

The Ombudsman: Reaping Benefits from Management Research: Lessons from the Forecasting Principles Project

J. Scott Armstrong • Ruth Pagell

The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104-6340

Goizueta Business Library, Emory University, Atlanta, Georgia 30322

armstrong@wharton.upenn.edu • rpagell@emory.edu

It is often claimed that managers do not read serious research papers in journals. If true, this neglect would seem to pose a problem because journals are the dominant source of knowledge in management science. By examining results from the forecasting principles project, which was designed to summarize all useful knowledge in forecasting, we found that journals have provided 89 percent of the useful knowledge. However, journal papers relevant to practice are difficult to find because fewer than three percent of papers on forecasting contain useful findings. That turns out to be about one useful paper per month over the last half century. Once found, papers are difficult to interpret. Managers need low-cost, easily accessible sources that summarize advice (principles) from research; journals do not meet this need. To increase the rate of progress in developing and communicating principles, researchers, journal editors, textbook writers, software developers, Web-site designers, and practitioners should make some changes. We offer some examples: Researchers should directly study forecasting principles. Journal editors should actively solicit papers; invited submissions were about 20 times better than standard submissions at producing useful findings that were often cited, and they do so at lower cost. Textbook writers should focus on principles so that readers can apply knowledge. Web-site and software developers should provide practitioners with low-cost ways to use principles. Practitioners should apply the principles that are currently available.

(Forecasting: applications. Professional: journal policies.)

Consider the plight of practitioners who wish to improve their procedures, say in forecasting. They seem to rely primarily on software, give some attention to books, courses, and consultants, and generally ignore journals. Should they read journals, as good physicians do, to learn about useful scientifically supported knowledge?

What Is Useful Knowledge for Forecasting?

By useful knowledge, we mean knowledge that can help people improve their forecasting. To be useful,

this knowledge must allow a forecaster to determine what actions to take in a given situation. We refer to such condition and action statements as forecasting principles, although they also might be called advice, rules, or guidelines. An example of a principle is, "When uncertainty is high, combine forecasts from two or more forecasting methods." When a principle has empirical support, we refer to it as a grounded principle.

The forecasting principles project, begun in 1997, was designed to assess all useful knowledge in forecasting and to formulate the knowledge as principles. It led to a book, *Principles of Forecasting* (Armstrong

2001), and a Web site, forecastingprinciples.com. Scott Armstrong invited researchers to participate in writing a book that would describe principles and all relevant evidence in their areas of expertise. Forty experts responded to this invitation. In addition, 123 outside reviewers, primarily but not exclusively academics, helped to identify principles and evidence. Appeals for evidence were sent out on e-mail lists and on forecastingprinciples.com. Academics and practitioners visit this site. We encourage continuing open peer review on this site to identify errors and omissions in the principles.

For the purposes of this paper, the operational definition of useful knowledge is any direct evidence related to the principles contained in *Principles of Forecasting*. The principles consist of the 139 in the summary chapter and many others described throughout the book.

In early 2001, we surveyed the authors of the 27 papers who provided the principles in the *Principles of Forecasting*. We asked them to indicate which references in their papers provided evidence on principles. They coded their references as “no evidence,” “indirect evidence,” “direct evidence,” and “cannot remember or not sure.” We left the definition of *evidence* to them. Presumably, the evidence could include logical deduction from known facts, expert opinion, and empirical evidence. As it turned out, they seemed to rely most heavily on empirical evidence. The authors coded the references in their own papers, except for three papers (which Armstrong coded). This process resulted in 545 referenced items that provided direct evidence, some of which were referred to numerous times. (They are listed on the researchers’ page at forecastingprinciples.com.)

Where Does Evidence on Principles Come From?

Papers in journals are the primary source of direct evidence on forecasting, with 89 percent of the 545 items (Table 1). An additional three percent were working papers, and two percent were proceedings papers. Books provided six percent of the items, although this percentage understates their impact because some

Source of Evidence	Percent
Academic journals	89
Working papers	3
Proceedings	2
Books (including chapters)	6

Table 1: Direct evidence on forecasting principles is provided mostly by journals ($n = 545$).

books were cited many times. Remarkably, commercial sources provided little evidence relevant to forecasting principles. Only one citation was to a trade journal, and practitioners were involved as authors on only seven percent of the journal papers.

The seven percent participation rate of practitioners seems low. For example, Hubbard and Lindsay (2002) estimated that about 14 percent of the papers published in major academic journals in marketing had practitioners as authors or coauthors. They also found that the participation of practitioners had dropped steadily from 44 percent prior to 1960.

The evidence came from a wide variety of journals. Eight of the journals had 10 or more items, and over half provided only one item. From 1978 to 2000, only three journals averaged more than one paper per year: the *International Journal of Forecasting* at 4.9, the *Journal of Forecasting* at 1.9, and *Organizational Behavior and Human Performance/Decision Process* with 1.1.

Finding and Interpreting the Evidence on Principles

Given the diversity of the sources, practitioners cannot feasibly track the literature relevant to their needs through journals, even with computer searches. Consider our experience in searching for papers on combining forecasts. An electronic search of the *Social Science Citation Index* (SSCI) using *combining* and *forecasts* from 1988 to 2000 produced 115 papers. But only nine of these (about eight percent) contained evidence on principles. Suggestions from key researchers and references in papers on combining led us to 57 relevant papers, among them the nine found in the computer searches. Thus, by using computer searches we found only 16 percent of the relevant papers. So even if a

practitioner were to have access to the SSCI, computer searches would be neither efficient nor sufficient in identifying useful studies. Furthermore, once papers are identified, it is still an expensive process for practitioners to obtain them.

When one has obtained papers, one must translate the authors' findings into principles. This was the task faced by the authors of the *Principles* book, and it was difficult. Journal papers are typically written in an arcane, often incomprehensible fashion and show little direct concern for principles. And, as is so often the case with studies in management science, many researchers failed to describe the conditions under which their findings hold. For example, in a study of 1,700 empirical papers published by six leading marketing journals from 1984 to 1999, Armstrong et al. (2001) found that only 11 percent of the papers had hypotheses that included conditions.

In view of the uncertainties over interpretation, the principles for nine papers in *Principles of Forecasting* (those written by Armstrong) were sent to the authors whose works were cited. While these authors typically agreed with Armstrong's interpretation, they did not always do so. Their suggestions led to improvements in the principles and to the identification of additional evidence.

Improving the Development of Principles on Forecasting

The 545 papers with direct evidence represent about one paper per month over the latter half of the 20th century, and they constituted a small percentage of the papers published on forecasting. For example, we searched for *forecasting* in titles, key words, or abstracts in the SSCI and located an average of 17 articles per month from 1986 to 2002. However, many papers relevant to forecasting do not have *forecasting* in the title, key words, or abstracts. For example, only about 42 percent of the 545 papers with direct evidence had the word *forecast* (or a variant) in the title. Thus, if this ratio applied to other areas of forecasting, only about one in 40 (that is, 17 divided by 0.42) papers on forecasting would contain useful findings.

This small percentage (2.5) of useful papers is especially disappointing when one considers that, on average, there are only about four papers for each of the

139 principles. Worse, the distribution of studies is unequal. Some principles have been studied extensively, while many have been ignored.

Much can be done to increase knowledge about forecasting principles. We discuss what can be done by researchers, journal editors, textbook writers, Website designers, software developers, and practitioners.

Researchers

To learn more about forecasting principles, researchers should undertake studies designed to provide evidence related to forecasting principles. If researchers depend on serendipity, as has been common in the past, progress over the next half century will be slow.

To aid researchers in selecting topics, we provide a list of research needs on the researchers' page at forecastingprinciples.com. Research needs are high for 23 principles. Some examples are "Obtain information from similar (analogous) series or cases" and "Damp seasonal factors for uncertainty."

Researchers should also take responsibility for communicating their findings. For example, those who obtain useful findings should, in the article abstract, explain specifically what they found and how. That would help others doing metaanalyses to judge whether the paper provided useful information. Unfortunately, academic papers often omit such essential information. For example, Armstrong coded abstracts from 2000 to 2002 (issues 16.2 to 18.2) of the *International Journal of Forecasting* and found that only 17 percent (12 of 69) contained findings and methods. The others provided only vague introductions about what was studied. Similarly, in coding the last two years of the *Journal of Forecasting*, Armstrong showed that only nine percent (six of 68) of the abstracts contained results. On average then, researchers referring to these journals would have easy access to only 13 percent of the findings.

We examined whether papers that provided useful findings receive more attention than those that did not. To do this, we analyzed papers published under the auspices of the International Institute of Forecasters. This included the *Journal of Forecasting (JoF)* from 1981 to 1985, and the *International Journal of Forecasting (IJF)* from 1985 to 1999. During that time, 109 papers

in these journals contained useful findings. We then compared citation rates for these papers to rates for a sample of papers that were published in the same journal issue immediately before or after each “useful paper.” (We used a probability scheme to select the papers before or after; if another paper from our useful list was selected, it was replaced and another selection was made.) As expected, papers with useful findings were cited more frequently. The mean citation rate for the 109 useful papers, 1.4 cites per year, was roughly twice that for the 84 benchmark papers at 0.7.

