



JAMES R CARPENTER, JONATHAN W BARTLETT,
TIM P MORRIS, ANGELA M WOOD,
MATTEO QUARTAGNO, MICHAEL G KENWARD

Multiple Imputation and its Application

Second Edition

STATISTICS IN PRACTICE

WILEY

Multiple Imputation and its Application

Statistics in Practice

Founding Editor

Vic Barnett

Nottingham Trent University, UK

Statistics in Practice is an important international series of texts which provide detailed coverage of statistical concepts, methods and worked case studies in specific fields of investigation and study.

With sound motivation and many worked practical examples, the books show in down-to-earth terms how to select and use an appropriate range of statistical techniques in a particular practical field within each title's special topic area.

The books provide statistical support for professionals and research workers across a range of employment fields and research environments. Subject areas covered include medicine and pharmaceuticals; industry, finance and commerce; public services; the earth and environmental sciences, and so on.

The books also provide support to students studying statistical courses applied to the above areas. The demand for graduates to be equipped for the work environment has led to such courses becoming increasingly prevalent at universities and colleges.

It is our aim to present judiciously chosen and well-written workbooks to meet everyday practical needs. Feedback of views from readers will be most valuable to monitor the success of this aim.

Multiple Imputation and its Application

Second Edition

James R. Carpenter

*London School of Hygiene & Tropical Medicine, London, UK
MRC Clinical Trials Unit, University College London, UK*

Jonathan W. Bartlett

London School of Hygiene & Tropical Medicine, London, UK

Tim P. Morris

MRC Clinical Trials Unit, University College London, UK

Angela M. Wood

University of Cambridge, Cambridge, UK

Matteo Quartagno

MRC Clinical Trials Unit at University College London, UK

Michael G. Kenward

Ashkirk, UK

WILEY

This edition first published 2023
© 2023 John Wiley & Sons Ltd

Edition History: John Wiley & Sons Ltd (1e, 2013)

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by law. Advice on how to obtain permission to reuse material from this title is available at <http://www.wiley.com/go/permissions>.

The right of James R. Carpenter, Jonathan W. Bartlett, Tim P. Morris, Angela M. Wood, Matteo Quartagno and Michael G. Kenward to be identified as the authors of this work has been asserted in accordance with law.

Registered Offices

John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, USA
John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, UK

For details of our global editorial offices, customer services, and more information about Wiley products visit us at www.wiley.com.

Wiley also publishes its books in a variety of electronic formats and by print-on-demand. Some content that appears in standard print versions of this book may not be available in other formats.

Trademarks: Wiley and the Wiley logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

Limit of Liability/Disclaimer of Warranty

While the publisher and authors have used their best efforts in preparing this work, they make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives, written sales materials or promotional statements for this work. This work is sold with the understanding that the publisher is not engaged in rendering professional services. The advice and strategies contained herein may not be suitable for your situation. You should consult with a specialist where appropriate. The fact that an organization, website, or product is referred to in this work as a citation and/or potential source of further information does not mean that the publisher and authors endorse the information or services the organization, website, or product may provide or recommendations it may make. Further, readers should be aware that websites listed in this work may have changed or disappeared between when this work was written and when it is read. Neither the publisher nor authors shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

Library of Congress Cataloging-in-Publication Data applied for

ISBN: 9781119756088 (Hardback); 9781119756095 (Adobe PDF); 9781119756101 (ePub)

Cover design: Wiley
Cover image: Courtesy of Harvey Goldstein

Set in 10/12pt TimesLTStd by Straive, Chennai, India

Contents

Preface to the second edition	xiii
Data acknowledgements	xv
Acknowledgements	xvii
Glossary	xix
PART I FOUNDATIONS	1
1 Introduction	3
1.1 Reasons for missing data	5
1.2 Examples	6
1.3 Patterns of missing data	7
1.3.1 Consequences of missing data	9
1.4 Inferential framework and notation	10
1.4.1 Missing completely at random (MCAR)	12
1.4.2 Missing at random (MAR)	13
1.4.3 Missing not at random (MNAR)	17
1.4.4 Ignorability	21
1.5 Using observed data to inform assumptions about the missingness mechanism	21
1.6 Implications of missing data mechanisms for regression analyses	24
1.6.1 Partially observed response	24
1.6.2 Missing covariates	27
1.6.3 Missing covariates and response	30
1.6.4 Subtle issues I: the odds ratio	30
1.6.5 Implication for linear regression	32
1.6.6 Subtle issues II: sub-sample ignorability	33
1.6.7 Summary: when restricting to complete records is valid	34
1.7 Summary	34
Exercises	35

