Instructor Guide

to

Principles of Business Forecasting

2ND EDITION

Keith Ord
Georgetown University

Robert Fildes
Lancaster University

Nikolaos Kourentzes
Lancaster University

Wessex Press, Inc.
www.wessexlearning.com

TABLE OF CONTENTS

ABOUT THE AUTHORS

Keith Ord is Professor Emeritus in the Operations and Information Management group at the McDonough School of Business at Georgetown University. He completed his graduate work at the University of London and held faculty positions at the Universities of Bristol and Warwick before moving to The Pennsylvania State University in 1980 and then to Georgetown University in 1999. His research interests include time series and forecasting, spatial modeling and the statistical modeling of business processes. He is a co-author of the 2008 research monograph *Forecasting with Exponential Smoothing: The State Space Approach* and also co-authored Kendall's *Advanced Theory of Statistics*. He has served as an editor of the *International Journal of Forecasting* and is currently on the editorial boards of several other journals. Keith is a Fellow of the American Statistical Association and of the International Institute of Forecasters.

Robert Fildes is Distinguished Professor of Management Science in the Management School, Lancaster University and Founding Director of the Lancaster Centre for Marketing Analytics and Forecasting. He has a mathematics degree from Oxford and a Ph.D. from the University of California in Statistics. He was co-founder in 1981 of the *Journal of Forecasting* and in 1985 of the *International Journal of Forecasting* (IJF). For ten years from 1988 he was Editor-in-Chief of the IJF and remains an associate editor. He was president of the International Institute of Forecasters between 2000 and 2004. His research interests are concerned with the comparative evaluation of different forecasting methods, the implementation of improved forecasting procedures in organizations and the design of forecasting systems. In 1976 he wrote one of the earliest business forecasting textbooks. Though long out-of-print, many of its core ideas have survived the test of time to surface again here in a more modern guise. Robert is a Fellow of the International Institute of Forecasters and of the UK Operational Research Society. In 2014 he was awarded the Beale Medal from the UK OR Society, its highest accolade.

Nikolaos Kourentzes is an Associate Professor in the Department of Management Science at Lancaster University Management School. His background is in Strategic Management, but quickly changed his interests to Management Science, with a Ph.D. from Lancaster University in forecasting with neural networks. He is on the editorial board of the *International Journal of Forecasting* and founding member of the Forecasting Society. Nikos' primary research interest is modeling uncertainty in a business forecasting context, whether that concerns model specification and selection, or ways to make forecasts more reliable and robust. His research addresses forecasting issues of aggregation and hierarchies, model combination, promotional modeling, and supply chain collaboration. He has published multiple forecasting related open-source packages for R, in his attempt to bring current forecasting research to practice.

This Guide replicates some of the material in the Preface but also elaborates on many of the items just covered briefly there. In addition, we provide guidance on the materials that are available on the website.

1 Objectives

Forecasting is not an armchair activity, nor is it an exercise in mathematical formalism, a one-click-and- you're-done computer project, or an uncritical appeal to past experience. Rather, the modern forecaster must be a creative thinker who is able to use available information wisely, draw on the experience of others, use technical arguments when needed and, finally create a computer-based forecasting system that allows management to plan effectively. Probably no such paragon exists, but we should at least aim for an appreciation of all these skills and the ability to work in a team to achieve success.

Our main objectives in developing the materials for the revised edition of this book were as follows:

1. To provide an introduction to both standard and advanced approaches to forecasting, with an emphasis on applications but with sufficient (optional) coverage of more technical aspects for the reader seeking to "look under the hood" [or under the bonnet for UK readers].
2. To develop the material so that it is accessible to an advanced undergraduate in business and to MBA students, while providing additional optional material for those seeking a more theoretical treatment of the major topics.
3. To act as a reference book that will guide those needing to study a particular forecasting topic in depth
4. To provide coverage of the various approaches to forecasting: extrapolative methods, causal modeling and judgmental methods and to develop an appreciation when each approach (or combination thereof) is appropriate.
5. Incorporate the latest research findings to help practicing forecasters carry out their job and to enable students to prepare for a managerial or an analytical career.
6. To identify and present the general principles that underlie good forecasting practice. In this respect we build unashamedly on Scott Armstrong's reference book, *Principles of Forecasting* (2001), although our focus here is on putting these principles to work to produce forecasts.
7. To recognize that there can never be just one approach to forecasting that meets all needs, but that we must invest in "horses for courses."[1]

1 For U.S. readers, this British racing term means that certain horses run better on certain courses. In the context of this book, it means that certain forecasting approaches are better for certain purposes.