Journal Editors

To take a simple first step, journal editors should refuse to publish papers until the authors write adequate abstracts explaining the findings and how they were obtained. Readers should be able to identify and understand the key findings without reading the entire paper.

In a review of empirical research, Armstrong (2003) showed that the traditional reviewing process screens out the more useful findings. As a result, the publication of important papers depends to a large extent on active intervention by editors. They should actively seek papers relevant to principles. The reviewer’s role would change from deciding what to publish to improving papers that might be published.

Using the papers published under the auspices of the International Institute of Forecasters, we divided the papers into two groups: those submitted in the traditional manner; and those that received special treatment because they were invited, written by members of the editorial board, or accepted by editors who overruled reviewers recommending rejection. Armstrong coded the papers as having received special treatment or not. In most cases, there was no uncertainty about the coding (for example, editorial board members are listed in each issue). It is likely, however, that he was unable to identify all such papers and these errors would weaken the test.

Using his experience as an editor, Armstrong estimated that about one third of the published papers in the *JoF* and *IJF* received special treatment, while two thirds came through the traditional channels. Given that over 90 percent of special-treatment papers

were published, while about 20 percent of traditional papers were published, traditional papers came from a pool that was about nine times as large as the special-treatment pool. However, about half (53 of 109) of the useful papers published in the *IJF* and *JoF* came from the special-treatment pool. Thus, although special-treatment papers represented about one tenth of the papers submitted, they led to half of the useful papers published.

Among the useful papers, the special-treatment papers were superior to those published through traditional channels. The 53 special-treatment papers were cited 2.0 times per year on average, versus 0.8 for the 56 traditional papers. Based on citation rates, published special-treatment papers were about 2.5 times better than published traditional papers.

A usefulness index can be developed on the two factors: whether the paper has useful findings and its average citation rate. On this index, invited papers were more than 20 times as effective as standard-submission papers (9 times 2.5). In addition, they were obtained at lower cost to the journal because they required much less reviewing and administration.

Thus, by actively intervening in the submission and refereeing process, editors improved the publication rate for papers with important findings. Similar results have been found in other fields. Laband and Piette (1994), in a study of 28 economics journals, found higher citation rates for papers published by authors who had obvious connections to the editors. Similarly, in a study of 15 accounting journals, Smith and Laband (1995) found that citations for articles for which they could identify an author-editor connection were more than triple those for articles for which they saw no connection. In the five years after publication, 69 percent of authors with editorial connections were cited, whereas 63 percent of those with no editorial connections were never cited. Campanario (1996) found a positive relationship between the percentage of papers by authors connected to a journal and the journal’s citation impact factor.

Whom should editors invite to publish papers? The obvious group is authors who have contributed to principles in the past. In addition, editors might provide open invitations for researchers willing to study

specified topics. (For example, an invitation to submit work on forecasting principles has been posted on the researchers' page at forecastingprinciples.com.)

Why does special treatment work? Editors can invite papers from researchers who have been successful in the past, perhaps asking them to study designated topics. The authors are free to approach the problem in any way they like, and they can ask their peers to review their work.

The suggestion that the editors actively intervene is not new. In a survey of editors of 28 education journals, Rodman and Mancini (1977) found that 89 percent of them published "inside track submissions." The successful *Journal of Economic Perspectives* publishes only invited papers. Academics, however, commonly regard special treatment as unfair, as shown in the survey by Sherrell, Hair, and Griffin (1989). But should fairness play a role in editorial decisions if the aim is scientific progress?

Journal editors might also intervene regarding the length of papers. We divided the 109 *IJF* and *JoF* papers with useful findings into five nearly equal-size groups. We compared the two quintiles with the shortest papers (a median of about 10 pages) to the two with the longest papers (with a median of about 20 pages). We then calculated the median annual citation rates for each group. We assumed that longer papers would be more important, and they did average 0.95 citations per article compared to 0.79 for the shorter papers. So doubling the pages increased citation rates by only 20 percent. While other explanations may exist for the small effect size, editors could try to increase overall citation rates for their journals by publishing more short papers (for example, by asking authors to reduce the lengths of their papers). Such a strategy is especially relevant now that journals can post the details needed for replications (for example, data and methods) on a Web site. Editors should be able to identify papers whose length is justified, such as review papers and papers reporting a number of experiments.

Textbook Writers

Many books claim to contain principles when they do not. Armstrong and Schultz (1993) examined nine books claiming to have marketing principles; none contained useful marketing principles.

While books containing grounded principles are rare, some exist, and they have been favorably received. These include March and Simon's (1958) book on principles for the design of organizations and Maier's (1963) book on principles for conducting problem-solving meetings.

Unfortunately, forecasting textbooks do not contain many principles (Cox and Loomis 2001). To remedy this omission, authors can write books specifically to report principles. Indeed, the motivation behind the *Principles of Forecasting* book (Armstrong 2001) was to provide a comprehensive set of principles for forecasting.

Web-Site Designers

Web sites can provide users with low-cost access to principles. This implies the need for a central authoritative location that summarizes principles. This has been done in forecasting through the forecasting-principles.com Web site, supported by the Wharton School, Kluwer Academic Publishers, and the International Institute of Forecasters. This site, which is free, has won numerous awards since it was created in 1997.

Is there demand for information about principles? We examined traffic at the forecastingprinciples.com Web site (Figure 1). Traffic from mid-1998 through 2002 grew rapidly. The site is first for the term *forecasting* on Google and other search engines.

Reward systems in academic research are designed around inputs (published papers). Little concern is given to whether the knowledge is useful or used. The Internet allows for measures of usefulness (continuing peer review along with commentary from practitioners), quality (continuing peer review and corrections), and use (for example, site traffic and links). It is becoming easier to assess the positive impact of research. As a result, reward systems might slowly evolve to recognize researchers who provided useful evidence on the principles in a field.

Web sites can provide easy accessibility to useful knowledge in a field. People who believe that principles are in need of revision or that useful principles have been overlooked can submit revisions or add new useful principles to the site. Continuing peer

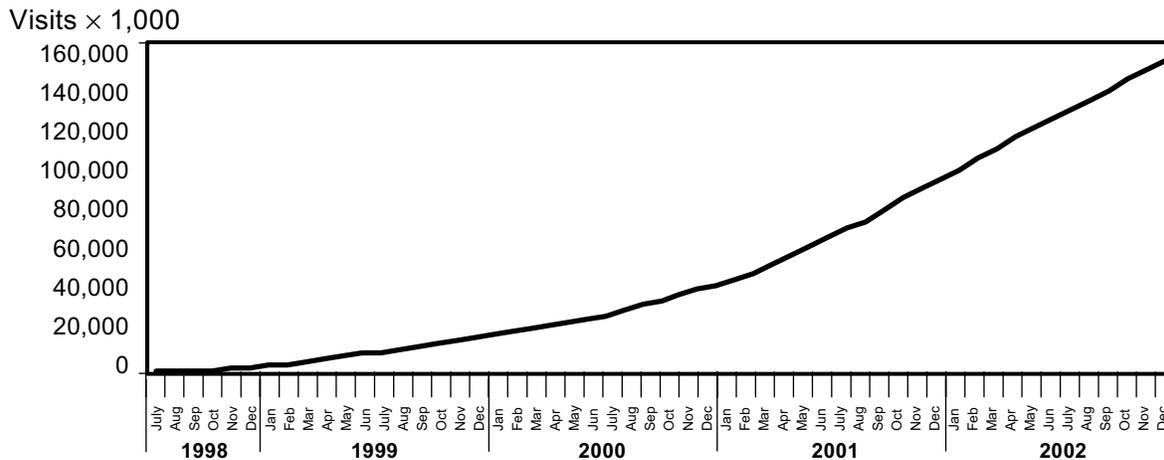


Figure 1: Cumulative visits to the forecastingprinciples.com Web site.

review can be published on the site. Errors and omissions can be used to update the principles. These benefits of Web sites can help to ensure that the site and the related books and papers contain all useful knowledge in the field.

Software Developers

The diffusion of principles will depend in part on how easy they are to implement. Software can play a major role. If the principles are incorporated into the software, they will be used.

Of the 139 major principles from the *Principles of Forecasting* handbook, only about 20 percent are currently used in software packages. No single program uses all 20 percent. The first item on the software page at forecastingprinciples.com is a list of needs for forecasting software. For example, current packages do not help with the principle, "For prediction intervals, incorporate the uncertainty associated with the prediction of the explanatory variables," and they provide little help with principles related to judgmental aspects of forecasting. We believe that software could help to apply most forecasting principles.