2	The multiple imputation procedure and its justification	39
2.1	Introduction	39
2.2	Intuitive outline of the MI procedure	40
2.3	The generic MI procedure	45
2.4	Bayesian justification of MI	48
2.5	Frequentist inference	50
2.5.1	Large number of imputations	50
2.5.2	Small number of imputations	51
2.5.3	Inference for vector β	53
2.5.4	Combining likelihood ratio tests	54
2.6	Choosing the number of imputations	55
2.7	Some simple examples	56
2.7.1	Estimating the mean with σ^2 known by the imputer and analyst	57
2.7.2	Estimating the mean with σ^2 known only by the imputer	59
2.7.3	Estimating the mean with σ^2 unknown	59
2.7.4	General linear regression with σ^2 known	61
2.8	MI in more general settings	64
2.8.1	Proper imputation	64
2.8.2	Congenial imputation and substantive model	64
2.8.3	Uncongenial imputation and substantive models	65
2.8.4	Survey sample settings	71
2.9	Constructing congenial imputation models	72
2.10	Discussion	73
	Exercises	73

PART II MULTIPLE IMPUTATION FOR SIMPLE DATA STRUCTURES **79**

3	Multiple imputation of quantitative data	81
3.1	Regression imputation with a monotone missingness pattern	81
3.1.1	MAR mechanisms consistent with a monotone pattern	83
3.1.2	Justification	84
3.2	Joint modelling	85
3.2.1	Fitting the imputation model	85
3.2.2	Adding covariates	89
3.3	Full conditional specification	90
3.3.1	Justification	91
3.4	Full conditional specification versus joint modelling	92
3.5	Software for multivariate normal imputation	93
3.6	Discussion	93
	Exercises	94

4	Multiple imputation of binary and ordinal data	96
4.1	Sequential imputation with monotone missingness pattern	96
4.2	Joint modelling with the multivariate normal distribution	98
4.3	Modelling binary data using latent normal variables	100
4.3.1	Latent normal model for ordinal data	104
4.4	General location model	108
4.5	Full conditional specification	108
4.5.1	Justification	109
4.6	Issues with over-fitting	110
4.7	Pros and cons of the various approaches	114
4.8	Software	116
4.9	Discussion	116
	Exercises	117
5	Imputation of unordered categorical data	119
5.1	Monotone missing data	119
5.2	Multivariate normal imputation for categorical data	121
5.3	Maximum indicant model	121
5.3.1	Continuous and categorical variable	123
5.3.2	Imputing missing data	125
5.4	General location model	125
5.5	FCS with categorical data	128
5.6	Perfect prediction issues with categorical data	130
5.7	Software	130
5.8	Discussion	130
	Exercises	131
PART III MULTIPLE IMPUTATION IN PRACTICE		133
6	Non-linear relationships, interactions, and other derived variables	135
6.1	Introduction	135
6.1.1	Interactions	137
6.1.2	Squares	137
6.1.3	Ratios	138
6.1.4	Sum scores	139
6.1.5	Composite endpoints	140
6.2	No missing data in derived variables	141
6.3	Simple methods	143
6.3.1	Impute then transform	143
6.3.2	Transform then impute/just another variable	143
6.3.3	Adapting standard imputation models and passive imputation	145

6.3.4	Predictive mean matching	146
6.3.5	Imputation separately by groups for interactions	148
6.4	Substantive-model-compatible imputation	152
6.4.1	The basic idea	152
6.4.2	Latent-normal joint model SMC imputation	157
6.4.3	Factorised conditional model SMC imputation	160
6.4.4	Substantive model compatible fully conditional specification	161
6.4.5	Auxiliary variables	162
6.4.6	Missing outcome values	163
6.4.7	Congeniality versus compatibility	163
6.4.8	Discussion of SMC imputation	164
6.5	Returning to the problems	165
6.5.1	Ratios	165
6.5.2	Splines	165
6.5.3	Fractional polynomials	166
6.5.4	Multiple imputation with conditional questions or ‘skips’	169
	Exercises	172
7	Survival data	175
7.1	Missing covariates in time-to-event data	175
7.1.1	Approximately compatible approaches	176
7.1.2	Substantive model compatible approaches	181
7.2	Imputing censored event times	186
7.3	Non-parametric, or ‘hot deck’ imputation	188
7.3.1	Non-parametric imputation for time-to-event data	189
7.4	Case-cohort designs	191
7.4.1	Standard analysis of case-cohort studies	192
7.4.2	Multiple imputation for case-cohort studies	193
7.4.3	Full cohort	193
7.4.4	Intermediate approaches	193
7.4.5	Sub-study approach	194
7.5	Discussion	197
	Exercises	197
8	Prognostic models, missing data, and multiple imputation	200
8.1	Introduction	200
8.2	Motivating example	201
8.3	Missing data at model implementation	201
8.4	Multiple imputation for prognostic modelling	202
8.5	Model building	202
8.5.1	Model building with missing data	202
8.5.2	Imputing predictors when model building is to be performed	204

