysis) may serve the decision-support modelling process well. However, as is discussed above, caution must be exercised. Parameters must be sufficient in number to support a good fit with a history-matching dataset. Furthermore, with dwindling information content of this dataset with respect to a prediction of interest, (that is, as we move from right to left across the graph), greater and greater care must be taken to ensure that

- prior probability distributions associated with different parameter types are “correct” (to the extent that they can be correct when applied to adjustable, upscaled and somewhat abstract parameters such as pilot points);
- aspects of the system that are imperfectly known (including boundary conditions) are represented as adjustable, and parameterised accordingly.

At the same time, it must be ensured that the model is not burdened with numerical problems that lengthen its run time or compromise calculation of finite-difference derivatives or

ensemble-based covariances.

As we move further to the left of the graph, we arrive at a point where maintenance of the integrity of parameter-based information harvesting becomes impossible. This results from the necessity to increase structural and parameterisation complexity, and stochastic properties associated therewith, in order to forestall problems associated with extraction and separation of a diminishing amount of prediction-salient information from a history-matching dataset. The perils of doing this are detailed in earlier sections of this epilogue.

Figure E5.2 schematises the situation that has been described so far.

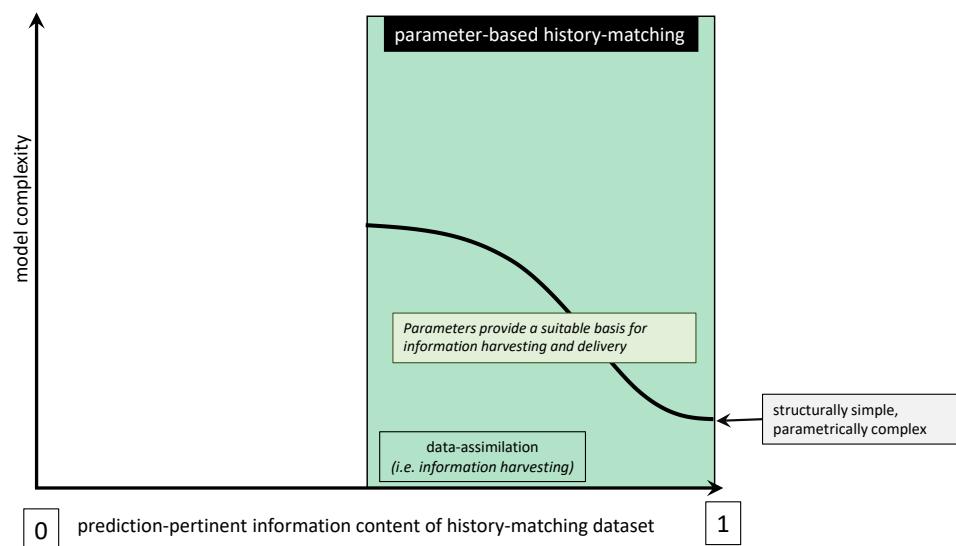


Figure E5.2. Region where parameter-based history-matching is appropriate.

E5.4 Two Options

Where the return on parameter-based history-matching no longer justifies the numerical and other costs of maintaining its integrity, a modeller is left with two options. He/she may turn to decreased or increased model structural complexity as the principal design characteristic of his/her model. The former path leads to worst case scenario analysis wherein information harvesting is eschewed in favour of ensuring that the worst predictive case is not understated. Alternatively, a modeller may decide to include greater structural and hydraulic property complexity in his/her model while dispensing with adjustable (and hence somewhat abstract) parameters. Of course, representation of a high level of hydraulic property complexity in a model must necessarily be stochastic. Hence a decision to take this modelling path is also a decision to quantify uncertainty using some variant of Monte Carlo analysis. See figure E5.3.