SAS and ForecastPro have stated that they will respond to client requests to add principles. This is a good policy. However, software developers are likely to have substantially more expertise than their clients. They should take the lead, and they should post notices when they update their software to implement

new forecasting principles or improve the implementation of existing principles. Tashman and Hoover plan to monitor software developers' notices about additions and revisions. Software companies should notify their customers when they add features that address principles.

In addition to commercial software, we urge people to post freeware. Freeware could be used to guide forecasters through judgmental forecasting procedures and to help them to select appropriate forecasting methods, to name two possibilities. For example, the forecastingprinciples.com site contains freeware that guides users through an audit of their forecasting procedures; it also contains Delphi software to help in the application of judgmental forecasting principles.

Practitioners

To the extent that practitioners rely on current practice, little progress will occur. People will continue to rely on unstructured judgment to make important forecasts. This will lead to the types of forecasts that lead to disasters, such as the dot.com crash.

Practitioners should ask developers to incorporate principles in their software, especially well grounded principles. Perhaps the easiest way for practitioners to identify these needs is to use the forecasting audit (at forecastingprinciples.com) for their forecasting problems. Of course, practitioners can use many of the principles without recourse to software.

Most important is that practitioners take advantage of the many principles developed for forecasting. This can be aided by guides on the forecastingprinciples.com Web site.

Conclusions

Researchers are slowly obtaining useful findings even though research on important problems is not well rewarded. Useful findings are difficult to publish, and once published, they are difficult to locate. In the area of forecasting, we found that only 13 percent of the abstracts in the two leading journals provided an adequate description of the papers. Once located, the papers may be difficult for practitioners to obtain. Those papers that are understandable often lack descriptions of the relevant conditions, thus making it difficult to see how they can be used. In view of these problems, it is understandable that practitioners do not rely on journals, even though that is where useful findings are reported. These problems are expected to be similar in other fields; as was shown by Armstrong (2003), for example, the problems are probably even more serious in the field of marketing.

In forecasting as in other areas of management science, we need to conduct research on principles, do metanalyses to summarize knowledge as principles, and make principles easily available through books, Web sites, and software.

Acknowledgments

In addition to the commentators, Robert Fildes and Richard Franke provided useful suggestions.

References

- Armstrong, J. S., ed. 2001. *Principles of Forecasting: A Handbook for Researchers and Practitioners*. Kluwer Academic Publishers, Norwell, MA.
- . 2003. Discovery and communication of important marketing findings: Evidence and proposals. *J. Bus. Res.* 56(1) 69–84 (commentaries and reply, pp. 85–92). Full text at jscottarmstrong.com.
- , R. L. Schultz. 1993. Principles involving marketing policies: An empirical assessment. *Marketing Lett.* 4(3) 253–265. Full text at jscottarmstrong.com.
- , R. Brodie, A. Parsons. 2001. Hypotheses in marketing science: Literature review and publication audit. *Marketing Lett.* 12(2) 171–187. Full text at jscottarmstrong.com.
- Campanario, J. M. 1996. The competition for journal space among referees, editors, and other authors and its influence on journals' impact factors. *J. American Soc. Inform. Sci.* 47(3) 184–192.
- Cox, J. E., D. G. Loomis. 2001. Diffusion of forecasting principles through books. J. S. Armstrong, ed. *Principles of Forecasting: A Handbook for Researchers and Practitioners*. Kluwer Academic Publishers, Norwell, MA.
- Hubbard, R., R. M. Lindsay. 2002. How the emphasis on 'original' empirical marketing research impedes knowledge development. *Marketing Theory* 2(4) 381–388.
- Laband, D. N., M. J. Piette. 1994. Favoritism versus search for good papers: Empirical evidence regarding the behavior of journal editors. *J. Political Econom.* 102 194–203.
- Maier, Norman R. F. 1963. *Problem Solving Discussions and Conferences*. McGraw-Hill, New York.
- March, J. G., H. S. Simon. 1958. *Organizations*. John Wiley and Sons, New York.
- Rodman, H., J. A. Mancini. 1977. Errors, manuscripts, and equal treatment. *Res. Higher Ed.* 7 369–374.
- Sherrell, D. L., J. F. Hair, Jr., M. Griffin. 1989. Marketing academicians' perceptions of ethical research and publishing behavior. *J. Acad. Marketing Sci.* 17(4) 315–324.
- Smith, K., D. N. Laband. 1995. The role of editors' professional connections in determining which papers get published: Evidence from accounting research journals. *Accounting Perspectives* 1(1) 21–30.
- Tashman, L. J., J. Hoover. 2001. Diffusion of forecasting principles through software. J. S. Armstrong, ed. *Principles of Forecasting: A Handbook for Researchers and Practitioners*. Kluwer Academic Publishers, Norwell, MA.

Comment: Some Biases in the Publication Process That Could Threaten the Validity of the Forecasting Principles Project

Juan Miguel Campanario

Grupo de Investigación sobre el Aprendizaje de las Ciencias, Departamento de Física, Universidad de Alcalá, 28871 Alcalá de Henares, Madrid, Spain, juan.campanario@uah.es

The approach to summarizing knowledge in forecasting implemented by J. Scott Armstrong and Ruth Pagell (2003) is useful and could also be used in some other fields such as mine (science education), in which practitioners and researchers are different. However, in this short comment I will focus on two biases in the publication process that could threaten the validity of the approach. The problem is that most of the

evidence on the principles comes from articles published in academic journals, and the process of academic publication may be biased.

The first bias is the bias towards positive results. In short, journals tend to publish papers in which authors report statistically significant results (Hubbard and Armstrong 1992, Beyer et al. 1995, Campanario 1998a, 1998b, Weller 2001). Parapsychologists were among the first to become aware of the problem. As a consequence, in 1975 the Parapsychological Association Council tried to solve the problem by adopting a policy against the selective reporting of positive outcomes (Bem and Honorton 1994). In the field of biomedicine, researchers have demonstrated the bias towards positive results in papers presented at the International Congress on Peer Review in Biomedical Publications (published in the *Journal of the American Medical Society*, 287, 21, June 5, 2002 and the *Journal of the American Medical Society* 280, 3, July 15, 1998, available at <http://www.jama-peer.org>).

Some surveys clearly show the editors' and referees' preferences towards manuscripts reporting positive results. For example, Kerr et al. (1977) surveyed 429 reviewers associated with 19 leading management and social science journals, asking them about reasons for manuscript acceptance and rejection. Of the referees surveyed, 44 percent stated that they probably or almost surely would reject a paper because of a lack of statistically significant results. Atkinson et al. (1982) studied the ratings psychology referees gave to the methodological rigor of three versions of a fictitious manuscript with identical methodologies but with results reported as statistically significant, marginally significant, or nonsignificant. They observed that reviewers' ratings on methodological rigor were more positive when a study reported significant results.

Authors are usually aware that manuscripts describing negative results are likely to be rejected. Thus, they probably avoid submitting manuscripts reporting such results. This is the well-known file-drawer problem: many studies are never published because of the statistical significance of the results (Scargle 2000).

We can discover this selective reporting by using metaanalysis. For example, Glass (1982) discovered

that in 10 metaanalyses the average experimental effect from studies published in journals was greater than the corresponding effect estimated from theses and dissertations. Dickersin (1990) reviewed some other studies that demonstrated a similar occurrence in the medical sciences.

Some indirect evidence suggests that articles in which authors report significant results have greater visibility than those in which they report nonsignificant results. For example, Easterbrook et al. (1991) reviewed 285 research projects approved by the Central Oxford Research Ethics Committee between 1984 to 1987. They discovered that the studies with statistically significant results had better chances of being published and were also more likely to be published in journals with high citation and impact factors than were studies with nonsignificant results.

Another problem is the bias against replication. It is very hard to obtain funds for replication alone, and journal editors prefer to publish new discoveries and to reject manuscripts that merely replicate previous findings. Many journals' guidelines for authors enforce this policy. Novelty is a necessary condition for manuscripts' acceptance most of the time. Editors often cite scarcity of space as another reason for avoiding "wasting" journal space by publishing replications. This problem is especially acute in the social sciences and the humanities fields, in which rejection rates in journals can surpass 60 percent (Weller, 2001) and paradigms are weaker than in the natural sciences. Curiously, these soft areas are in great need of replications that confirm theoretical approaches and paradigms. To confirm the existence of bias against replication, Hubbard and Vetter (1996) reviewed 4,270 empirical studies published in leading journals of business disciplines (accounting, economics, finance, and management) and found that only about six percent were replications and extensions.

These biases confirm a scenario in which summarizing knowledge from journal articles could yield some biased findings. Thus, I would like to stress the importance of considering (and pondering) the negative findings reported in such documents as abstracts and dissertations to minimize the bias towards posi-

tive results. Another basic precaution is to minimize the importance of findings that have not been replicated. In addition, it would be beneficial to persuade editors to allot more space to publishing replications and confirmations so that we obtain solid findings confirmed by subsequent research.