8.6	Model performance	204
8.6.1	How should we pool MI results for estimation of performance?	205
8.6.2	Calibration	205
8.6.3	Discrimination	206
8.6.4	Model performance measures with clinical interpretability	206
8.7	Model validation	206
8.7.1	Internal model validation	207
8.7.2	External model validation	208
8.8	Incomplete data at implementation	208
8.8.1	MI for incomplete data at implementation	208
8.8.2	Alternatives to multiple imputation	210
	Exercises	212
9	Multi-level multiple imputation	213
9.1	Multi-level imputation model	213
9.1.1	Imputation of level-1 variables	216
9.1.2	Imputation of level 2 variables	219
9.1.3	Accommodating the substantive model	223
9.2	MCMC algorithm for imputation model	224
9.2.1	Ordered and unordered categorical data	226
9.2.2	Imputing missing values	227
9.2.3	Substantive model compatible imputation	227
9.2.4	Checking model convergence	229
9.3	Extensions	231
9.3.1	Cross-classification and three-level data	231
9.3.2	Random level 1 covariance matrices	232
9.3.3	Model fit	234
9.4	Other imputation methods	234
9.4.1	One-step and two-step FCS	234
9.4.2	Substantive model compatible imputation	235
9.4.3	Non-parametric methods	236
9.4.4	Comparisons of different methods	236
9.5	Individual participant data meta-analysis	237
9.5.1	Different measurement scales	239
9.5.2	When to apply Rubin's rules	239
9.5.3	Homoscedastic versus heteroscedastic imputation model	240
9.6	Software	241
9.7	Discussion	241
	Exercises	242
10	Sensitivity analysis: MI unleashed	245
10.1	Review of MNAR modelling	246
10.2	Framing sensitivity analysis: estimands	249

10.2.1	Definition of the estimand	249
10.2.2	Two common estimands	250
10.3	Pattern mixture modelling with MI	251
10.3.1	Missing covariates	256
10.3.2	Sensitivity with multiple variables: the NAR FCS procedure	258
10.3.3	Application to survival analysis	260
10.4	Pattern mixture approach with longitudinal data via MI	263
10.4.1	Change in slope post-deviation	264
10.5	Reference based imputation	267
10.5.1	Constructing joint distributions of pre- and post-intercurrent event data	268
10.5.2	Technical details	269
10.5.3	Software	271
10.5.4	Information anchoring	275
10.6	Approximating a selection model by importance weighting	279
10.6.1	Weighting the imputations	281
10.6.2	Stacking the imputations and applying the weights	282
10.7	Discussion	289
	Exercises	290
11	Multiple imputation for measurement error and misclassification	294
11.1	Introduction	294
11.2	Multiple imputation with validation data	296
11.2.1	Measurement error	297
11.2.2	Misclassification	297
11.2.3	Imputing assuming error is non-differential	299
11.2.4	Non-linear outcome models	299
11.3	Multiple imputation with replication data	301
11.3.1	Measurement error	302
11.3.2	Misclassification	306
11.4	External information on the measurement process	307
11.5	Discussion	308
	Exercises	309
12	Multiple imputation with weights	312
12.1	Using model-based predictions in strata	313
12.2	Bias in the MI variance estimator	314
12.3	MI with weights	317
12.3.1	Conditions for the consistency of $\hat{\theta}_{MI}$	317
12.3.2	Conditions for the consistency of \hat{V}_{MI}	318
12.4	A multi-level approach	320
12.4.1	Evaluation of the multi-level multiple imputation approach for handling survey weights	322
12.4.2	Results	325

