

instance, the X11 adjustment is said to be “robust” (p. 1), but readers are not told what this robustness relates to. When discussing model-based seasonal adjustment, the authors then question the “robustness of the [model-based] estimation methods in the case of very noisy series” (p. 9), still without explanation. They also state “the relative weakness of statistical theory for non-stationary series” (p. 9), forgetting the important work of Bell (1984) on signal extraction in non-stationary series. These statements made without justification aroused my curiosity about their motivation.

Of course, when reading the description of a program that was developed several decades ago, one can regret that the use of statistical theory is not always up to date. It has to be kept in mind that X-11 was built by highly skilled researchers who had to solve a quantity of problems related to seasonal adjustment in large-scale applications, and that the majority of these problems were well beyond the statistical knowledge available at the time. Without any doubt, the success of X-11 shows that the solutions that have been implemented were relevant. With this book, D. Ladiray and B. Quenneville detail for us all the different operations and they provide the user community of X-11, and derived products, with an exhaustive manual that I believe will be very useful.

References

- Bell, W. (1984). Signal extraction for nonstationary time series. *Annals of Statistics*, 12, 646–664.
- Bell, W. R., & Hillmer, S. C. (1984). Issues Involved with the Seasonal Adjustment of Economic Time Series. *Journal of Business and Economic Statistics*, 2, 291–320.
- Khintchine, A. (1934). Korrelationstheorie der stationäre stochastischen Prozesse. *Mathematische Annalen*, 109, 604–615.
- Wiener, N. (1930). Generalised harmonic analysis. *Acta Mathematica*, 55, 117–258, Reprinted in Selected Papers of Norbert Wiener, MIT Press, Cambridge, Massachusetts, 1964.

Christophe Planas
 Joint Research Centre of
 European Commission
 Ispra
 Italy

PII: S0169-2070(02)00035-3

Principles of Forecasting: A Handbook for Researchers and Practitioners, J. Scott Armstrong (Ed.), (2001), Boston: Kluwer Academic Publishers, 849 pages. Hardback: ISBN: 0-7923-7930-6; \$190, £133, €210.00, Paperback: ISBN: 07923-7401-0; \$95; £66.50, €105.

1. Introduction

At the start of a new millennium it seems appropriate that we should reflect upon and summarize our current knowledge of forecasting. The *Principles of Forecasting* handbook aims to provide this summary. More precisely, its objective is to: ‘summarize knowledge in forecasting as a set of principles. These “principles” represent advice, guidelines, prescriptions, condition-action statements and rules’. The production of the book was a huge project involving 40 experts, who wrote and peer reviewed the book’s 30 chapters, with a further 123 experts providing additional advice. A comprehensive dictionary of forecasting is also included in the handbook. Because of the wide range of forecasting methods covered, four reviewers were asked to discuss: i) the book in general and ii) the areas of the book that were closest to their specialist interests. Their reviews are set out below.

2. Quantitative non-causal methods (Reviewed by Keith Ord)

The first attempt to provide a systematic set of statements on forecasting perhaps took place in the cave of the Oracle at Delphi. By all accounts, it was a somewhat chaotic affair as the wind constantly disturbed the leaves on which the forecasts were written. By contrast, *Principles* provides a structured and well-conceived framework for forecasting that is firmly anchored by past research and therefore is resistant to the changeable breezes of prognosticating fashion. Scott Armstrong and his army of co-authors deserve the sincere thanks of the forecasting community for the successful completion of an enormous project.

Principles is not, and was not intended to be, a “how-to” book. Plenty of those exist for specific approaches to forecasting, and the reader needs a good grounding from that literature before consulting this book. As is freely admitted by the editor, support for the various principles is uneven. Some are demonstrated in the text, others are supported by citations and some are plausible assertions. The forecaster making use of these principles needs to take the time to evaluate them, both in general and in the context of the task in hand, before attempting applications.

2.1. Quantitative methods

My assignment was to examine those parts of the *Principles* relating to “quantitative non-causal methods” so a natural place to start was the section on “Selecting Methods” by Scott Armstrong. This chapter might more usefully have appeared right after the introduction and focused exclusively on *selecting an approach to forecasting*. The researcher would then be able to answer the broader questions before confronting detailed implementation issues. Scott does consider this issue, and there is a ‘route map’ on

the inside front cover, but greater emphasis would have been desirable. I found *Principles* to be rather like a travel guide: very good at describing and evaluating the attractions once you have arrived, but less effective at telling you whether you should be there in the first place.

