


Two Notable New Forecasting Texts

Principles of Business Forecasting by Keith Ord & Robert Fildes

Forecasting: Principles and Practice

by Rob Hyndman & George Athanasopoulos

Reviewed by Stephan Kolassa



Principles of Business Forecasting (PoBF)

Keith Ord of Georgetown University and Robert Fildes of Lancaster University are well-known figures in forecasting. Both have published extensively and have considerable experience in educating tomorrow's forecasters and in hands-on projects in business forecasting. Now they have joined forces to write a new textbook: *Principles of Business Forecasting* (PoBF; Ord & Fildes, 2013), a 506-page tome full of forecasting wisdom.

Coverage and Sequencing

PoBF follows a commonsense order, starting out with chapters on the why, how, and basic tools of forecasting. Exponential Smoothing comes next, the workhorse of business forecasters, as described in the methods tutorial in this issue. The treatment of smoothing is spread over three chapters: one for non-seasonal data, one on seasonal smoothing (other approaches to decomposing seasonal time series are also touched upon here), and one on the relationship between smoothing and state space models. State space models, which allow prediction intervals to be calculated for exponential smoothing procedures, are a relatively recent development, and the authors succeed in presenting this normally complex topic simply and clearly.

PoBF then moves on to regression-type forecasting methods, with a chapter on ARIMA (Box-Jenkins models), a chapter on simple linear regression, a chapter on multiple linear regression, and another on model building.

The model-building chapter provides useful extensions of regression modeling to cover indicator (dummy) variables; lagged variables and errors; procedures for variable selection (such as stepwise regression); the problems of multicollinearity, changing variance, and structural change; nonlinear terms and models; and, finally, outliers and leverage points.

We next find a chapter on advanced methods of forecasting (predictive classification, neural networks, and vector autoregression). For this complex material, the authors' presentation is down-to-earth and informative, with helpful numerical examples of the concepts.

The final three chapters give the text a true practitioner focus. They examine judgmental forecasting, forecasting in the larger business context, and the processes and politics of forecasting. An extensive discussion of forecasting support systems offers valuable perspectives to the business forecaster. In these chapters, Ord and Fildes have given us an up-to-date overview of the way forecasting is conducted in the business organization.

Level of Sophistication

PoBF is not always an easy book – but then again, neither is forecasting easy. It contains a fair amount of mathematics and statistics in explaining the forecasting methods, but admirably avoids the intricate details for which we rely on modern statistical forecasting packages, and provides the interested reader with pointers to other literature.

The book is not meant to be a compendium of academic knowledge on forecasting. Rather,

it is aimed at practicing business forecasters, who should have a certain knowledge about the underlying statistics in their forecasting package, but who will rarely need to use mathematical formulations to write new procedures (e.g., to maximize likelihoods by hand).

That said, how much math and statistics will a reader need to profit from *PoBF*? Actually, surprisingly little. Higher mathematics (linear algebra and calculus) are used very sparingly, and one can skip these derivations with little loss. Most of the book – specifically, everything apart from the ARIMA, regression, and model-building chapters – is very accessible to readers without a background in statistics. For the methodological chapters, some background in statistics is extremely helpful, at least to the extent of understanding Q-Q plots, p-values, null hypothesis significance testing, and various distributions (e.g., z, t, F, χ^2). When operational, the book's website will contain a primer on basic statistical concepts along with the data sets used and a glossary.

I found *PoBF* to be extremely readable. In particular, I like the way it spreads the more technical topics over multiple chapters. The discussion questions, minicases, and exercises are very helpful in understanding each chapter. Each chapter closes with a list of principles that succinctly summarize the takeaways. They credit this idea to Scott Armstrong's *Principles of Forecasting* handbook (2002), a volume they refer to frequently.

Summary

PoBF is perfect for the practicing (as opposed to the academic) forecaster. It draws attention to the fact that forecasts, targets, and plans are three very different things. It describes biases within organizations due to varying incentives of the different functional areas. It discusses surveys about the biggest problems forecasters face, and offers recommendations on how to address them. It explains how to select forecasting software and how to marry judgment with statistical forecasts.