12.5	Further topics	328
12.5.1	Estimation in domains	328
12.5.2	Two-stage analysis	328
12.5.3	Missing values in the weight model	329
12.6	Discussion	329
	Exercises	330
13	Multiple imputation for causal inference	333
13.1	Multiple imputation for causal inference in point exposure studies	333
13.1.1	Randomised trials	335
13.1.2	Observational studies	335
13.2	Multiple imputation and propensity scores	338
13.2.1	Propensity scores for confounder adjustment	338
13.2.2	Multiple imputation of confounders	340
13.2.3	Imputation model specification	342
13.3	Principal stratification <i>via</i> multiple imputation	343
13.3.1	Principal strata effects	344
13.3.2	Estimation	345
13.4	Multiple imputation for IV analysis	346
13.4.1	Instrumental variable analysis for non-adherence	346
13.4.2	Instrumental variable analysis <i>via</i> multiple imputation	348
13.5	Discussion	350
	Exercises	351
14	Using multiple imputation in practice	355
14.1	A general approach	355
14.1.1	Explore the proportions and patterns of missing data	356
14.1.2	Consider plausible missing data mechanisms	356
14.1.3	Consider whether missing at random is plausible	356
14.1.4	Choose the variables for the imputation model	357
14.1.5	Choose an appropriate imputation strategy and model/s	357
14.1.6	Set and record the seed of the pseudo-random number generator	357
14.1.7	Fit the imputation model	358
14.1.8	Iterate and revise the imputation model if necessary	358
14.1.9	Estimate monte carlo error	358
14.1.10	Sensitivity analysis	359
14.2	Objections to multiple imputation	359
14.3	Reporting of analyses with incomplete data	363
14.4	Presenting incomplete baseline data	364
14.5	Model diagnostics	365
14.6	How many imputations?	366
14.6.1	Using the jack-knife estimate of the Monte-Carlo standard error	368

14.7	Multiple imputation for each substantive model, project, or dataset?	369
14.8	Large datasets	370
14.8.1	Large datasets and joint modelling	371
14.8.2	Shrinkage by constraining parameters	372
14.8.3	Comparison of the two approaches	375
14.9	Multiple imputation and record linkage	375
14.10	Setting random number seeds for multiple imputation analyses	377
14.11	Simulation studies including multiple imputation	377
14.11.1	Random number seeds for simulation studies including multiple imputation	377
14.11.2	Repeated simulation of all data or only the missingness mechanism?	378
14.11.3	How many imputations for simulation studies?	379
14.11.4	Multiple imputation for data simulation	380
14.12	Discussion	381
	Exercises	381
Appendix A Markov Chain Monte Carlo		384
A.1	Metropolis Hastings sampler	385
A.2	Gibbs sampler	386
A.3	Missing data	387
Appendix B Probability distributions		388
B.1	Posterior for the multivariate normal distribution	391
Appendix C Overview of multiple imputation in R, Stata		394
C.1	Basic multiple imputation using R	394
C.2	Basic MI using Stata	395
References		398
Author Index		419
Index of Examples		429
Subject Index		431

Preface to the second edition

No study of any complexity manages to collect all the intended data. Analysis of the resulting partially collected data must therefore address the issues raised by the missing data. Beyond simply estimating the proportion of missing values, the interplay between the substantive questions and the reasons for the missing data is crucial. There is no simple, universal, solution.

Suppose, for a substantive question at hand, the consequences of missing data in terms of bias and loss of precision are non-trivial. Then the analyst must make a set of assumptions about the reasons, or mechanisms, causing data to be missing, and perform an inferentially valid analysis under these assumptions. In this regard, analysis of a partially observed dataset is the same as any statistical analysis; the difference is that when data are missing we cannot assess the validity of these assumptions in the way we might do in a regression analysis, for example. Hence, sensitivity analysis, where we explore the robustness of inference to different assumptions about the reasons for missing data, is important.

Given a set of assumptions about the reasons data are missing, there are a number of statistical methods for carrying out the analysis. These include the expectation-maximization (EM) algorithm, inverse probability (of non-missingness) weighting, a full Bayesian analysis and, depending on the setting, a direct application of maximum likelihood. These methods, and those derived from them, each have their own advantages in particular settings. We focus on multiple imputation for its practical utility, broad applicability, and relatively straightforward application. Since the first edition was published ten years ago, new applications of multiple imputation have continued to emerge and we have had to be selective in what we cover. The topics included are those we have found most relevant for our research and teaching.

Like the first edition, the book is divided into three parts. Part I lays the foundations, with an introductory chapter outlining the issues raised by missing data, followed by a chapter describing the theoretical foundations of multiple imputation. Part II describes the application of multiple imputation for standard regression analyses, explaining how MI can be used for continuous, categorical, and ordinal data. Part III describes how to apply MI in a range of practical settings, specifically analysis with non-linear relationships, analysis of survival data, development and validation of prognostic models, analysis with multilevel data structures, sensitivity analysis, handling measurement error, analysis involving weights, and causal inference. We conclude with a chapter outlining some broad practical points on the application of

multiple imputation. While readers may wish to read only specific relevant chapters in Part III, Chapter 14 is intended to be relevant to all readers. We illustrate ideas with a range of examples from the medical and social sciences, reflecting the wide application that MI has seen in recent years.