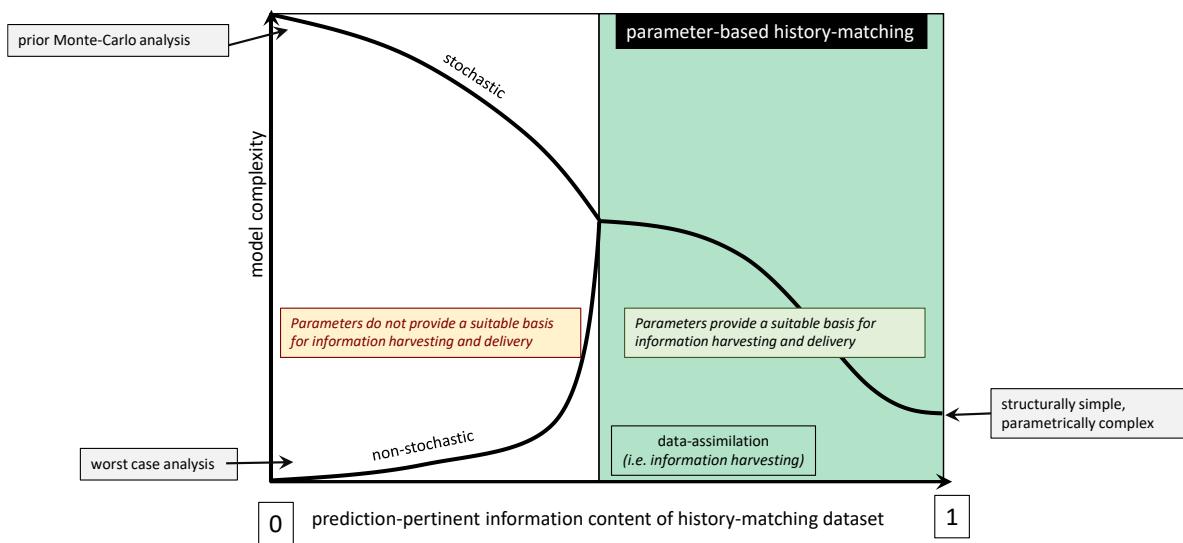


Figure E5.3. Extension of decision-support modelling into regions where parameter-based history-matching becomes inappropriate.

E5.5 Data Space Inversion

The decision to adopt probabilistic expression of complex hydraulic features does not preclude harvesting of information from a history-matching dataset. However this can no longer be accomplished using adjustable parameters. Data space inversion (DSI) becomes the only practical option for extraction of prediction-pertinent information from measurements of past system behaviour and for delivery of that information to predictions of management interest.

Use of DSI reduces information vulnerability during its journey from data to decisions. At the same time, DSI is numerically cheap. However a modeller who accepts the advantages offered by DSI must also accept the responsibility of giving full stochastic expression to subsurface hydraulic features that may influence the value of a decision-critical model prediction.

DSI is represented in figure E5.4. Note that the shaded area where DSI can be used extends into the area where parameter-based history-matching is also useable. Given the light model run burden of DSI, it can thereby be used to check the integrity of parameter-based information harvesting and delivery where the two are used together.

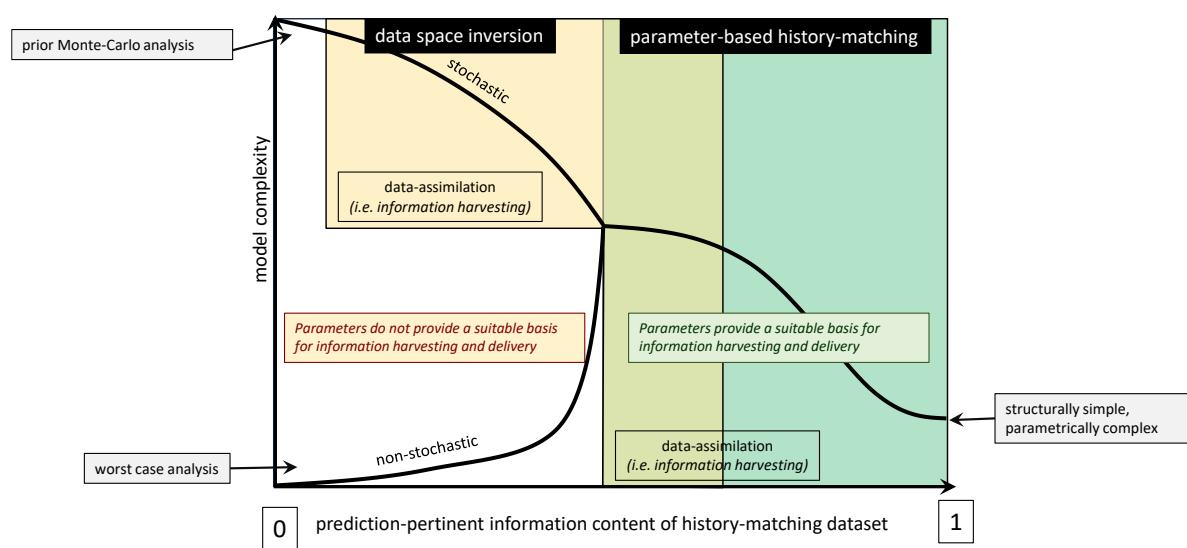


Figure E5.4. The completed map of decision-support modelling appropriateness.