The starting point for simple quantitative methods is the question “Should we use domain knowledge or statistical analysis to aid model selection?” My review is based on the assumption that you would answer yes. Unfortunately, *Principles* does not always agree. Scott Armstrong’s paper on “Extrapolation of Time-Series and Cross Sectional Data” does a good job of enunciating general principles relevant to simpler methods, but does not recognize the benefits of model-based inference as a basis for selecting forecasting procedures. The work of Box and Jenkins is discussed only very briefly (page 231) and the contributions of Andrew Harvey and Jeff Harrison and their co-authors on structural (or state-space) modeling not at all. Granted, one of Scott’s cornerstone beliefs is that statistical modeling has been oversold. Nevertheless, the lack of any effective discussion of how to combine domain knowledge and statistical insights is a missed opportunity.

2.2. Specific themes

Fred Collopy, Monica Adya and Scott Armstrong do a nice job in their discussion of rule-based forecasting, although some of their principles (e.g. on decomposition) could usefully be considered at earlier stages of the overall forecasting process. The strength of the rule-based approach is the use of domain knowledge; the method has the potential to become even more valuable if researchers treat statistical inference and modeling as complementary to their efforts rather than as competition.

Bill Remus and Marcus O’Connor provide a balanced evaluation of artificial neural networks

(ANN). As is evident from their chapter and other work in the area (e.g. Balkin and Ord, 2000) the uncritical application of ANN can produce inferior forecasts and such methods must be used selectively.

The chapter by George Duncan, Wil Gorr and Janusz Szczypula gives an effective guide to Bayesian pooling methods, primarily in the context of cross-sectional studies. The title of this section “Analogies” is something of a misnomer and indeed, *Principles* does not appear to contain a discussion of forecasting by analogy as such. Nevertheless, pooling is an important topic that should receive increasing attention, and these authors provide a useful and thoughtful evaluation.

Finally, the chapter by Nigel Meade and Towhidul Isham examines forecasting diffusion processes. They do an excellent job of marshalling the evidence in support of their stated principles and use these principles to guide the prospective user towards good practice.

2.3. Related issues

Several chapters cover more general areas but a number of these deserve particular attention as they include topics that have been considered in some depth in the context of extrapolation: forecast evaluation, measurement of uncertainty and combining forecasts.

Scott Armstrong provided the chapters on evaluation and on combining. The evaluation chapter starts with a number of strong principles, but some of the later principles are debatable. For example, “Do not use root mean square error for comparisons across series” is certainly correct but it is later rendered in Exhibit 10 (page 465) as “Do not use RMSE” a much more questionable statement. The chapter on combining is strong on empirical evidence and deserves careful consideration from practitioners.

Chris Chatfield assesses the current state of

measuring forecast uncertainty and quite properly concludes that this is a particularly weak spot in current practice. His principles provide a constructive road map for future activity both in research and in software development.

2.4. How to read the Principles

Few readers will have the time to read the *Principles* from cover to cover, nor is that to be recommended. After reading the Introduction, I would recommend that the reader move on to the chapter on software by Len Tashman and Jim Hoover. Their evaluation of the diffusion of principles into forecasting packages leads to several conclusions: (1) many spreadsheet systems are inadequate for most purposes, (2) the best business forecasting packages do quite well but still have a way to go, and (3) that the set of principles deserves serious consideration when developing a forecasting system. The companion chapter by James Cox and David Loomis on the diffusion of principles in forecasting books reinforces point (3).

Once convinced of the need to be principled, the chapter on method selection should be reviewed, to provide a signpost for the next stage of development. At that point, readers will consult the chapters most directly relevant to their chosen approach. In implementing a forecasting system, the practitioner will want to examine the summary chapter written by Scott Armstrong and the accompanying checklist. The associated website will considerably enhance the value of *Principles* in applications.

2.5. Conclusion

Not everyone will support every principle, nor is that to be expected. Indeed, I cannot resist suggesting that some of the principles, such as “Use all relevant data, especially for long-term forecasts” [page 220] are sufficiently Delphic to make the old Oracle smile quietly to herself.

However, most of the principles are well documented and provide guidelines for either action or introspection as ways to improve business forecasting. The *Principles* and its associated website should have a substantial impact on forecasting practice for many years to come.