Each chapter concludes with a range of exercises, designed to consolidate and deepen understanding of the material. The computer-based exercises have been designed with R and Stata users in mind. The book's home page at <https://missingdata.lshtm.ac.uk> contains both (i) hints for the exercises (including suggestions for R and Stata code) and (ii) full solutions where applicable.

We welcome feedback from readers. Please email `james.carpenter@lshtm.ac.uk` in the first instance.

*James R. Carpenter, Jonathan W. Bartlett, Tim P. Morris,
Angela M. Wood, Matteo Quartagno and Mike G. Kenward*
September 2022

Data acknowledgements

We are grateful to the following:

1. AstraZeneca for permission to use data from the 5-arm asthma study in examples in Chapters 1, 3, 6, and 10;
2. GlaxoSmithKline for permission to use data from the dental pain study in Chapter 4;
3. Mike English (Director, Child and Newborn Health Group, Nairobi, Kenya) for permission to data from a multi-faceted intervention to implement guidelines and improve admission paediatric care in Kenyan district hospitals in Chapter 9;
4. Peter Blatchford for permission to use data from the Class Size Study (Blatchford *et al.*, 2002) in Chapter 9;
5. Sara Schroter for permission to use data from the study to improve the quality of peer review in Chapter 10.

In Chapter 12, we used data from the Millennium Cohort Study available through the UK Data Service (ukdataservice.ac.uk), study number 2000031.

In Chapters 1, 5, 9, and 10, we have analysed data from the Youth Cohort Time Series for England, Wales, and Scotland, 1984–2002 First Edition, Colchester, Essex, published by and freely available from the UK Data Archive, Study Number SN 5765. Thanks to Vernon Gayle for introducing us to these data.

In Chapter 6, we have analysed data from the Alzheimer’s Disease Neuro-imaging Initiative (ADNI) database (adni.loni.ucla.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this book. A complete listing of ADNI investigators can be found at http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf.

The ADNI was launched in 2003 as a public–private partnership, led by Principal Investigator Michael W. Weiner, MD. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimer’s disease (AD). For up-to-date information, see www.adni-info.org.

Data collection and sharing for this project was funded by the ADNI (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer's Association; Alzheimer's Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organisation is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer's Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

In Chapter 11, we have analysed data from the US National Health and Nutrition Examination Survey (NHANES), which is freely available from the US Centers for Disease Control and Prevention. We thank the participants and survey organisers of NHANES.

Acknowledgements

No book of this kind is written in a vacuum, and we are grateful to many friends and colleagues.

In particular, we would like to thank the following for discussions and/or inputs that have contributed to this book or shaped our thinking through the years: John Carlin, Orlagh Carroll, Stephen Cole, Suzie Cro, Rhian Daniel, Dan Jackson, Ruth Keogh, Katherine Lee, Clemence Leyrat, Geert Molenberghs, Tra My Pham, Freya Roberts, James Roger, Jeremy Saxe, Shaun Seaman, Jonathan Sterne, Kate Tilling, and Ian White.

Specific thanks to Tetiana Gorbach for checking the exercises in Chapter 2.

James would like to acknowledge many years' collaboration with Harvey Goldstein, who died in 2020.

James would like to thank Mike Elliott, Rod Little, Trivellore Raghunathan, and Jeremy Taylor for facilitating a visit to the Institute for Social Research and Department of Biostatistics at the University of Michigan, Ann Arbor, in Summer 2011, when the first draft of the first edition was written.

Thanks to Tim Collier for the anecdote in Section 1.3.

We also gratefully acknowledge funding support from the ESRC (three-year fellowship for James Carpenter, RES-063-27-0257, and follow-on funding RES-189-25-0103) and MRC (grants G0900724, G0900701 and G0600599, MC_UU_12023/21, MC_UU_12023/29, MC_UU_00004/07, MC_UU_00004/09 and MC_UU_00004/06, MR/T023953/1).

We would also like to thank Richard Davies and Kathryn Sharples at Wiley for their encouragement and support with the first edition, and Kimberly Monroe-Hill and Alison Oliver for their support and patient encouragement with the second edition.

Lastly, thanks to our families for their forbearance and understanding over the course of this project.

Despite the encouragement and support of those listed above, the text inevitably contains errors and shortcomings, for which we take full responsibility.