References

- Armstrong, J. Scott, Ruth Pagell. 2003. Reaping benefits from management research: Lessons from the forecasting principles project. *Interfaces* 33 91–111.
- Atkinson, D. R., M. J. Furlong, B. W. Wampold. 1982. Statistical significance, reviewer evaluations and the scientific process: Is there a (statistically) significant relationship? *J. Counseling Psych.* 29(2) 189–194.
- Bem, D. J., C. Honorton. 1994. Does Psi exist? Replicable evidence for an anomalous process of information transfer. *Psych. Bull.* 115(1) 4–18.
- Beyer, J., R. G. Chanove, W. B. Fox. 1995. The review process and the fates of manuscripts submitted to AMJ. *Acad. Management J.* 38(5) 1219–1260.
- Campanario, J. M. 1998a. Peer review for journals as it stands today—Part 1. *Sci. Comm.* 19(3) 181–211. Available at <http://www.uah.es/otrosweb/jmc>.
- . 1998b. Peer review for journals as it stands today—Part 2. *Sci. Comm.* 19(4) 277–306. Available at <http://www.uah.es/otrosweb/jmc>.
- Dickersin, K. 1990. The existence of publication bias and risk factors for its occurrence. *J. Amer. Medical Association* 263(10) 1385–1389.
- Easterbrook, P. J., J. A. Berlin, R. Gopalan, D. R. Matthews. 1991. Publication bias in clinical research. *The Lancet* 337(8746) 867–872.
- Glass, G. V. 1982. Meta-analysis: An approach to the synthesis of research results. *J. Res. Sci. Teaching* 19(2) 93–112.
- Hubbard, R., J. S. Armstrong. 1992. Are null results becoming an endangered species in marketing? *Marketing Lett.* 3(2) 127–136.
- , D. E. Vetter. 1996. An empirical comparison of published replication research in accounting, economics, finance, management and marketing. *J. Bus. Res.* 35(2) 153–164.
- Kerr, S., J. Tolliver, D. Petree. 1977. Manuscript characteristics which influence acceptance for management and social science journals. *Acad. Management J.* 20(1) 132–141.
- Scargle, Jeffrey D. 2000. Publication bias: The “file-drawer” problem in scientific inference. *J. Sci. Exploration* 14(1) 91–106. Retrieved April 4, 2003, <http://www.scientificexploration.org/jse/articles/pdf/scargle.pdf>.
- Weller, A. C. 2001. *Editorial Peer Review. Its Strengths and Weaknesses*. ASIST Monograph Series, Medford, NJ.

Comment: Research on Forecasting Principles: Formulation, Feedback, and Testing

J. Keith Ord

McDonough School of Business, Georgetown University,
Washington, DC 20057, ordk@georgetown.edu

Armstrong and Pagell (2003) are to be congratulated for raising some important issues and then providing constructive criticisms and suggestions backed by objective information and analysis. In particular, they outline a process of knowledge transmission for management studies and indicate ways in which the present process could be improved. They use the field of forecasting as a source of examples, and since that is my own area of interest, I also will focus on that area.

I will consider the following points:

—How should researchers transmit findings on principles to practitioners?

—How might users contribute to research on principles?

—How do we structure statements of principles to maximize their impact?

The authors begin with the suggestion that forecasters should read the journals, as do good physicians. This statement seems reasonable, but warrants further examination. Physicians cover the spectrum from researchers to general practitioners, as do forecasters, but most physicians are full time, whereas many individuals responsible for forecasting do that as only a (small?) part of their overall responsibilities. Thus, we may need somewhat different channels of communication. O’Neill and Kanyamibwa (1995) found that basic ideas (as measured by citations) tend to flow from the more theoretical to the more applied journals. A reasonable conclusion is that new developments tend to emerge in the more theoretical journals, and the empirical worth of these new methods is then tested in more applied settings. Thus, new treatments may be announced in such periodicals as the *New England Journal of Medicine*, but they are backed by findings reported earlier in discipline-focused research journals. Of course, one journal may on occasion perform both functions.

I agree that we should encourage research on principles, but we should recognize that principles typi-

cally emerge from empirical testing rather than from the Archimedes' bathtub approach, so practitioners will search accordingly. Likewise, papers that contribute to the validation of principles are more likely to be found in journals that emphasize empirical testing, but that is not necessarily a criticism of theoretical journals. Returning to forecasting, good examples of such empirical testing are the forecasting competitions published in the *International Journal of Forecasting*; the most recent being that by Makridakis and Hibon (2000). The same issue of that journal contains several related articles, and a detailed commentary appeared in the same journal the next year (Ord 2001). Earlier research developed the methods explored in the competition, as most practitioners look primarily for empirical validation, not for new theoretical developments. That is, we need information flows that address forecasters' most critical concerns. I am not in any basic disagreement with the authors here.

A second concern is that the information flow Armstrong and Pagell described is largely one way. By all means, we should encourage practitioners to lobby software developers to include well-grounded principles in their products. Also, we should test those principles that are not well grounded. However, these steps miss part of the feedback loop. The practitioner community is often a source of new ideas and problems, such as the need for greater emphasis on forecasts for short time periods (weekly sales, or even minute-by-minute changes in the stock market or in the demand for electricity). The identification of such problem areas leads to new research and ultimately to new principles, but the path may be indirect. Indeed, researchers will often develop several solutions before identifying best practice.

My third point concerns the formulation of principles. Some ideas are straightforward, but others are quite sophisticated and assume the almost mystical quality of delphic utterances. For example, near the end of their paper, Armstrong and Pagell refer to the principle, "For prediction intervals, incorporate the uncertainty associated with the prediction of the explanatory variables." I agree, but implementation is by no means straightforward. We must distinguish pure forecasts from what-if forecasts. In the first case, we seek *unconditional* forecasts as a

basis for planning when typically we cannot control the inputs (for example, company forecasts based on gross domestic product). In the second case, we often have control over (some of) the inputs, such as the advertising budget, and we seek to evaluate the potential impact of different policies (*conditional* forecasts). Thus, some variables have genuine uncertainty attached, for which we should make allowances, but others may not. In addition, we need to account for a number of other factors when constructing prediction intervals (Chatfield 1993). Again, I do not disagree with the authors' aims in encouraging forecasters to develop forecasts in accordance with key principles. However, the statement of such principles may become very complex. Do we perhaps need to start with simpler principles, such as "Use prediction intervals," and turn to more advanced issues later?

Armstrong and Pagell have made extensive use of meta-analyses to support their arguments, and their comments regarding improved abstracts and the use of keywords deserve serious consideration. I suggest that we ask the publication committees of our major professional organizations to consider these issues and develop improved guidelines, rather than leave individual editors to carry on the struggle.

In summary, I support the suggestion that we should encourage further research on forecasting principles, and I agree that our current journals could do more to support such efforts. However, we must recognize that contributions to this objective may take diverse forms and encourage research accordingly.

References

- Armstrong, J. S., R. Pagell. 2003. Reaping benefits from management research: Lessons from the forecasting principles project. *Interfaces* 33 91–111.
- Chatfield, C. 1993. Calculating interval forecasts. *J. Bus. Econom. Statist.* 11(2) 121–144.
- Makridakis, S., M. Hibon. 2000. The M3-Competition: Results, conclusions and implications. *Internat. J. Forecasting* 16(4) 451–476.
- O'Neill, L., F. Kanyamibwa. 1995. A journal citation analysis using MDS. *Proc. Northeast Decision Sci. Inst.* 377–379.
- Ord, J. K., ed. 2001. Commentaries on the M3-Competition. *Internat. J. Forecasting* 17(4) 537–584.

Comment: Too Few Principles or Too Few Grounded Principles?

Mark Uncles

*School of Marketing, University of New South Wales,
Sydney NSW 2052, Australia,
m.uncles@unsw.edu.au*

The forecasting principles project (Armstrong and Pagell 2003) offers many lessons, especially for forecasters and analysts. These lessons are also a wake-up call for those in allied disciplines, such as business statistics, econometrics, marketing, and information systems. Many believe that we have too few principles and that the ones we have are not communicated effectively. This shortage of principles has implications for researchers, journal editors, textbook writers, software developers, Web-site designers, and practitioners. Here I comment on whether there are in fact too few principles, why this might be, and how we can begin to improve processes for the discovery, validation, and communication of grounded principles.

Do We Have Too Few Principles?

This question prompts at least three responses. One is that few principles are to be discovered in business, and those that exist may not be very useful in the context of competitive, dynamic markets. The business world is unpredictable and unknowable, our tools and techniques are necessarily imperfect, and if we muddle through and survive, we are doing well. This response is a throwback to the antiplanning stance of Lindblom (1959). Since then we have developed many useful analytical tools. Nevertheless, we still hear echoes of the antiplanning argument. Hamel (2000), for instance, talks of the need for revolution and for radical, nonlinear thinking in business. Principles that are widely known, he argues, are of little value because they fail to offer a competitive edge. The one worthwhile principle is to question orthodox principles. To succeed in business you must go against the grain, challenge the way things are, instill chaos, and thrive on uncertainty.