Is there anything here I would have done differently? Of course. I found it unclear why certain topics (unit roots, cointegration, simulations to assess uncertainty) ended up in appendices or nontechnical chapters when they would have fit well into the advanced-methods chapter. I would have liked more attention paid to three topics most forecasters seem to agree are useful to improving forecasts: using simple methods as benchmarks, combining forecasts, and dampening trends. All three are mentioned, but the casual reader can be excused for missing them. Lastly, cost of forecast error and forecast value-added concepts are lacking.

These are extremely minor criticisms. I wholeheartedly recommend *Principles of Business Forecasting* to every business forecaster, to everyone who aspires to become a business forecaster, and to everyone who is training aspiring business forecasters. (I am hoping that forecast managers will read this book as well, but I'm probably being overly optimistic.) You will certainly profit from reading it. I know I did, and I know I will do so again when I reread it.



Forecasting: Principles and Practice (FP&P)

Rob Hyndman and George Athanasopoulos, both at Monash University, are also well known in the forecasting community. They have collaborated to produce an online textbook, *Forecasting: Principles and Practice (FP&P)*, which is explicitly intended as a replacement for *Forecasting: Methods and Applications*, the earlier textbook by Makridakis, Wheelwright & Hyndman (1998). *FP&P* is available free of charge at <http://otexts.com/fpp/>. It is not finished as yet, but the authors hope to finalize it by the

end of 2012 and to publish a print version in 2013.

Organization

The book is at first conventionally organized, beginning with “Getting Started” followed by “The Forecaster’s Toolbox,” which supplies the basic graphics, numerical data summaries, simple forecasting methods, and accuracy metrics. Next comes a chapter on judgmental forecasts, which at this point in time is only a placeholder. Most academic forecasters would take up judgmental forecasting after having dealt with statistical approaches such as ARIMA and Exponential Smoothing. However, the first thing a practicing forecaster witnesses in his company is the frequency of judgmental adjustments to statistical forecasts. Thus, I applaud the ordering of chapters in this way and hope that in addition to the discussion of “purely” judgmental approaches (such as Delphi), they will examine judgmental adjustments to statistical forecasts as well.

The book continues with simple regression, multiple regression, and then time-series decomposition, exponential smoothing, and ARIMA. All of these chapters are complete (though exercises have not yet been added), but the next chapters – on advanced methods such as dynamic regression, neural networks, further forecasting methods, demand forecasting, and using forecasting methods in practice – have yet to be finished.



Reliance on R

As with the Ord/Fildes text, *FP&P* requires only a modest mathematical and statistical background. Some sections use matrices or

calculus, but these parts are clearly marked and can be skipped. An introductory statistics course would be necessary in order to derive the full benefit from the ARIMA and regression chapters. Since forecasting is inherently a computational science, we see a lot of equations in some places, especially the sections on Exponential Smoothing and ARIMA. Throughout *FP&P*, the authors use the statistical environment R, a free, open-source statistical computing environment that can be used for forecasting. In fact, Hyndman and Athanasopoulos wrote an R-package (*fpp*) as a companion to this e-book, and they provide an appendix with a tutorial, “Using R.” Please note that R is a viable alternative to costly proprietary forecasting software, especially for students or forecasters on a budget, as well as having appeal to real power users since every function can, in principle, be accessed, examined, or modified. However, R lacks the polished graphical user interface of proprietary software and can be intimidating for the casual user (Kolassa & Hyndman, 2010). Yet the possibility of copying the R code snippets from the e-book and directly pasting them into an R console to see the output and to experiment with R is extremely appealing, and will certainly invite more hands-on practicing; additionally, one can easily use the R code snippets as blueprints to apply to one’s own data.

Unfortunately, the “Using R” tutorial will be of limited help for R novices; for instance, it inexplicably does not mention the “?” command for getting help, or “??” for free search, or the various ways of getting help online that are crucial for the new user. One exercise in the book asks the user to explore some of the other time series in certain R packages, without mentioning how one would find out what data sets there are in a package. When asked about this, Rob Hyndman commented that it was a conscious decision in order to force the student to learn how to use the index; but I am afraid that raising such hurdles even before the user can start on the

actual exercise (which, remember, was to examine some time series, not to learn about the help section in R, or the R-help mailing list or online R tutorials) will only serve to frighten away anyone who has not already worked with R before. R already has a steep learning curve; there really is no need to add to it.