3. Econometric methods (Reviewed by Lars-Erik Öller)

3.1. *Econometric Forecasting*

Econometrics is a quite technical subject these days. If you open a textbook in econometrics you will find an account of the methods that are recommended for use. Criteria for what is a good or bad model are given according to the rules of statistics, such as unbiasedness, consistency, minimum error variance, uncorrelated and homoscedastic errors, constant structure etc. However, what really yields accurate forecasts is rarely discussed¹. The only books I know of that come close to this are the two recent works by Michael Clements and David Hendry (Clements and Hendry, 1998, 1999). The chapter on Econometric Forecasting by P. Geoffrey Allen and Robert Fildes partly relies on some of the results reported in these two books.

A point that the authors of this chapter want to stress is that not much of what is generally considered as good econometrics has been shown to really lead to much improvement in forecast accuracy. They discuss questions that can be raised when constructing forecasting models and try to answer them, based on, alas often scant, evidence of effect on forecast accuracy. The discussion is illuminating and covers a huge literature, ending up in a practical list of 23 advisory rules for econometricians,

together with conditions for applying the rules and evidence in favor and against it.

The authors have a point when they stress the scarcity of forecast evidence in favor of econometrically adequate models. They recommend economic theory, especially if the number of observations is small. One would like to add that good econometric modeling practices would also be advisable, even if they are not shown to be superior in any currently available forecast comparison. It is a welcome achievement by the forecasting society that more and more econometricians test their model outside the sample. On the other hand, one should not exaggerate forecasting performance at the expense of statistical considerations. As long as we are talking about macroeconomic data, it is a well known fact that their reliability is limited. To rely exclusively on a few forecast errors may not be a wise strategy.

With this clarification I'm ready to accept, and indeed welcome, all but one of the practical rules. Rule 14 (not numbered in the text) says: "Initially estimate equations in levels, not in first differences". As admitted by the authors, this is a controversial area. Given that, one would like not to see too categorical rules. The first reservation to make here is that, if there are unit roots in the variables, the rule makes sense only if we are talking about autoregressive models. As the principle stands now, it could be understood as endorsing spurious regressions. Even the evidence the authors put forward is rather mixed on this point. It is true that unit root tests are still rather weak, but my advice would be: start by looking at graphs of the time series and ask yourself if they may have unit roots. If so test, using the procedures mentioned here. If they have unit roots, test if they are cointegrated, but do not stop looking at graphs and using your good judgement. If they are, proceed by defining an error-correction model, which certainly can be estimated in levels. According to the results in Clements and Hen-

¹Another purpose of a model may be to be used as a tool for policy scenarios, in which case the design of the model should be different than if forecasting is the main purpose.

dry (1995), this strategy poses a smaller risk than not imposing cointegration and allowing the presence of spurious levels terms. Many economic time series, such as production and prices, can be distinguished as integrated processes just by optical inspection and then the categorical advice of *always* starting in levels seems not to be warranted.

Talking about autocorrelation in the error term, one might add that such a forecast is not rational, borrowing a term from the expectations literature. Removing autocorrelation by introducing new variables and lags illustrates the trade-off between adequacy and overparametrization, a trade-off that appears with all adequacy statistics, and shows how hard it is to give categorical advice—and how much econometrics is an art. In this connection one might have wanted to see a rule on lag parsimony in VAR models, as compared to single equations. In general, more lags are needed in the latter than in VARs, due to feed-backs in VARs, see Enders (1995, p. 368).

Summarizing, one would wish that the knowledge invested in this chapter would reach all econometricians, young and old, who work with time series. You will not find it in this compact form anywhere else.

3.2. *Econometric Models for Forecasting Market Share*

For business management, keeping a close watch on market share (MS) seems to be as important as monitoring sales and profits. Conditional forecasts of MSs are expected to produce decision scenarios, where at least some of the causal variables should be discretionary policy alternatives (price on own products, advertising, discounts, etc.). The way to produce such forecasts is to model the market using econometric techniques. If the only thing that matters is *forecasts*, a univariate autoprojective model is the main competitor, the simplest case being a naïve forecast.