A second and contrary view is that we have lots of principles. Many have been found, and many more will be found. Armstrong and Pagell simply haven't

been looking in the right places. Almost all how-to books in business and management are full of guidelines, hints, and tips. Invariably, the authors of these books see themselves as providing principles. For example, Trout and Rivkin's (1996) comment that "over-simplification is the only way around over-communication" is stated as a principle. Their book is packed full of guidance, cases, and home-truths, but devoid of formal evidence. "Get up close and personal, tap basic human needs, mythologize the new, cultivate authenticity, and build communities of interest" are just some of the rules listed by Grant (1999) for building successful brands in the 21st century. Here we have two instances of practitioners setting down principles—two among a list that runs into hundreds across all business disciplines.

Cutting across these viewpoints is a third, that we can discover principles if we are prepared to look, but to be useful they must be empirically grounded. Inspired by the forecasting principles project, a group of marketers have been examining the nature of marketing principles. They stress the importance of discovering and validating grounded principles (see the debate between Rossiter (2001, 2002) and Uncles (2002)). Useful knowledge—be it a concept, theory, framework, law, or principle—is grounded, meaning we can say where it will and will not hold. If we know this, we can make use of the principle. By contrast, a principle based on anecdote may be suggestive but is flawed as a guide because we do not know whether to treat it as an aberration or as a generalization. Its status is unknown. We appear to have too few grounded principles and too many anecdotal principles.

Why Are There So Few Grounded Principles?

To answer this question, we must look at the incentives and rewards associated with finding, validating, and disseminating grounded principles. Armstrong and Pagell hint at the issues but do not point them out strongly enough.

Consider the incentives and rewards for practitioners. Developing formal knowledge does not top any practitioner's list of activities, even though interest is growing in organizational learning. Indeed, many organizations have internal barriers that prevent effective learning (McIntyre and Sutherland

2002). Nor are practitioners expected to test the general validity of their know-how; instead, they seek specific answers to specific problems. Moreover, practitioners have little incentive to publish. Some write for personal fulfillment and derive intrinsic rewards from doing so, but writing for publication is rarely part of a practitioner's job description. In fact, practitioners have reasons not to publish: concerns about confidentiality, protection of intellectual property, and the prospect of commercializing their knowledge. Ironically the so-called knowledge economy intensifies these pressures; knowledge is now something to be commercialized and traded, not something to share. The success of business consultancies, such as Accenture, McKinsey, and BCG, and market research companies, such as Millward Brown, A.C. Nielsen, and Taylor Nelson Sofres, testifies to the value of capitalizing on knowledge. The incentives and rewards for publishing knowledge, therefore, simply do not exist for many practitioners.

By contrast, academics must publish. But in business disciplines at least they are apparently under no obligation to ground a finding thoroughly. Often, they perform only a simple one-off test of their hypothesis. This hypothesis becomes an established "truth" when it is written about in textbooks, usually with no cautionary footnotes to explain the limited testing undertaken. The review criteria used by leading journals and grant-awarding bodies tend to make the problem worse by emphasizing originality and innovation. Better to be new and wrong than old and right! As Hubbard and Lindsay (2002) point out, these practices and underlying incentives impede knowledge development. The incentives and rewards are not aligned with the goal of developing grounded principles.

What Can We Do To Find, Validate, and Communicate Grounded Principles?

Armstrong and Pagell make many suggestions for finding, validating, and communicating grounded principles, although their list is not exhaustive. For instance, if we are to find and validate more principles, practitioners must have more incentives to think about principles and to share their ideas. Prizes? Awards? The Deming awards celebrate quality, and the Clio awards, advertising creativity. What do we

have, or need, to foster principles? Among academics, we must recognize those who put in the effort to undertake exact and differentiated replication studies as part of the process of validating principles. Providing research funding for this kind of activity might be a good place to start.

We can improve the communication of principles by insisting on clear abstracts and executive summaries; by disseminating executive summaries to practitioners (for example, via e-mail); by supporting and encouraging those who distill and interpret journal material for practitioners; by requiring professional training (for example, requiring that certified and chartered practitioners know the principles upon which their business disciplines are based); and by appealing to textbook writers and educators to make explicit which principles are grounded and which are not. Armstrong has demonstrated that we can make progress, notably through the forecasting principles project and the associated Web site (forecastingprinciples.com), and through a related project to provide decision support in advertising (advertisingprinciples.com).

Some journal articles will always be principles sparse. How-to books with no depth will continue to sell at airport bookshops. However, we have grounds for optimism, and Armstrong and Pagell have made constructive suggestions. Raising our consciousness of these issues is a start.

References

- Armstrong, J. Scott, Ruth Pagell. 2003. Reaping benefits from management research: Lessons from the forecasting principles project. *Interfaces* 33 91–111.
- Grant, John. 1999. *The New Marketing Manifesto*. Texere, London, U.K.
- Hamel, Gary. 2000. *Leading the Revolution*. Harvard Business School Press, Boston, MA.
- Hubbard, Raymond, R. Murray Lindsay. 2002. How the emphasis on 'original' empirical marketing research impedes knowledge development. *Marketing Theory* 2(4) 381–402.
- Lindblom, Charles E. 1959. The science of "muddling through." *Public Administration Rev.* 19(2) 79–88.
- McIntyre, Shelby H., Max Sutherland. 2002. A critical analysis into the accumulation of marketing knowledge at the level of the firm. *Marketing Theory* 2(4) 403–418.
- Rossiter, John. 2001. What is marketing knowledge? Stage 1: Forms of marketing knowledge. *Marketing Theory* 1(1) 9–26.

- . 2002. The five forms of transmissible, usable marketing knowledge. *Marketing Theory* 2(4) 369–380.
- Trout, Jack, Steve Rivkin. 1996. *The New Positioning*. McGraw-Hill, New York.
- Uncles, Mark. 2002. From marketing knowledge to marketing principles. *Marketing Theory* 2(4) 345–353.
-

Comment: Breaking Barriers Between Forecasting Research and Practice

Len Tashman
Institute for Forecasting Education,
6681 Spear Street, Charlotte, Vermont 05445,
lentashman@forecastingeducation.com

In the summer of 2002, I joined the board of directors of the International Institute of Forecasters (IIF) and assumed responsibility for improving practitioner services. The organization was concerned about shortcomings in its ability to reach out to forecasting practitioners, whether in business, not-for-profits, or government, enabling them to benefit from the forecasting knowledge disseminated through the usual academic channels—books, scholarly journals, and annual conferences, such as the International Symposium on Forecasting.

Their concern was well rooted. In surveys taken at practitioner conferences, very few participants indicate awareness of the traditional sources of forecasting knowledge. This result reinforces the oft-repeated criticism of academics in applied fields, that they concentrate on communicating with each other, giving scant attention to nonresearchers who merely practice the profession.

The principles of forecasting project, which has culminated in a major new book and comprehensive new Web site, was motivated by the need to more effectively organize and synthesize useful knowledge in forecasting. The book and the Web site have been enormously helpful steps.

In addition, they may be breaking down some barriers between practitioners and researchers. I have seen copies of *Principles of Forecasting* in corporate offices and at practitioner conferences, places where traditional forecasting texts and journals rarely seem to tread.

However, according to Scott Armstrong and Ruth Pagell (2003), we still face many challenges. They

think that practitioners are being let down virtually everywhere: The forecasting textbooks cover methods but are not explicit in providing useful principles and practices. The scholarly journals contain few papers worth reading and often fail to provide abstracts that indicate the value of the research. Forecasting software embeds a minority of the principles of forecasting. These omissions slow the diffusion of innovation—what researchers are discovering, practitioners are not learning.

The root of the problem may lie in the conflicting needs and incentives of academic researchers and business practitioners. Academic research takes a long time—five years from conception to publication is not unusual—and speed merits no rewards (unless the writer is seeking tenure or promotion), nor do clarity or practicality. Business practitioners need rapid solutions to forecasting problems and clear explanations of how new methods and procedures can help them.

Moreover, journal editors consider their responsibility to lie in the intellectual advance of professional knowledge, and they are reluctant to dilute the content of their journals with papers that provide useful knowledge but lack state-of-the-art methodology. Neither of the two scholarly forecasting journals provides incentives for quality research specifically on practical topics. In one of these journals, an attempt to incorporate a practitioner section never got off the ground. The other journal never tried.

One of Armstrong and Pagell's recommendations does seem readily implementable, to improve the content of journal abstracts so that they explicitly include the practical value of the research. I would like to see journals mandate three-part abstracts: (a) what was investigated, (b) the results, and (c) the value to forecasting practice. Absent the last, I'd like the abstract to say, "No practical value apparent," which would be refreshingly honest. Editors would have to allow authors to increase the length of abstracts, possibly offsetting this increase by demanding greater brevity in the bodies of articles.