Graphics

FP&P heavily emphasizes graphical methods. One section in the “Toolbox” chapter is entirely devoted to different ways of plotting data in time plots, seasonal plots, seasonal subseries plots, scatterplots, and scatterplot matrices – and all of these come with the corresponding R code. This section is one of the most helpful parts of the book, and every forecaster would profit from reading it. Looking at one’s data from various angles always aids in understanding the underlying dynamics. Are there trends? Are seasonal patterns stable over time? Are there outliers? Or entire patterns of suspicious data points? One example involves multiple consecutive weeks of zero airplane passengers flying from Melbourne to Sydney because of a labor dispute. Such periods should be accounted for before using these historical data to forecast future passenger levels, and this pattern clearly shows up in the time plot. The next thing I would like to see included in the section on graphics would be ways to examine entire data sets of hundreds or thousands of time series simultaneously, and the e-book format would in fact allow such an updating.

Omissions

FP&P supplies almost no references at all, which is a pity. A lot of material is only touched on, and the reader could certainly use guidance on how and where to delve more deeply. Not even the IIF, the IJF, or *Foresight* is mentioned, but this can easily be rectified in an e-book.

As of this writing (September 2012), the book is not yet comprehensive and there are points where it obviously could be expanded. For

instance, the section on numerical data summaries explains the mean and the median, the standard deviation, and the interquartile range – but it does not mention that the mean and the standard deviation are strongly influenced by extreme values and that the median and the interquartile range are much more robust. Similarly, there is a good summary of what multicollinearity is in the chapter on multiple regression, but no numerical example or guidance on how to actually detect it, e.g., using Variance Inflation Factors or other established collinearity diagnostics. In addition, some chapters are apparently finished but lack exercises; this will hopefully be remedied by the end of 2012.



eBook Format

The e-book format has advantages and disadvantages. For one, I don’t need to lug around a big hard-copy text – I can just open my laptop or my smartphone. The authors can easily correct errors or add material. I can simply copy and paste R code from *FP&P* to my R console. References are implemented as direct hyperlinks to online versions of papers or to a book’s page at Amazon.com. You can search for a term in *FP&P* using Google with a “site:otexts.com/fpp” search.

Unfortunately, no downloadable version of the book is available for those long train rides with limited data connectivity, but the

authors promise that an offline version will be offered together with the print version. And it is hard to make notes in an online e-book the way one could do in the margins of a paper textbook, but an offline e-book in a standard format may allow this.

Summary

Does *Forecasting: Principles and Practice* deliver on its stated goal of replacing Makridakis, Wheelwright & Hyndman's 1998 book? Upon opening the older book, one immediately notices that much of the material and also the data sets have been transferred to the new book. Some material has been cut, e.g., the chapter on long-range forecasting and (unfortunately) the references and the pointers to resources and forecasting organizations. On the other hand, much material is updated based on what forecasters have learned over the last 15 years. Examples include the sections on forecast-accuracy measurement, neural networks, and everything on R. Thus, *FP&P* really is an (almost) perfect successor to the older book.

The bottom line: this is the book for you if you are happy with short, to-the-point explanations, if you are able to find your own way through the literature to clarify unclear points, and if you are comfortable

using R. Everyone else would certainly also profit from leafing through *FP&P* – the price is hard to beat, after all – but should prepare for a strenuous read. I have already book-marked *Forecasting: Principles and Practice* and will definitely return to it.

COMPARISON

We now have two new books on the market, which both aim to attract the forecasting student and the practitioner. How do these books compare?

The first note of comparison is the price tag. *Principles of Business Forecasting* retails for about \$150 U.S. on Amazon, whereas *Forecasting: Principles and Practice* is free in its online version. We don't know yet how much the paper or offline versions of *FP&P* will cost.