The section on econometric models for forecasting market share by Roderick J. Brodie and Peter J. Danaher first outlines the problem by introducing the reader to the competitive market system. The MS models used by practitioners are either linear, half-logarithmic or a ratio of two multiplicative expressions. The average reader will hardly get a feel for these models only by reading this chapter, because they are not explained verbally and the mathematical definitions are unclear and include typos. But then, the purpose of the text is not to be a general introduction into the MS subject. Other methods are briefly mentioned: conjoint analysis (cf. Ch. 5) and neural networks (NNW) (Ch. 8). About NNW it is said that they have not been shown to produce consistently better forecasts than more conventional models, except in a few cases, one of them outlining non-linear behavior. But NNW are designed precisely to model nonlinearities! It is also said that NNW are computer intensive, which is true, but a smaller problem than the fact that, like all nonlinear methods, it is *observation-intensive*, thus excluding applications where only a modest number of observations are available.

The authors refer to the principles laid down by Allen and Fildes (Ch. 11) (AF), applying them to this particular field and adding the special features that cannot be contained in a general set of rules. The main recommendation is to use econometric models when degrees of freedom allow it and the current market effect of discretionary variables is strong. It is worth noting that they disagree with AF² on the choice of estimation method. For MS applications, the SUR method, when applicable, is recommended in place of OLS, which was preferred by AF as a general rule. In this connection one could add that SUR becomes identical to OLS if the

²The authors also disagree with AF on using constant coefficient models.

explanatory variables are the same in all equations.

A problem often mentioned, but not solved in this chapter, is how to model the reaction of competitors. An outsider would suggest a simultaneous model or, if countermoves are delayed by at least one period, a VAR model. Non co-operative game theory may also be useful.

As in AF, the chapter ends in a list of principles with arguments pro and con, mainly based on empirical studies. The authors are well aware of the fact that there are very few forecast studies on which to base conclusions. Often, naïve no-change forecasts seem to outperform those generated by econometric models. What one doesn't know is how these econometric models were specified. In older studies, and areas in the periphery of mainstream econometrics, one sometimes finds dubious model specifications. If nonstationary variables were modeled in levels, the regressions may be spurious, in which case the naïve random walk is in fact closer to a correctly specified model than that of an OLS in levels. Often diagnostics are so poor that it is impossible to judge on the adequacy of a suggested econometric model. Finally, if the data have structural breaks the difference adapts very fast, again favoring the naïve forecast.

Despite some shortcomings, this is a practical chapter that should help both applied researchers and practitioners involved in forecasting market shares.

4. Judgmental forecasting (Reviewed by Janet A. Sniezek)

Although much forecasting actually takes place in groups (Armstrong, 1985), the Handbook has much more to say on individual than group judgment processes. Several of the Hand-

book's authors do advise obtaining forecasts from multiple persons. Collopy, Ayda, & Armstrong report that few developers of expert systems use multiple experts though this is what they prescribe. In the "Combining Forecasts" chapter, Armstrong reviews the advantages of combining forecasts from several experts. Similarly, MacGregor recommends averaging a set of judgments in the application of decomposition to improve forecasting accuracy. The suggestion to rely on multiple forecasters is good advice, as averaging multiple forecasts will generally improve accuracy. Stewart identifies another advantage of multiple forecasters: assessing agreement among them provides an indicator of forecast reliability.

However, there are two additional—and arguably more important—reasons for researchers and practitioners to be interested in forecasting by multiple persons than the gains in accuracy achieved by averaging their respective forecasts. For starters, actual organizational practices typically involve behavioral aggregation by the group of forecasters themselves, making it imperative to study interacting groups. It is not possible to fully understand group behavior by studying the behavior of individuals (cf. Sniezek, 1992), and the principles derived for individual forecasters may not apply to groups.

The second reason for caring about group forecasting is to tackle the more significant problem of reducing systematic bias. If individual judgments are systematically biased, so too will be the average of the judgments from multiple persons. Only through some form of communication among group members is it possible for them to produce group forecasts that are sufficiently different from their average forecast that systematic error is reduced. Rowe and Wright produce evidence that the use of the Delphi technique with groups can successfully reduce bias. They explain the logical basis for this success, and detail the conditions necessary for it. Their chapter is a most welcome break

from a confusing literature all too full of atheoretical competitions among techniques.