1.1 PIVASE

To achieve the last of these objectives, we consider a series of steps summarized by the mnemonic **PIVASE** (pronounced "pi-vase,"). The letters stand for the five stages of forecasting that must be incorporated into the forecasting task:

- ***Purpose*:** What do we hope to achieve by generating the forecast? That is, what plans are dependent upon the results of the forecasting exercise? Over what horizon do we need the forecast to link with planning
- ***Information:*** What do we know and when will we know it? Detailed information is only useful if it is available in timely fashion.
- ***Value:*** How valuable is the forecast? What would you pay to have perfect knowledge of the future event?
- ***Analysis:*** What forecasting methods are suitable for the exercise and do they, capture the data characteristics effectively? How do these methods perform on hold-out out-of-sample data.?
- ***System:*** How can we ensure the methods we have developed are used in planning by those who rely on forecasts? How do we choose between methods?
- ***Evaluation:*** How do we know whether a particular forecasting exercise was effective and valuable?

In short, good forecasting is a challenge. This observation means that the reader must engage in working on real data with all its aberrations. Wherever possible, we have used information from actual governmental or business sources, so that it is possible to go back to the source and update the database. Not only does this activity make forecasting more realistic, but we also get to see how well our forecasts did with the new "out-of-sample" data. With this second edition, we can examine the effects of the 'Great Recession'.

1.2 Principles

At the end of each chapter, we list a set of principles that the forecaster should keep in mind. Some of these—such as "Check the data for outliers"—may seem obvious, but if they are overlooked, the result may be a forecasting disaster. Just as the driver of an automobile must pay simultaneous attention to multiple indicators, the effective forecaster must internalize these principles in everyday practice.

As noted above the impetus for this structure comes from Scott Armstrong's book *Principles of Forecasting*. This project has been enhanced and constantly updated. Details may be found on the website *www.forecastingprinciples.com*.

2 Resources

Active participation in the forecasting process requires a variety of ancillary material. We provide the following:

In the book:

- **Exercises** are provided at the end of all the chapters other than the first and last. These range from simple mechanical problems to aid understanding to real data forecasting exercises, some of which are quite complex and afford rich opportunities for discussion.

- **Minicases** are provided at the end of key chapters and enable more open-ended problem solving. Outline suggestions for solutions are included in the Solutions Manual, but by their nature the minicases do not have neat "answers".
- **Discussion Questions:** The questions are provided throughout the book and on some of the PowerPoint slides to encourage active learning.
- **Class Assessment (Chapter 9):** A challenge for any instructor is to ensure students answer their assignments independently. We provide software to ensure every student faces a unique problem.
- **Glossary**: included in the print version for this second edition, provides definitions of all the technical terms used in the text.
- **Additional data sources**: new and ever more accessible data sources become available all the time. In addition to those listed in Chapter 1, a new resource for world-wide macroeconomic data is *http://www.qlik.com/us/products/qlik-data-market*.
- **Other online forecasting resources:** A listing is provided in the Preface.

On the instructor portal:
https://wessexlearning.com/pages/principles-of-business-forecasting-2nd-ed-instructors-material

- **Solutions Manual:** Solutions are provided for all the exercises. Instructors are encouraged to be selective in the solutions they make available to students.
- **Microsoft PowerPoint™ Slides:** A set of PowerPoint slides is provided for each chapter. Instructors are encouraged to modify these slides to meet their own pedagogical needs.
- **News Bulletin and Errata Sheet:** We've tried hard to avoid errors but no doubt there are some in there. Likewise, there are bound to be occasional news items, specific to the book but of general interest. We will keep this sheet up-to-date so please check the textbook companion site from time to time. If you detect any errors, please send details to *ordk@georgetown.edu*.