PoBF offers more detailed explanations and illustrations of the methods, and, while longer, examines the methods from different angles. Overall, it seems to be more gentle in its didactics.

FP&P relies exclusively on the free software package R. *PoBF* unfortunately does not mention this free computing environment, only referring to commercial statistical or forecasting packages such as Forecast Pro, Minitab, SAS, SPSS, EViews, and PcGive/OxMetrics (except for a brief mention of the free specialized econometric package TRAMO-SEATS).

Even apart from the price differential, it is much easier for the reader to load the “fpp” package and copy R code from the book to the R console than it is to go to the *PoBF* website, download a data set, fire up his or her favorite forecasting program, and load the dataset. Thanks to R, the hurdle for actually earning some hands-on experience with forecasting is a lot lower with *FP&P* than for *PoBF*.

PoBF gives references for further reading after each chapter, which allows one to learn more about specific topics. *FP&P* gives very few references, and readers who would like



to learn more about, say, invertibility of an ARIMA process would need to hunt for relevant literature themselves.

In comparing an e-book and a paper textbook, differences are inherent in the two media. *FP&P* can and will change, and even those sections that are already finished can be rewritten or expanded. Errors and typos can be corrected (whereas *PoBF* will publish a list of errata at its website), new research can be worked into the text, and readers can suggest specific topics to be addressed. There are no page limits as there were for *PoBF*. On the other hand, as long as no offline version is available for download, the user cannot annotate his *FP&P* copy and will need an Internet connection to read it. In contrast, I was perfectly able to read my copy of *PoBF* during my vacation in an idyllic – and blissfully Internet-free – seaside cottage.

One issue on which the two texts do not differ is the level of math/stat knowledge they demand. Most of both books is accessible with no particular math/stat background at all; parts of both books are more easily read if the reader knows matrices and calculus; and readers will need a bit of inferential statistics knowledge to fully profit from the chapters on regression and ARIMA.

A book for the practicing forecaster also needs to cover the nonscientific “softer” aspects of forecasting: what kind of larger processes make use of forecasts, what are the advantages and disadvantages of locating the forecasting function in various organizational units, how about judgmental adjustments to statistical forecasts, and what kinds of incentives do people around the forecaster face? *PoBF* has very helpful chapters on all this, while *FP&P* as yet only has placeholder chapters on these issues.

Each book has its strengths, and I would keep both close to hand; every forecaster should at least look at both of them. You may want to start out with *Forecasting: Principles and Practice* if you are a cash-strapped student or a forecaster with no budget for a book and

a software package, if you are comfortable with rather terse exposition, if you are happy using R, if you like to play around with your data and your forecasts while you read, and if you have a solid Internet connection up and running. Conversely, *Principles of Business Forecasting* may be the better choice if you have enough budget (or a library copy of the book and a campus license for a forecasting package), if you prefer a gentler exposition, if you like to just read your book and/or have the discipline to work with data and software (even though this is slightly more of a hassle), or if your favorite reading/studying haunt has no Internet connection. But whichever book you start with, it makes a lot of sense to read the other book, too.

REFERENCES

- Armstrong, J. Scott, ed. (2002). *Principles of Forecasting: A Handbook for Researchers and Practitioners*. New York, NY: Springer
- Hyndman, Robert J. and Athanasopoulos, George (2012). *Forecasting: Principles and Practice*. <http://otexts.com/fpp/>. Accessed 09/04/2012 to 09/07/2012
- Kolassa, Stephan and Hyndman, Robert J. (2010). Free Open-Source Forecasting Using R. *Foresight*, 17, 19-24
- Makridakis, Spyros; Wheelwright, Steve and Hyndman, Robert J. (1998). *Forecasting: Methods and Applications*. 3rd edition. Hoboken, NJ: Wiley
- Ord, Keith and Fildes, Robert (2013). *Principles of Business Forecasting*. Mason, OH: Cengage Learning.



Stephan Kolassa is a Senior Research Expert at the SAP Center of Excellence Forecasting and Replenishment and an Associate Editor for *Foresight*.
stephan.kolassa@saf-ag.com