As well, the forecasting journals should ask authors to write executive summaries, briefly describing the value of the papers being published to both practitioners and researchers. The role model is *Principles*

of *Forecasting* (Armstrong 2001), in which each chapter concludes with the pair of sections: "Implications for researchers" and "Implications for practitioners." These pages should be printed in an initial section of the journal ("Executive summaries of the papers in this issue") and posted on the journal's Web site.

Armstrong and Pagell are insufficiently critical of software developers. I am constantly disappointed at how little attempt developers make to help users understand the statistical results. They have made enormous progress in automating methodologies, which I applaud, but they have paid scant attention to helping the users teach themselves how to assess their models and forecasts. They seldom warn users that the validity of their calculations, especially for prediction intervals, rests on assumptions that have not been tested, may not hold, and possibly cannot be tested. The software seems to be telling the user, "Trust me but don't question me."

Troubling as well is that most software developers refuse to participate in evaluations of forecasting software. The half-dozen or so exceptions dominate the published software reviews. Unfortunately, their packages are invariably the low-priced forecasting engines. The producers of the demand planning packages at the other end of the price scale have not cooperated in attempts at software review and evaluation. I suspect the reason is that their forecasting engines are behind the times, and not that they are ahead of the pack with leads that need to be guarded. It is lamentable that we lack a *Consumer Reports'* capability of independent evaluation.

I agree with Armstrong and Pagell that the main avenue of progress lies in developing Web sites for transmitting useful forecasting knowledge. *Forecastingprinciples.com* is a spectacular beginning and, at present, the main portal for forecasting information. It seems to me that even if journals stick to their intellectual guns with respect to content, they can offer more practical guidance on their Web sites. Textbook publishers can do the same; practitioners today can hardly choose an appropriate text based on the information publishers provide. Provision of accurate guidance might require a level of honesty and objectivity beyond what marketing researchers accept as good promotional strategy.

Armstrong and Pagell do not offer suggestions on what practitioners can do to inform themselves. Directors of forecasting do go to conferences to dispense advice; but they do not often write analytically about their organizations' experiences with forecasting approaches, systems, and software. It would be in the collective interest of forecasting practitioners if they did so. For example, many companies evaluate software vendors and their products and then choose one system for implementation. Their evaluations could be of great help to colleagues at other companies.

Another development offers promise for diffusing forecasting knowledge: the emergence of certification programs for professional forecasters. Surveys over the last decade attest to corporate interest in opportunities to recognize and reward forecasting study and experience.

Certification has been a robust success in the allied fields of production management and finance. A well-conceived program for forecasters can provide incentives for upgrading forecasting skills, offer tangible evidence of accomplishment in the field of forecasting, and build bridges between the creators and practitioners of forecasting knowledge.

In summary, the principles of forecasting project has energized the branch of the forecasting profession that values diffusion of knowledge to forecasting practitioners. Substantial challenges, however, remain rooted in the divergence of motivations between researchers and practitioners, in the failure of forecasting textbooks to teach forecasting principles and best practices, in the inattention of scholarly journals to disseminating the practical implications of published papers, in the fascination of software developers with automating method selection at the expense of enhancing forecaster know-how, and in a very sluggish initiation of certification programs for forecasters.

But Armstrong and Pagell have contributed a road map for progress. To all of us on that mission—Bon voyage.

References

- Armstrong, J. Scott, ed. 2001. *Principles of Forecasting: A Handbook for Researchers and Practitioners*. Kluwer Academic Publishers, Norwell, MA.

——, Ruth Pagell. 2003. Reaping benefits from management research: Lessons from the forecasting principles project. *Interfaces* 33 91–111.

Comment: Do Practitioners Care About Findings from Management Research?

Kesten C. Green

*Victoria Management School, Victoria University of Wellington, New Zealand, and Decision Research Ltd.,
kesten@paradise.net.nz*

In their paper, Armstrong and Pagell (2003) suggest how the supply of research findings to practitioners might be improved, but what about the demand for these findings?

Above all, practitioners wish to prosper in their businesses and careers. In my business, we help managers with decision-making problems and, as part of our service, we often conduct surveys. To our ongoing surprise, we are frequently unable to persuade clients that surveys they wish to commission would benefit from the application of principles derived from research findings.

Here is an example: A client wished to estimate demand for a proposed new service. As it happens, we are world experts in estimating demand; one of my colleagues, Don Esslemont, played a major role in the development of the Juster Scale (Morwitz 2001). The client, however, rejected use of the Juster Scale. He refused to believe that an 11-point scale would work in practice and declined to consider the evidence. Although talk of best practice is fashionable among clients, ignorance of what best practice is, is the norm. Although clients hire practitioners for their expertise, many are not able to judge the expertise of practitioners directly. Some realize this. Others do not and, like our client, make decisions with misplaced confidence.

I assume other practitioners have experiences with clients that are similar to ours. Such experiences discourage practitioners from seeking to improve their practices. Consequently, when I conducted a survey, I did not expect to find evidence of strong demand by practitioners for new findings from academic research.

Evidence of Demand for Useful Findings

I surveyed 41 MBA students who were attending a financial management lecture at the Victoria Management School. Prior to commencing their studies, 10 percent of the respondents worked as administrators, 24 percent as managers, and 42 percent in the professions, as advisors or as analysts. The remainder worked in various other occupations or did not provide clear responses. One had not been working. Forty-eight students were registered for the course. Few, if any, of the students attending the lecture failed to take part. I asked respondents whether, in the year before commencing their studies, they had done anything different in their work as a result of learning about findings from academic research (Appendix). I also asked how many management-related academic papers they had read in that year.

I hate being told that my findings are obvious and that I should not have bothered to conduct my research. With that in mind, I gave six experts (two business school academics and four business school alumni, one of whom was a senior education sector leader) information on the survey and asked them to predict the proportion of respondents who answered “yes” to the first question. I also asked them to predict the median number of papers read by those who answered “yes” to that question and the equivalent number for those who answered “no.”

Like me, and like Armstrong and Pagell, the experts were sceptical about the extent of any diffusion of academic findings from management research. On average, they expected 19 percent of the respondents to have done something different as a consequence of learning about academic research findings; the responses ranged from 10 to 32 across respondents. They expected those who had done something different to have read a median of four papers (1–7) and those who had not, to have read a median of 1.5 papers (0–3). The figures in parentheses are the ranges of the six experts’ responses.

Of the 41 MBA students, 37 percent claimed to have done something different—twice as many as the experts had predicted.

Those survey respondents who had done something different had read a median of 5.5 academic papers in the course of a year. Further, all of those who had

done something different had read at least one paper. As many as 74 percent of those who had read at least one paper and 88 percent of those who had read six or more papers had done something different.

The experts underestimated the number of journals read by those who did something and overestimated the number read by those who did nothing. In other words, they had been overly sceptical about the impact of journals on practice.

Using a questionnaire similar to the one I had used for MBA students, Scott Armstrong surveyed 98 practitioners at the Business Forecast Systems conference in Boston on September 23, 2002. Fifty-five percent of the respondents to this survey claimed to have done something different as a result of learning about findings from academic research, a substantially greater proportion than among the MBA student respondents.

The Boston survey respondents who had done something different had read a median of four academic papers in a year. Most (85 percent) of those who had done something different had read at least one paper. Sixty-five percent of those who had read at least one paper and 79 percent of those who had read six or more papers had done something different. These findings are similar to the MBA student findings, although the link between reading papers and doing something different appears somewhat weaker.

The findings from both surveys support Armstrong and Pagell's (2003) contention that journals are an important source of useful knowledge. The findings are tentative because, in the context of the surveys, useful knowledge was simply information that led the respondents to do something different. The survey responses appear to show that many of the respondents used information arising from academic research, but we have no information on whether what they did turned out to be useful.

The apparent adoption of findings from journals by the surveyed MBA students was higher than any of the experts predicted and the adoption claimed by the Boston practitioners was even higher. Does this mean that the experts, Armstrong and Pagell (2003), and I are overstating the case for poor transmission of academic management research into practice? Perhaps. But, as two of the experts suggested to me (after learning of the findings), the responses might have been subject to social desirability bias. In other words, the respon-

dents were likely to have believed that making use of research findings was a good thing, and that belief might have led them to overreport such behaviour.

The difference between the MBA student survey (37 percent claimed to have done something different) and the Boston practitioner survey (55 percent) might arise from such a bias. In the former survey, I asked respondents who claimed to have done something different to provide an example, while in the latter, Armstrong did not. The requirement for an example might have reined in a tendency to overstate the influence of research findings. Nevertheless, however much practitioners actually care about findings from management research, increasing the adoption of useful findings is likely to be both possible and desirable.

Increasing the Demand for Useful Findings

Armstrong and Pagell (2003) proposed a three-part regime for increasing practitioners' adoption of research findings. First, establish principles. Second, ensure that the principles are accessible to all those with a stake in the discipline—practitioners in particular—through Web sites and books. Third, require that journal editors elicit and select papers, and that researchers choose research projects, with the aim of filling gaps in the principles of their disciplines.