On the site for students:
http://www.wessexlearning.org/pobf2e/

- *R* routines and tutorials: In this second edition, we have decided to support the use of *R* as it is freely available and includes many forecasting routines. An online appendix provides the tutorials and access to the programs.
- **Excel Data Sets:** All the data sets used in the book, both in the examples and the exercises, are available (in .xlsx format). These files should read into most forecasting, statistical and econometric packages; if there are any difficulties the file can be opened in Excel and edited as needed.
- **Excel Macros:** Unlike basic statistical and regression programs, forecasting software for extrapolative methods (e.g. exponential smoothing) will often produce different results for the same data set, due to variations in the assumptions and estimation procedures. To enable students to replicate the results in the text we have provided two macros. The first, *Exponential Smoothing Macro* (Exponential Smoothing Macro_v2) enables the user to use exponential smoothing techniques under a variety of conditions involving different trend and seasonal patterns. Please see the accompanying ESM Manual for details. The second, the Trend curves macro (*TrendCurvesMacro_v2*) enables the user to fit a range of non-linear market response models (such as the Gompertz); further details are available in the Trend Manual.
- **Basic Statistical Concepts**: Appendix A (on the website) provides an overview of basic statistical methods used in the text but not explicitly developed there.
- **Forecasting software**: The principal classes of programs available for forecasting, whether general purpose statistical or econometric programs or specialized forecasting software are briefly described in Appendix B (on the website).

3 Structure of the Book

3.1 Introduction (Chapters 1 and 2)

Chapter 1 provides an overview of forecasting, including a variety of situations in which both time series and cross-sectional data can help the forecaster's understanding. The mini-cases provide a forum for discussing the general issues relating to the PIVASE framework.

Chapter 2 introduces the basic statistical tools that are needed in later chapters. The chapter begins with the usual graphical and numerical approaches to summarizing data but then turns to more specialized topics such as the use of transformations, measures of forecasting accuracy and the construction of prediction intervals. The material in this chapter may be supplemented with the topics in Appendix A as needed.

3.2 Extrapolative Methods (Chapters 3–6)

The focus of these chapters is on using extrapolative methods for a single time series to forecast in the short to medium term. The advantage of such methods is that they require only the past values of the variable of interest, so they are particularly useful when a large number of individually low-value items must be forecast. The manager may then focus on those cases that appear to exhibit unusual behavior (exception reporting).

In Chapter 3, we consider regular (nonseasonal) time series and introduce exponential smoothing methods for series without a trend (Simple Exponential Smoothing, SES) or with a trend (Linear Exponential smoothing, LES, or Holt's method). These extrapolative methods remain the "industry-standard" approach to forecasting for large numbers of stock items despite their longevity. The focus is upon the forecasting methods and any theoretical developments are deferred to Chapter 5.

The discussion in Chapter 4 builds upon the previous chapter to include seasonal series, particularly the so-called Holt–Winters methods, although the Exponential Smoothing Macro (ESM) allows a broader mix of seasonal and trend patterns. In this chapter we also consider seasonal adjustment and decomposition procedures that are important for macroeconomic series. Again, the focus is on methods and not theory.

The discussions in Chapters 3 and 4 are limited to forecasting *methods*, which provide point forecasts but do not produce measures of uncertainty. Thus, in Chapter 5 we consider the class of state-space models, which provides a natural framework for exponential smoothing and allows the creation of prediction intervals. The particular state-space framework we describe allows a straightforward mapping from the models to the forecasting methods and addresses the assumptions that are needed to create prediction intervals.

Chapter 6 describes the Autoregressive Integrated Moving Average (ARIMA) models developed by George Box and Gwilym Jenkins, two of the most famous names in forecasting. This methodology is often known as the Box-Jenkins approach to forecasting. The selection of an ARIMA model is based upon data analysis using the autocorrelations. We also develop prediction intervals for these models. There are strong links between the state-space and ARIMA approaches and these are explored. Although the primary focus has been upon forecasts of the level of a time series, the related measures of uncertainty also vary. Hence, we also consider models for changing variances (ARCH/GARCH models), which are widely used in financial analysis. Just as point forecasts change, so do the related measures of uncertainty. These too need to be forecast.

3.3 Statistical Model Building (Chapters 7–9)

Business forecasting involves both the analysis of potentially related time series and the use of cross- sectional databases (e.g., the decisions by banks and companies to extend credit to new customers). Purely extrapolative methods are of limited value in this context and we need to build models with a causal framework. We begin these developments by considering simple linear regression in Chapter 7, where we examine the use of a single predictor variable to assist in explaining the variations in the dependent variable to be forecast. The statistical framework for inference is developed, along with the methods for forecasting using a regression model.

This discussion leads naturally to the topic of multiple regression—the use of two or more explanatory variables—in Chapter 8. This extension requires additional machinery because we must validate both the overall model and its individual components. At this stage it becomes possible to engage in extensive model development and four minicases are provided for students to develop their modeling skills.