*James R. Carpenter, Jonathan W. Bartlett, Tim P. Morris,
Angela M. Wood, Matteo Quartagno, and Mike G. Kenward*
September 2022

Glossary

Indices and symbols

i	indexes units, often individuals
j	indexes variables in the dataset
n	total number of units in the dataset
p	depending on context, number of variables in the dataset or number of parameters in a statistical model
W, X, Y, Z	random variables
$Y_{i,j}$	i th observation on j th variable, $i = 1, \dots, n; j = 1, \dots, p$
R	response indicator, where $R_i = 1$ denotes observed data for unit i and $R_i = 0$ indicates missing data
θ	generic parameter
θ	generic parameter column vector, typically p by 1
$\alpha, \beta, \gamma, \delta$	regression coefficients
β	column vector of regression coefficients, typically p by 1.
*	a random draw, typically from a probability distribution but sometimes data/a datum
$U(\cdot)$	a generic score statistic

Matrices

Ω	Matrix, typically of dimension $p \times p$
$\Omega_{i,j}$	i, j th element of Ω
Ω^T	Transpose of Ω , so that $\Omega_{i,j}^T = \Omega_{j,i}$
$Y_j = (Y_{1,j}, \dots, Y_{n,j})^T$	n by 1 column vector of observations on variable j
$\text{tr}(\Omega)$	Sum of diagonal elements of Ω , i.e. $\sum \Omega_{i,i}$ known as the trace of the matrix

Abbreviations

CAR	censoring at random
CNAR	censoring not at random
EM	expectation maximisation
FCS	full conditional specification

FEV ₁	forced expiratory volume in one second (measured in litres)
FMI	fraction of missing information
IPW	inverse probability weighting
IV	instrumental variable
MAR	missing at random
MCAR	missing completely at random
MI	multiple imputation
MNAR	missing not at random
PMM	predictive mean matching

Probability distributions

$f(\cdot)$	Probability distribution function
$F(\cdot)$	Cumulative distribution function
	to be verbalised ‘given’ or ‘conditional on’, as in $f(Y X)$ ‘the probability distribution function of Y given [conditional on] X ’

Miscellaneous

We use the terms *complete records* and *incomplete records* rather than *complete cases* and *incomplete cases*, respectively. For generic regression of Y on \mathbf{X} , we use the terms *outcome* or *dependent variable* for Y and *covariates* or *independent variables* for X .

PART I

FOUNDATIONS

1

Introduction

Collecting, analysing, and drawing inferences from data are central to research in the medical and social sciences. Unfortunately, for any number of reasons, it is rarely possible to collect all the intended data. The ubiquity of missing data, and the problems this poses for both analysis and inference, has spawned a substantial statistical literature dating from 1950s. At that time, when statistical computing was in its infancy, many analyses were only feasible because of the carefully planned balance in the dataset (for example the same number of observations on each unit). Missing data meant the available data for analysis were unbalanced, thus complicating the planned analysis and in some instances rendering it infeasible. Early work on the problem was therefore largely computational (e.g. Healy and Westmacott, 1956, Afifi and Elashoff, 1966, Orchard and Woodbury, 1972, Dempster *et al.*, 1977).

The wider question of the consequences of non-trivial proportions of missing data for inference was neglected until the seminal paper by Rubin (1976). This set out a typology for assumptions about the reasons for missing data and sketched their implications for analysis and inference. It marked the beginning of a broad stream of research about the analysis of partially observed data. The literature is now huge and continues to grow, both as methods are developed for large and complex data structures, and as increasing computer power and suitable software enables researchers to apply these methods.

For a broad overview of the literature, a good place to start for applied statisticians is Little and Rubin (2019). They give a good overview of likelihood methods and an introduction to multiple imputation. Allison (2002) presents a less technical overview. Schafer (1997) is more algorithmic, focusing on the expectation maximisation (EM) algorithm and imputation using the multivariate normal and general location model. Molenberghs and Kenward (2007) focus on clinical studies, while Daniels and Hogan (2008) focus on longitudinal studies with a Bayesian emphasis.

The above books concentrate on the parametric approaches. However, there is also a growing literature based around using inverse probability weighting, in the spirit of Horvitz and Thompson (1952), and associated doubly robust methods. In particular, we refer to the work of Robins and colleagues (e.g. Robins and Rotnitzky, 1995, Scharfstein *et al.*, 1999). Vansteelandt *et al.* (2009) give an accessible introduction to these developments. A comparison with multiple imputation in a simple setting is given by Carpenter *et al.* (2006). The pros and cons are debated in Kang and Schafer (2007) and the theory is brought together by Tsiatis (2006).

This book is concerned with a particular statistical method for analysing and drawing inferences from incomplete data called *multiple imputation* (MI). Initially proposed by Rubin (1987) in the context of surveys, increasing awareness among researchers about the possible effects of missing data (e.g. Klebanoff and Cole, 2008) has led to an upsurge of interest (e.g. Sterne *et al.* (2009), Kenward and Carpenter (2007), Schafer (1999a), Rubin (1996)), fuelled by the increasing availability of software and computing power.