As Armstrong and Pagell describe the situation, forecasters have implemented this regime to some extent. It is early, but the findings from the two surveys could indicate that the regime has increased forecasting practitioners' adoption of research findings. A greater proportion of the Boston-conference forecasters than of the MBA students claimed to have used findings from academic research. Moreover, among those who had done something different, those forecasters who made use of academic findings had typically read fewer academic papers than the MBA students. Assuming social desirability bias does not explain all of this difference, it might be explained in part by the finding that 21 percent of the 98 Boston-conference forecasters had visited forecastingprinciples.com site—all but two of them more than once. None of the MBA students had visited the site.

Of the forecasters who had visited forecastingprinciples.com, 67 percent had also done something

different, compared to 52 percent of forecasters who had not visited the site (Fisher's exact test, one-tailed, $P = 0.17$). In particular, 33 percent of forecasters who had visited the site had also used the method of causal forces in extrapolation, whereas 17 percent of those who had not visited the site had used the method (Fisher's exact test, one-tailed, $P = 0.11$).

The foregoing comparison suggests that Armstrong and Pagell (2003) may be right in their assumption that making access to useful findings easier will lead to greater adoption by practitioners. However, given the typical experience of practitioners with clients, to dramatically improve the efficiency with which we transmit useful findings into practice, we will likely need to increase the demand for useful findings.

Measures to increase demand should lead clients to value practitioners who are familiar with new findings over practitioners who are not. Clients typically are not competent to judge the relative expertise of practitioners, and we cannot reasonably expect them to become experts to do so. Clients rely on such indirect measures as the size and longevity of practitioners' businesses to assess expertise. While it is appropriate for clients to rely on indirect measures, the current measures do not inevitably lead to demand for useful findings from academic research. A well-implemented certification scheme might both satisfy clients' needs for an indirect measure of practitioner expertise and increase demand by practitioners for research findings.

The success of a certification scheme for a management discipline would likely be aided by the prior implementation of the measures Armstrong and Pagell (2003) recommended. In particular, the discipline would need a comprehensive set of well-founded principles.

Despite its apparent benefits, I have reservations about certification. For example, as in medicine and law, might the owners of large consulting firms take control of a discipline's certification to restrict entry for their own benefit? We should consider alternative policies and, in sympathy with Armstrong and Pagell (2003), I recommend proceeding cautiously.

Appendix

Practitioner use of findings from academic research

This short questionnaire asks you about your experience, as a practitioner, of the influence of academic research. I'm interested here in

the influence of findings from academic research in management-related disciplines such as marketing, operations research, organisational behaviour, industrial relations, etc. I'm also interested in the influence of findings from the social sciences in general, for example economics, sociology, social psychology, etc. I'm NOT asking about findings from the natural sciences. As you are currently taking part in a programme of academic study, I'd like to know about your experiences BEFORE you started on this programme.

...

1. During the year before you started this programme, did you do anything different in your work as a result of learning about findings from academic research?

Yes No

4. Roughly how many management-related academic research papers (e.g. journal articles) did you read in the year before you started this programme?

papers

Acknowledgments

I am grateful for the expert opinions of Fred Collopy, Urs Daelenbach, Margot Rothwell, Amardeep Sandhu, Philip Wrigley, and one other who asked to remain anonymous. I am also grateful to John Bradshaw for access to MBA students, and Scott Armstrong, Fred Collopy, Don Esslemont, Mary Haight, and Zane Kearns for suggestions on exposition.

References

- Armstrong, J. S., R. Pagell. 2003. Reaping benefits from management research: Lessons from the forecasting principles project. *Interfaces* 33 91–111.
- Morwitz, V. G. 2001. Methods for forecasting from intentions data. J. S. Armstrong, ed. *Principles of Forecasting: A Handbook for Researchers and Practitioners*. Kluwer Academic Publishers, Norwell, MA, 33–56.

Comment: Technology Transfer in an Academic Setting: Finding the Needle

David N. Laband

Forest Policy Center, School of Forestry and Wildlife Sciences, Auburn University, Auburn, Alabama 36849-5418, labandn@auburn.edu

Armstrong and Pagell (2003) make a compelling case, based on informational search costs, that it simply doesn't pay the typical manager who might benefit from additional knowledge of forecasting principles to search the massive professional literature in the hopes of identifying the very occasional useful

contribution. The manager is looking for the proverbial needle in the haystack. While the benefits may be substantial, the costs likely are prohibitive. It seems to me that those interested in facilitating transmission of forecasting principles from producers (principally academic scholars) to consumers (managers) should identify mechanisms whereby the needle finds the manager.

To accomplish this result, Armstrong and Pagell offer several suggestions: (1) scholars should increase production of "studies that are designed to provide evidence related to forecasting principles," (2) journal editors should find and publish papers that focus on empirical applications of useful principles, and (3) we should develop Web sites that specialize in bringing the same information to interested parties. While these suggestions may seem excellent, the authors beg the question of why we still need to do these things. The implicit indictment Armstrong and Pagell make is that, in their opinion, not enough of activities (1) through (3) is taking place. To the extent that others share their concerns, a logical first question to ask is, why is there underproduction of practitioner-ready research on forecasting principles? My response as an economist is to identify the private incentives (or lack thereof) motivating people to actively provide practitioners with scientific information related to forecasting principles.

Armstrong and Pagell argue, albeit quietly, that scholars have a natural incentive to do so, because forecasting research that touches on principles is cited more frequently than forecasting research that does not touch on principles. Whether this incentive is great enough to induce the level of commitment in scholars that they believe is appropriate is an open question. Presumably, the lure of citations (and the enhanced scientific reputation accorded to the journal that publishes the highly cited articles) also provides an incentive for current or prospective journal editors to focus on publishing papers that touch on forecasting principles.

However, maybe a stronger case can be made. If Armstrong and Pagell are correct, a currently unexploited market opportunity awaits an entrepreneur. They argue that the current supply chain is meeting

the demand for knowledge about forecasting principles inefficiently. It might be worthwhile, therefore, for someone to create a journal that specializes in publishing articles devoted to empirical evidence on forecasting principles (or refocuses an existing journal to specialize in such articles). If the supply of such articles is too thin to justify this specialization, then someone might find it profitable to publish occasional edited volumes collecting already-published papers that are useful to practitioners.

More generally, if the insights and contributions researchers make to forecasting theory (or any other discipline) are of potentially great importance to practitioners, then entrepreneurs who are willing to do the searching and synthesizing (and in some cases the application itself) for the practitioners can make profits. Clearly, some already take advantage of such profit opportunities: a number of academic scholars and other experts provide private consulting services based on their superior knowledge of forecasting principles. Rather than developing the expertise in house, practitioners contract out for it. This is one option that I suspect is used widely in some disciplines and less so in others. The consulting money to be made provides scholars with an incentive beyond citations to focus their research efforts on forecasting principles.

Judging from their analysis, Armstrong and Pagell believe that the existing private incentives are inefficient, in that too little transmission of vital information about forecasting techniques is taking place. However, the advantage of encouraging more transmission through scholarly journals or Web sites instead of through private consulting is unclear. As with other commodities and services, production of useful information is costly. The problem of cost can be overcome if people are willing to compensate the producers of the information. Unless researchers, journal editors, and Web site organizers can profit, broadly speaking, from producing information about forecasting principles, they will not produce it. It certainly is not clear that such information should be regarded as a public good, with its provision subsidized implicitly or explicitly.

So, I am not convinced that the information transfer through the current mix of private consultants,

Web sites, and scattershot journal articles is inefficient. However, if an editor wanted to encourage additional production of research related to forecasting principles and report on this research in a journal specializing in such content, I have several suggestions: (1) find a funding mechanism for keeping the journal financially viable, (2) to motivate scholars to produce good research in this area and to submit it to your journal, award substantial cash prizes (\$25,000 to \$50,000) each year to the authors of the most outstanding or influential articles published in the journal, and (3) pay ad hoc reviewers for providing timely and high-quality reviews.

To cover the costs of suggestions (2) and (3), and perhaps (1), charge a substantial submission fee (say \$500) or page charges for publication or both.

Permit authors to submit their articles to other publication outlets at the same time your journal is considering them.

On a more fundamental level, if managers could benefit greatly from additional information about forecasting principles, then perhaps we should refocus our teaching in undergraduate and graduate business courses accordingly. At least in theory, a market exists for a good forecasting text that emphasizes the identified principles. Normally, a pretty good market also exists for continuing-education short courses. An entrepreneur might find it worth the time and effort needed to explain the principles (and various applications of them) in short courses targeted to those who might really benefit.

Reference

Armstrong, J. S., R. Pagell. 2003. Reaping benefits from management research: Lessons from the forecasting principles project. *Interfaces* 33 91–111.