Despite the impressive results reported by Rowe and Wright, studies using the Delphi technique tend to show less change from average individual estimates than can occur with interacting groups. For example, Sniezek and Henry (1989) found over a 20% reduction in standardized bias for judgments from interacting groups compared to the average of members' individual judgments. Overall, 30% of group judgments were better than the judgment of the most accurate member. Perhaps most surprisingly, 15% of the groups were more accurate than their best member because they produced "out-of-range" judgments, i.e., group judgments that were higher (or lower) than all members' individual judgments. Such dramatic improvements in judgment accuracy are simply not possible from merely combining multiple judgments, whether with unit or unequal weights.

Unfortunately, the potential for radical improvement in accuracy via group interaction is far from guaranteed, and there is always the risk that the more extreme changes by interacting groups will result in large error. However, a safe bet is that groups will be more confident than individuals (Sniezek, 1992). Arkes offers suggestions about reducing overconfidence in groups. Nevertheless, in light of the popularity of group forecasting, a review of group forecasting processes remains a significant omission from the Handbook.

Yet there is no question that the Handbook has an astonishing wealth of scientifically supported advice to improve judgmental forecasts. Armstrong's chapter on "Judgmental Bootstrapping" is as delightful as it is informative. The scholarship is superb, with a careful and even fun presentation of the literature on judgmental bootstrapping. There are two particularly important messages for the manager who relies on the judgments of experts. First, bootstrapping typically improves on judgmental forecasts.

Second, some people do not care about the first point. Thus if the goal is to improve forecasting accuracy, bootstrapping is a technique worth considering. This chapter instructs on the fundamental and finer points of successful bootstrapping applications in a clear and balanced manner.

One of the many highlights in *Principles of Forecasting* is the excellent review of the empirical literature on overconfidence in judgmental forecasting by Arkes. This chapter explains various causes of overconfidence and how they can be remedied. Drawing on studies of human judgment in various domains, Arkes identifies general principles that can either increase or limit confidence. This work has special significance in the field of judgmental forecasting as confidence in judgment can explain when forecasts are made via human judgment vs. other methods, and how seriously they are taken once made. Perhaps the first step in improving judgmental forecasting is improving the appropriateness of confidence in judgmental forecasts. So, read this chapter first.

One point concerning confidence assessments in Rowe and Wright deserves correction: the data from judgment and choice tasks in our lab do not support the conclusion that relative frequency judgments are more accurate than subjective probabilities. In the study of students by Sniezek, Paese, & Switzer (1990), global estimates of the number of correct choices did imply lower confidence than the subjective probabilities of the individual choice items, but *not* better accuracy. In fact, *the magnitude of the bias was actually greater for the frequency estimates*; they showed a larger *underconfidence bias* ($-.18$) than the subjective probabilities showed *overconfidence bias* ($+.09$). Further, the Sniezek & Buckley (1991) study of managers also showed lower confidence in global confidence estimates than in individual item confidence assessments. But this was true when both forms of confidence assessments were

expressed on rating scales, thereby ruling out any explanation based on differences between frequency and probability judgments. The persistent finding is that confidence in a set of judgments as a whole is lower than the average confidence for the individual judgments comprising the set, regardless of the measure used for the expression of confidence.

Those who are open to the notion that human judgment can be unreliable will be well educated by Stewart's chapter. Unlike many places in the literature, his use of the term "reliability" is perfect. Especially useful for both researchers and practitioners is the distinction between unreliability in information acquisition vs. processing. Readers should understand that the term "accuracy" has a special meaning in this chapter. In the context of forecasting, accuracy typically connotes some form of distance metric (such as mean absolute deviation, percent error, etc.). But in reference to the lens model, the measure of judgment quality—often called "achievement"—is given by the correlation between judgments and the actual criterion values over a set of cases. As a correlation index, achievement is unaffected by the addition or subtraction of an error constant. Even with a very large mean absolute deviation (MAD) error score, achievement can be near perfect. (For details on the relationship between a distance measure of accuracy and achievement, see Sniezek & Reeves, 1986).

Good insights into the proper use judgment in forecasting are provided in chapters by Webby, O'Connor, & Lawrence and Sanders & Ritzman. A well-known problem is deciding how much to rely on historical data vs. one's own knowledge of the phenomenon of interest. One notion that emerges from a reading of these chapters (as well as that by Stewart); is the importance of being able to assess environmental predictability. Knowing the extent to which the phenomenon of interest is predictable allows one to follow guidelines for the use of judgment

in making or adjusting forecasts, and provides useful warnings about the tendency to be more inconsistent in judgments. Yet, as Arkes makes clear, the uncertainty surrounding judgments can be as difficult to judge as the judgments themselves are to make. We need good methods of assessing environmental predictability, but we do not have them yet.