MI is attractive because it is both practical and widely applicable. Well-developed statistical software (see, for example, issue 45 of the Journal of Statistical Software) has placed MI within the reach of most researchers in the medical and social sciences, whether or not they have undertaken advanced training in statistics. However, the increasing use of MI in a range of settings beyond that originally envisaged has led to a bewildering proliferation of algorithms and software. Further, the implications of the underlying assumptions in the context of the data at hand are often unclear.

We are writing for researchers in the medical and social sciences with the aim of clarifying the issues raised by missing data, outlining the rationale for MI, explaining the motivation and relationship between the various imputation algorithms and describing and illustrating its application in various settings and to some complex data structures.

Throughout most of the book (with the partial exception of Chapter 8), we will assume that a key aim of analysis with incomplete data is to recover the information lost due to missing data. More specifically, we will take the ‘substantive model’ as the model that would be used with complete data. We can then define certain desirable properties of our estimator with incomplete data. First, it should be unbiased for the value of the parameter we would see with complete data. Second, it should have low variance. Third, we should have a reliable variance formula and a means of constructing confidence intervals with the advertised coverage.

In the context of multiple imputation, it is worth noting that these remain our aims; the aim of multiple imputation is not to accurately predict the missing values. Rubin (1996) describes it as follows:

‘Judging the quality of missing data procedures by their ability to recreate the individual missing values [...] does not lead to choosing procedures that result in valid inference, which is our objective’.

An objection may be that the ability to perfectly predict missing values *would* result in valid inference; however, in our view, this hypothetical scenario would be one in which data are not really ‘missing’.

Central to the analysis of partially observed data is an understanding of why the data are missing and the implications of this for the analysis. This is the focus of the remainder of this chapter. Introducing some of the examples that run through the book, we show how Rubin's typology (Rubin, 1976) provides the foundational framework for understanding the implications of missing data.

1.1 Reasons for missing data

In this section, we consider possible reasons for missing data, illustrate these with examples, and note some preliminary implications for inference. We use the word 'possible' advisedly, since we can rarely be sure of the mechanism giving rise to missing data. Instead, a range of possible mechanisms are consistent with the observed data. In practice, we therefore wish to analyse the data under different mechanisms to establish the robustness of our inference in the face of uncertainty about the missingness mechanism.

All datasets consist of a series of *units* each of which provides information on a series of *items*. For example, in a cross-sectional questionnaire survey, the units would be individuals, and the items their answers to the questions. In a household survey, the units would be households, and the items information about the household and members of the household. In longitudinal studies, units would typically be individuals, while items would be longitudinal data from those individuals. In this book, units

Figure 1.1 Detail from a senior mandarin's house front in New Territories, Hong Kong.

therefore correspond to the highest level in multi-level (i.e. hierarchical) data, and unless stated otherwise, data from different units are statistically independent.

Within this framework, it is useful to distinguish between units where all the information is missing, termed *unit non-response* and units who contribute partial information, termed *item non-response*. The statistical issues are the same in both cases and both can in principle be handled by MI. However, the main focus of this book is the latter.

Example 1.1 Mandarin tableau

Figure 1.1, which is also shown on the book's cover, shows part of the frontage of a senior mandarin's house in the New Territories, Hong Kong. We suppose interest focuses on characteristics of the figurines, for example their number, height, facial characteristics, and dress. Unit non-response then corresponds to missing figurines, and item non-response to damaged – hence, partially observed – figurines. □

1.2 Examples

We now introduce two key examples, which we return to throughout the book.

Example 1.2 Youth Cohort Study (YCS)

The Youth Cohort Study of England and Wales (YCS) is an ongoing UK government-funded representative survey of pupils in England and Wales at school-leaving age (School year 11, age 16–17) (UK Data Archive, 2007). Each year that a new cohort is surveyed, detailed information is collected on each young person's experience of education, and their qualifications, as well as information on employment and training. A limited amount of information is collected on their personal characteristics, family, home circumstances, and aspirations.

Over the life cycle of the YCS, different organisations have had responsibility for the structure and timings of data collection. Unfortunately, the documentation of older cohorts is poor. Croxford *et al.* (2007) have deposited a harmonised dataset that comprises YCS cohorts from 1984 to 2002 (UK Data Archive Study Number 5765 dataset). We consider data from pupils attending comprehensive schools from five YCS cohorts; these pupils reached the end of Year 11 in 1990, 1993, 1995, 1997, and 1999.