Model development and checking are the themes of Chapter 9. Although it sometimes appears that regression modeling is just a matter of downloading the database and running a suitable statistical program, genuine applications involve the careful selection of variables and database building. Even when those phases are complete, the resulting model must be checked to ensure that the final form satisfies—or at least approximates—the underlying statistical assumptions. These models also need to produce forecasts that are more accurate than those provided by simpler alternatives and there is always a risk of overfitting: producing needlessly complicated models that may possess numerically unstable coefficients. A further problem arises with unusual or outlying observations, which may distort parameter estimates and forecasts if not identified and dealt with appropriately.

Stability over time is an essential characteristic of a good forecasting model and the 'Great Recession' of 2007–2009 and the subsequent dramatic changes in crude oil price offers a dramatic example of some of the difficulties of forecasting. This second edition and the updated data sets we have used has allowed us to explore in more depth the important topic of model stability.

3.4 Advanced Methods and Forecasting Practice (Chapters 10–13)

The material in these chapters forms an essential knowledge base for the modern forecaster but might be beyond a first course in forecasting. As the forecasting literature expands, new methods emerge and solutions to new problems are developed. It is important to understand such methods and to recognize when they might offer an advantage over more traditional techniques.

Chapter 10 covers more advanced techniques and is the most difficult both conceptually and technically. Although the forecasting methods described therein may be less familiar than those of earlier chapters, they have clear implications for improved practice and are seeing increased use by modern forecasters. The techniques discussed in earlier chapters assumed a (near-) continuous variable and typically they are appropriate when the focus is upon aggregates (sales, profits, etc.). However, an important class of forecasting problems relates to the behavior of individuals, whether we consider purchasing decisions, credit worthiness or some other individual attribute. Appropriate methods in such contexts are classification and regression trees, and logistic regression.

Neural networks represent a general approach that has attracted a great deal of interest in recent years. We explore their uses in several settings so that their merits and limitations can be identified. *R* routines are provided since these are not typically available in standard software.

Hitherto, all the forecasting methods considered have assumed only a single dependent variable. The fourth principal topic in Chapter 10 extends the possibilities to cover vector autoregressive (VAR) models. Such models have become extremely popular in the econometrics area in recent years. Several appendices to the chapter cover specialized topics relating to the use of VAR models.

Forecasting practice often relies heavily on subjective judgments by experts—and non-experts—in the field, yet forecasting courses often ignore this important topic. In Chapter 11, we consider different approaches to judgmental forecasting and discuss when judgmental inputs can add value. We examine a range of judgmental methods, including Delphi and prediction markets. The choice between judgmental and quantitative forecasting methods is often a false dichotomy, because the correct response is often to use both in combination to capitalize on their respective strengths. This topic offers the instructor some light-hearted relief from the emphasis on statistical modeling and in our experience, it is greatly enjoyed by students – more importantly practical forecasters regard it as essential.

Business forecasting is more than a skill set: we need to develop frameworks for implementation and this is developed in the last two chapters. The first section of Chapter 12 considers the core question of how forecasts and forecasting methods should be evaluated. Procedures for comparing forecasting methods, often a "hot topic" for an organization, are fraught with difficulties. This is a core topic for any forecasting course. The second section examines how forecasts are prepared in organizations through a software- based forecasting support system. We consider what characterizes an effective support system and how such systems and the methods they contain should be evaluated. We then cover three important application areas in detail: operations, marketing, and modeling that is focused on individual customer behavior.

Because the purpose of forecasting is to aid planning, and the reason for planning is to ensure that "things don't just happen," the best models in the world will not help if the forecaster and the user are not communicating properly. Accordingly, Chapter 13 examines the construction of a forecasting system within an organization, concentrating on the interaction between the forecaster and the user of the forecasts. Unfortunately, forecasts are rarely as accurate as their users would like and this final chapter also considers the risks of inaccurate forecasts and how organizations can try to cope with the remaining uncertainties.

4 Course Outlines: Some Suggestions

Student backgrounds, instructor interests and course length all affect course planning in a major way, as they should. These outlines are intended to stimulate initial thinking, it being clearly recognized that none of them would exactly fit the needs of a particular case. The overheads, most of which have been tested in the classroom, show the approach we adopt but the instructor will need to modify many of them to fit the particular course. The following course "types" are fairly common and we provide suggestions for each. Approximate durations are specified to give an idea of the pace and depth of coverage.