Reply to Comments: Incentives for Developing and Communicating Principles: A Reply

J. Scott Armstrong
The Wharton School, University of Pennsylvania,
Philadelphia, Pennsylvania 19104-6340
armstrong@wharton.upenn.edu

The commentators raised many interesting ideas in response to Armstrong and Pagell (2003) from

which one general theme emerges: The commentators claim that management science lacks the incentives to encourage efforts to develop and communicate grounded principles. As a result, academics often conduct their research as an intellectual exercise with little concern as to whether their findings might eventually be of any practical use.

The problem extends beyond management science. Smith (1991), an editor of the *British Medical Journal*, concluded from a review that only about 15 percent of medical interventions are supported by solid scientific evidence. He attributes this disconnect to an estimate that only about one percent of articles in medical journals are scientifically sound. Such results indicate problems with incentives in research.

Disincentives

The discovery of new and useful grounded principles carries the message that the currently accepted procedures may not provide the best solutions. Some forecasting procedures, such as Box-Jenkins (Armstrong 2001) and game theory (Green 2002), have been shown to be of little value. Other procedures, such as unit roots testing in econometric forecasting, have little empirical grounding to demonstrate their value (Allen and Fildes 2001). Some researchers may be annoyed by these conclusions, especially given that they are based on empirical evidence.

The history of science shows that many disincentives to scientific advances exist. Those who have made major advances have often been treated poorly. Kuhn (1962) claimed that when innovative findings conflict with important beliefs, resistance is likely to be strong and long lasting. For example, Richard Harrison's quest for an accurate timepiece, which was needed to determine longitude at sea, put him at odds with the scientific community of his day (Sobel 1995). Barber (1961) describes the fierce resistance met by famous scientists.

Present-day examples are easy to find. The Danish statistician Bjørn Lomborg (Lomborg 2001) was denounced by the Danish Committee on Scientific Dishonesty and by other groups largely because they do not agree with his findings that the environment is improving (lomborg.com). This argument had been made initially by the late Julian Simon (1981), who

had also been denounced by many academics. Emotion usually runs high in such cases, even among those who have not read Lomborg's or Simon's publications.

The evidence for resistance to findings is more than anecdotal. Laboratory experiments summarized by Armstrong (1996) show that scientists commonly resist findings that challenge existing beliefs. Typically, they argue that the disconfirming findings are based on poor methodology.

I cannot claim to be a disinterested party in this matter, having devoted considerable energy to discovering and communicating grounded principles. I believe my findings have had a positive impact, and I am currently the most frequently cited professor in the Wharton School's Marketing Department. However, the route I have taken has had perils. Some people regard my findings as heresies (see findings at <http://jscottarmstrong.com>). Reviewers have nearly always rejected what I think to be my most important papers. Thanks to interventions by editors, however, most of my papers have eventually been published. I have described some of my adventures with reviewers (Armstrong 1996). I suspect my commitment to principles has affected my internal promotional reviews at the Wharton School, especially in the late 1970s and the 1980s. Thanks to tenure, I still have a job. Fortunately, the administration and my colleagues have been supportive and done much to help me in recent years. Even so, I am currently paid about half of the usual pay for faculty at my rank at Wharton. Perhaps I will live long enough to gain parity.

Incentives

Not all grounded principles are controversial, of course. But what is the motivation to work on refining accepted principles? In many cases, researchers go to great pains to show that their research is original, and journals look for "originality." Findings that derive from prior work are not held in such high regard. At the extreme, reviewers have a low regard for replications, as noted by Hubbard and Vetter (1996).

Working on principles is time consuming. It is difficult to locate all the relevant studies and to translate them into useful findings. The job often requires reconciling conflicting results. It calls for an under-

standing of the conditions in the principles. Establishing principles is valuable work and it should be encouraged.

The commentators described a number of ways to improve incentives. Their suggestions have spurred me to join with others to take two related actions to encourage useful research on forecasting:

Grants for Directed Research

SAS has announced that it will provide two annual \$5,000 grants to the International Institute of Forecasters to support research directed at developing and testing forecasting principles. The funding process will focus on the research needs found on the Researchers' page at forecastingprinciples.com, and the site will include details about how to apply.

Invited Papers

Recipients of an SAS grant will be invited to publish the resulting paper in the *International Journal of Forecasting*. The paper will also be subject to peer review in an effort to improve it. However, as is the case with the *Journal of Economic Perspectives*, which relies almost solely on invited papers, authors are expected to seek peer review. I will also seek approval from the *International Journal of Forecasting* to extend additional invitations to researchers.

In addition to these efforts to encourage research, the following steps are planned to improve the communication of useful findings:

Informative Abstracts

The *International Journal of Forecasting (IJF)* has begun to ask authors to provide abstracts that describe their findings and procedures. The commentaries by Ord, Uncles, and Tashman offered support for this proposal.

Reviews of Papers with Principles

To make new findings more accessible to others, the "Research on forecasting" section of the *IJF* is now encouraging reviewers to evaluate the papers in light of the forecasting principles summarized on forecastingprinciples.com. In addition, authors of recently published papers are invited to describe how their work contributes to forecasting principles, with the

descriptions to be posted at the forecastingprinciples.com site. The first posting, provided by Paul Goodwin, serves as a model.

Educational Materials

Books that contain information relevant to forecasting principles will be identified on the "Text and trade books" page at forecastingprinciples.com.

Special Interest-Group Pages

Special interest-group pages are being added at forecastingprinciples.com. The purpose is to provide the central source for those interested in forecasting in a particular area. Academics and practitioners are invited to host pages. Currently, Wil Gorr hosts a page on crime forecasting, and Kesten Green has a page on forecasting for conflict situations.

Future Prospects for Principles

As Green mentions, more attention needs to be paid to assessing demand. With respect to the forecasting principles project, it would be useful to learn what principles have been used and which of these have been useful. As a start, people who use the forecasting audit on forecastingprinciples.com can provide their e-mail addresses so that we can send them questionnaires regarding their use of principles.

The International Institute of Forecasters is considering practitioner certification. This could be used to signify that people understand the forecasting principles and have some experience with them.

The increasing popularity of meta-analyses and the Internet have reduced the cost of obtaining evidence on principles. In addition, the Internet has provided people with the opportunity to learn about principles.

It is encouraging to see that similar projects have been successful. The Cochrane Collaboration (cochrane.org) was formed many years ago for "Preparing, maintaining and promoting the accessibility of systematic reviews of the effects of health-care interventions." In 2000, the Campbell Collaboration (campbellcollaboration.org) was organized to find out "what helps, what harms, and based on what evidence" for problems in social, behavioral, and educational areas. In biology, the facultyof1000.com classifies studies as "novel finding," "technical advance,"

"interesting hypothesis," "important confirmation," or "controversial findings;" they then rate these papers so others can see which are recommended, must reads, or exceptional. Descriptions are also provided on why a paper is important.

The technology is now in place for a revolution that would emphasize principles in various fields. While the primary barriers are those relating to disincentives and the lack of incentives, I am optimistic that these will change.

Acknowledgments

I received much help from others, in particular, Monica Adya, Robert Fildes, Paul Goodwin, Kesten Green, Raymond Hubbard, and Mark Uncles. This is not to imply that these reviewers agreed with everything here.

References

- Allen, P. Geoffrey, Robert Fildes. 2001. Econometric forecasting. J. S. Armstrong, ed. *Principles of Forecasting: A Handbook for Researchers and Practitioners*. Kluwer Academic Publishers, Norwell, MA.
- Armstrong, J. Scott. 1996. Management folklore and management science—On portfolio planning, escalation bias, and such. *Interfaces* 26(4) 28–42.
- . 2001. Extrapolation of time series and cross-sectional data. J. S. Armstrong, ed. *Principles of Forecasting: A Handbook for Researchers and Practitioners*. Kluwer Academic Publishers, Norwell, MA.
- , Ruth Pagell. 2003. Reaping benefits from management research: Lessons from the forecasting principles project. *Interfaces* 33 91–111.
- Barber, Bernard. 1961. Resistance by scientists to scientific discovery. *Science* 134 596–602.
- Green, Kesten C. 2002. A comparison of forecasting decisions in conflict situations: Game theory, role-playing, and unaided judgment. *Internat. J. Forecasting* 18(3) 321–344 (commentary, pp. 345–395).
- Hubbard, Raymond, Daniel E. Vetter. 1996. An empirical comparison of published replication research in accounting, economics, finance, management, and marketing. *J. Bus. Res.* 35(2) 153–164.
- Kuhn, Thomas S. 1962. *The Structure of Scientific Revolutions*. University of Chicago Press, Chicago, IL.
- Lomborg, Bjørn. 2001. *The Skeptical Environmentalist: Measuring the Real State of the World*. Cambridge University Press, New York.
- Simon, Julian L. 1981. *The Ultimate Resource*. Princeton University Press, Princeton, NJ.
- Smith, Richard. 1991. Where is the wisdom: The poverty of medical evidence. *British Medical J.* 303(6806), 798–799.
- Sobel, Dava. 1995. *Longitude*. Walker Publishing, New York.