We explore relationships between Year 11 educational attainment (the General Certificate of Secondary Education) and key measures of social stratification. The units are pupils, and the items are measurements on these pupils, and a non-trivial number of items are partially observed. □

Example 1.3 Randomised controlled trial of patients with chronic asthma

We consider data from a five-arm asthma clinical trial to assess the efficacy and safety of budesonide, a second-generation glucocorticosteroid, on patients with

chronic asthma. Four hundred and seventy-three patients with chronic asthma were enrolled in the 12-week randomised, double-blind, multi-centre parallel-group trial, which compared the effect of a daily dose of 200, 400, 800, or 1600 mcg of budesonide with placebo.

Key outcomes of clinical interest include patients' peak expiratory flow rate (their maximum speed of expiration in litres/minute) and their forced expiratory volume, FEV₁ (the volume of air, in litres, the patient with fully inflated lungs can breathe out in one second). In summary, the trial found a statistically significant dose-response effect for the mean change from baseline over the study for both morning peak expiratory flow, evening peak expiratory flow, and FEV₁ at the 5% level.

Budesonide-treated patients also showed reduced asthma symptoms and bronchodilator use compared with placebo, while there were no clinically significant differences in treatment-related adverse experiences between the treatment groups. Further details about the conduct of the trial, its conclusions, and the variables collected can be found elsewhere (Busse *et al.*, 1998). Here, we focus on FEV₁ and confine our attention to the placebo and lowest active dose arms. FEV₁ was collected at baseline, then 2, 4, 8, and 12 weeks after randomisation. The intention was to compare FEV₁ across treatment arms at 12 weeks. However (excluding three patients whose participation in the study was intermittent), only 37 out of 90 patients in the placebo arm, and 71 out of 90 patients in the lowest active dose arm, still remained in the trial at 12 weeks. □

1.3 Patterns of missing data

It is very important to investigate the patterns of missing data before embarking on a formal analysis. This can throw up vital information that might otherwise be overlooked and may even allow the missing data to be traced. For example, when analysing the new wave of a longitudinal survey, a colleague's careful examination of missing data patterns established that many of the missing questionnaires could be traced to a set of cardboard boxes. These turned out to have been left behind in a move. They were recovered, and the data entered.

Most statistical software now has tools for describing the pattern of missing data. Key questions concern the extent and patterns of missing values, and whether the pattern is *monotone* (as described in the next paragraph), as if it is, this can considerably speed up and simplify the analysis.

Missing data in a set of p variables are said to follow a *monotone missingness pattern* if the variables can be re-ordered such that, for every unit i and variable j ,

1. if unit i is observed on variable j , where $j = 2, \dots, p$ it is observed on all variables $j' < j$, and
2. if unit i is missing on variable j , where $j = 2, \dots, p$ it is missing on all variables $j' > j$.

A natural setting for the occurrence of monotone missing data is a longitudinal study, where units are observed either until they are lost to follow up, or the study concludes. A monotone pattern is thus inconsistent with patterns of interim missing data, where some units are observed for a period, missing for the subsequent period, but then observed. Questionnaires may also give rise to monotone missing data patterns when individuals systematically answer each question in turn from the beginning till they either stop or complete the questionnaire. In other settings, it may be possible to re-order items to achieve a monotone pattern.

Example 1.2 Youth Cohort Study (*ctd*)

Table 1.1 shows the covariates we consider from the YCS. There are no missing data in the variables *cohort* and *boy*. The missingness pattern for General Certificate of Secondary Education (GCSE) score and the remaining two variables is shown in Table 1.2. In this example, it is not possible to re-order the variables (items) to obtain a monotone pattern due, for example to pattern 3 ($N = 697$). □

Example 1.3 Asthma study (*ctd*)

Table 1.3 shows the withdrawal pattern for the placebo and lowest active dose arms (all the patients are receiving their randomised medication). We have removed three patients with unusual interim missing data from Table 1.3 and all our analyses. The remaining missingness pattern is monotone in both treatment arms. □

Table 1.1 YCS variables for exploring the relationship between Year 11 attainment and social stratification.

Variable name	Description
Cohort	Year of data collection: 1990, '93, '95, '97, '99
Boy	Indicator variable for boys
Occupation	Parental occupation, categorised as managerial, intermediate, or working
Ethnicity	Categorised as Bangladeshi, Black, Indian, other Asian, Other, Pakistani, or White

Table 1.2 Pattern of missing values in the YCS data.

Pattern	GCSE score	Occupation	Ethnicity	No.	% of total
1	✓	✓	✓	55145	87%
2	✓	.	✓	6821	11%
3	.	✓	✓	697	1%
4	✓	.	.	592	1%