A key issue intimately connected with judgmental adjustments and production of forecasts is *acceptance*. The problem most often encountered is resistance to the forecast, making methods of increasing acceptance necessary. Gregory & Duran build on the suggestions of Armstrong (1985) and Shoemaker (cf. 1997a,b), producing a valuable set of strategies for constructing scenarios that will boost others' beliefs concerning the likelihood of pertinent events and their acceptance of the forecast. Their chapter does a superb job of reviewing research on scenarios, showing how scenarios are powerful tools for influencing others' beliefs about event likelihoods. But do consider two issues concerning the possibility of misusing scenarios to increase forecast acceptance. First, as with any technique of persuasion, scenarios can be used to the advantage of the sender and not the receiver. Because the processes by which scenarios change beliefs are subtle, the receiver is not likely to be aware of the manipulation. Another potential problem is creating overtrust, i.e., an acceptance that cannot be justified due to the uncertainty of the environment and expectations regarding forecast error. Given that forecasts concern the future, they always entail a degree of inherent uncertainty that makes some skepticism more appropriate than total acceptance. Perhaps the proper goal is for a matching between the true beliefs of the forecaster and those of the forecast recipient. That is, the recipient should place no more or less faith in the forecast than does the forecaster.

Overall, the Handbook is clearly successful in reviewing forecasting theory and research and

developing useful principles. Nevertheless, the Handbook is not for everyone—only for those who believe that their forecasts have room for improvement and want them to be more accurate. It is not always obvious to managers and experts that forecasts can be improved. Any attempt at organizational change meets resistance, but this may be especially true with efforts to change forecasting practices. As discussed by various authors (e.g. Harvey; Arkes; and Armstrong's chapter on "Evaluating Methods"), it is not easy to determine the success of forecasting practices due to delayed if not ambiguous or missing feedback.

A far tougher problem in many organizations is the lack of a genuine desire for accurate forecasts. Arkes chapter begins "Much of the research in forecasting concerns accuracy. After all, minimizing the discrepancy between the forecast and the eventual event is everyone's goal." Well, Arkes, your, and my goal, perhaps, but not everyone's. In working with a large organization to improve their group forecasts, we came across a manager who gave a credible interval for his forecast that failed to include his forecast inside the interval. Assuming that he did not grasp the notion of a credible interval, we proceeded to repeat the directions for constructing a range around one's forecast to express one's true beliefs about the expected range of the forecasted variable. His reply was that *we* did not understand; his true belief was that there was no way the actual value would be anywhere near his forecast. In further discussion, it was clear that his forecast was designed for strategic influence on his superiors, not accuracy. He recognized the costs of the error to the organization, but was comfortable that he personally would not be accountable for the anticipated forecast error. While such sentiments may not always be so explicit, they are a troubling source of forecast error. Those who believe that their organization is committed to accurate forecasting might do well to reflect on

the various other objectives their members have in the context of the forecasting task. Even where the desire for accurate prediction is strong, the desire for good outcomes may be even stronger. Rarely do we prefer the case where we accurately predict a severe loss to one where we incorrectly predict disaster and are happily surprised.

Those who do value accurate judgment but are offended by the thought that expert human judgment can be unreliable or biased will be offended by the Handbook, for it takes an unapologetic look at a multitude of factors that can limit the accuracy of judgmental forecasts. Those who understand that the human mind has limitations will find a rich set of strategies for overcoming them in the task of forecasting.

5. Diffusion of Principles (Reviewed by Mike Leonard)

Scott Armstrong and his fellow authors have written an excellent handbook on forecasting. The book, together with its companion Web site and *Forecasting Dictionary*, provides the professional forecasting community with a centralized reference on sound forecasting principles. The handbook's chapters are well organized, with each chapter addressing a particular forecasting principle. Each relatively self-contained chapter provides an overview of a forecasting principle, useful advice to both researchers and practitioners, and an excellent bibliography for further investigations. Each chapter discusses each principle in greater detail and breaks down each forecasting principle into its components in a hierarchical fashion. The forecasting principles are nicely summarized at the end of the handbook with an overall hierarchical breakdown that enables practitioners to audit their forecasting process. The hierarchical breakdown of the forecasting principles is consistently addressed throughout the handbook.