- *Short course*: aimed primarily at practitioners with an understanding of statistical methods; focus primarily methodological **[8 hours]**
- *Advanced short course, I*: aimed at experienced practitioners who have a background knowledge comparable to that outlined for the *short course;* focus primarily methodological **[8 hours]**

- *Advanced short course, II*: aimed at experienced forecasting practitioners and demand planners who have a background knowledge comparable to that outlined for the *short course;* focus primarily managerial **[8 hours]**
- *Half-semester course for MBA students*: a balance of technical and managerial considerations; basic knowledge of statistics assumed **[20 hours]**
- *Full-semester course for MBA students or advanced undergraduates in business*: an extended coverage of both technical and managerial issues **[40 hours]**
- *Undergraduate course in the management sciences and statistics*: Basic material in chapters 2 and 7 could be excluded. Chapter 9 would be emphasized as many courses just focus on the technical features of regression.
- *Undergraduate course in business analytics*: Forecasting is key component of any business analytics course. The course would emphasize data mining elements. Standard R code included here is a major benefit **[40 hours]**
- *Master's degree programs in applied statistics, economics, management science or business analytics:* greater focus on technical issues

Clearly there are many variations on these themes and an instructor would not necessarily cover all the topics listed under a particular heading. Rather, the following table is included to stimulate course planning. ("BR" denotes background reading prior to the course; "☑" denotes likely material for inclusion and a "?" denotes material for possible inclusion.)

Topic	Book	Short course	Advanced short course I	Advanced short course II	Half-semester MBA	Full-semester MBA/UG	UG MS/ Stat	UG Business Analytics	Applied MS
Introduction	Chapter 1	BR	BR	BR	☑	☑	BR	☑	☑
Basic Tools	Sections 2.1-2.6; Appendix A	BR	BR	BR	BR	BR	BR	☑	BR
Forecast evaluation	Sections 2.7-2.8	☑	BR	BR	☑	☑	☑	☑	☑
Exponential smoothing-nonseasonal	Sections 3.1-3.9	☑	BR	BR	☑	☑	☑	☑	☑
Exponential smoothing-seasonal	Sections 4.1, 4.2, 4.6, 4.7	☑	BR	BR	☑	☑	☑	☑	☑
Seasonal decomposition	Sections 4.3-4.5	?	☑				☑		☑
State-space models	Sections 5.1-5.5		☑			?	☑		☑
ARIMA models	Sections 6.1-6.8		☑			?	☑		☑
Simple linear regression	Sections 7.1-7.8	BR	BR	BR	BR	BR	BR	☑	BR
Multiple regression	Sections 8.1-8.6, 9.1-9.3, 9.4, 9.10	☑	☑	☑	☑	☑	?	☑	☑
Multiple regression-advanced	Sections 9.5-9.9	?	☑		?	☑	☑	☑	☑
Individual behavior	Sections 10.1-10.3		☑	☑		☑		☑	☑
Neural nets (NN)/VAR	Sections 10.4, 10.5		☑					?	☑
Judgment-based forecasting	Sections 11.1-11.7	?	?	☑	☑	☑		☑	?
Putting forecasting methods to work, I	Sections 12.1-12.3	☑	☑	☑	☑	☑	☑	☑	☑
Putting forecasting methods to work, II	Sections 11.6, 12.4, 12.5		☑	☑	?	☑			☑
Forecasting in practice	Sections 13.1-13.4		?	☑	☑	☑			?

5 Materials on the Website

In addition to this Instructor's Guide, the following materials are available on the book's website:

- **Data sets**: Each data set cited in the text is available as an Excel 2010 file (.xlsx)
- **PowerPoint:** Slides are provided for each chapter (.pptx)
- **Excel macros:** Exponential Smoothing Macro and Trend Curves Macro (macros are .xlsx and the manuals are .docx)
- **Solutions:** Provided by chapter (.docx)
- **Appendix A:** Covers Basic Statistical Concepts, available as PDF file (.pdf)
- **Appendix B:** Covers Forecasting Software, available as PDF file (.pdf)
- **Appendix C:** Forecasting in *R*: Tutorial and Examples
- **Errata sheet:** Word file that will be updated as needed (.docx)
- **News Bulletin:** A listing of occasional new developments that relate to material in the book and would be of general interest. Users are invited to contact the authors with suggestions.
- **E-version**: For students, each chapter is available on the web for purchase allowing the instructor to design their own individual course.