Best of all, the *Forecasting Dictionary* promotes a standard definition to facilitate communication between forecasting professionals, addressing a problematic issue.

As someone who researches and develops forecasting software for SAS Institute Inc., I find that the handbook offers sound advice and has provided direction for future software development. In fact, since the handbook's publication, SAS has already implemented some of the advice contained in the handbook. The handbook states that forecasting principles are often promulgated through forecasting software. This is undeniably true. Researchers uncover new forecasting principles, forecasting software vendors can implement these principles, and skilled practitioners can then use the software to better implement the forecasting principles. But, except for better education, this promulgation cannot stop less-skilled practitioners from buying bad software or misusing good software. From a software vendor's perspective, the promulgation direction of research-software-practitioner is sometimes reversed. Practitioners have needs and the forecasting software vendors must supply their customer's needs (or go out of business). However, though software vendors welcome the efforts of researchers, the researchers may or may not provide answers to the software vendors. Since software vendors are often closer to the practitioners, they are sometimes better tuned to the practitioners needs than the researcher's. For example, missing values and/or limited data, intermittent demand forecasting (with demand influenced by seasonal or causal factors) and multiple-seasonality, high-frequency forecasting (monthly, daily, and/or hourly Internet and electrical load forecasting) are very important to practitioners, but these topics receive insufficient treatment by researchers. The reverse-promulgation of the forecasting principles from software vendors to researchers is mainly a problem between software vendors and researchers; the practitioners just want

sound advice and a way to implement the sound advice with software. The practitioner is the ultimate customer of both software vendors and researchers. The software vendors and researchers need to improve communication and solve the practitioner's problems. In this regard, the handbook should prove to be invaluable to researchers, software vendors and practitioners of forecasting.

References

- Armstrong, J. S. (1985). *Long-range forecasting* (2nd ed.). New York: John Wiley.
- Balkin, S. D., & Ord, J. K. (2000). Automatic neural network modeling for univariate time series. *International Journal of Forecasting*, 16, 509–515.
- Clements, M. P., & Hendry, D. F. (1995). Forecasting in cointegrated systems. *Journal of Applied Econometrics*, 10, 127–146.
- Clements, M. P. & Hendry, D. F. (1998). *Forecasting Economic Time Series*, Cambridge: Cambridge University Press.
- Clements, M. P. & Hendry, D. F. (1999). *Forecasting Non-stationary Economic Time Series*, Zeuthen Lecture Book Series.
- Enders, W. (1995). *Applied Econometric Time Series*, New York: Wiley.
- Shoemaker, P. J. H. (1997a). When and how to use scenario planning: A heuristic approach with illustration. *Journal of Forecasting*, 10, 549–564.
- Shoemaker, P. J. H. (1997). Twenty common pitfalls in scenario planning. In: Fahey, L. & Randall, R. M. (Eds.), *Learning From the Future: Competitive Foresight Scenarios*. New York: Wiley.
- Sniezek, J. A. (1992). Groups under uncertainty: An examination of confidence in group decision making. *Organizational Behavior and Human Decision Processes*, 52, 124–155.
- Sniezek, J. A., & Buckley, T. (1991). Confidence depends on level of aggregation. *Journal of Behavioral Decision Making*, 4, 263–272.
- Sniezek, J. A., Paese, P. W., & Switzer, F. S. (1990). The effect of choosing on confidence in choice. *Organizational Behavior and Human Decision Processes*, 46, 264–282.
- Sniezek, J. A., & Reeves, A. P. (1986). Feature cues in probability learning: Data base information and judg-

ment. *Organizational Behavior and Human Decision Processes*, 37, 297–315.

Snizek, J. A., & Henry, R. A. (1989). Accuracy and confidence in group judgment. *Organizational Behavior and Human Decision Processes*, 43, 1–28.

Paul Goodwin
University of Bath
Bath
UK

J. Keith Ord
Georgetown University
Washington DC
USA

Lars-Erik Öller
Statistics Sweden
Sweden

Janet A. Snizek
University of Illinois at Urbana-Champaign
USA

Mike Leonard
SAS Institute Inc.
USA

PII: S0169-2070(02)00034-1
