

Management Science: What Does it Have to do with Management or Science?

J Scott Armstrong

This paper is an edited version of the College of Business Studies Silver Jubilee Commemoration lecture, given as part of celebrations recognising 25 years of teaching and research in the Faculty of Business Studies, now the College of Business, at Massey University.

Keywords: marketing, management science

Introduction

First, I want to tell you what I have in mind in talking about management and science. A broad view of management is that it involves procedures to forecast, plan, analyse, decide, motivate, communicate, and implement. On the scientific side, my definition is limited: science is the use of objective and replicable procedures to compare different approaches, techniques or theories.

Management science has delivered many useful things and it has made management more efficient, but it is capable of doing much more. Only a small percentage of the studies in management science are useful, and that proportion is declining. Meanwhile, useful findings that do occur are often unseen, or rejected, or ignored. So, in addition to talking about improving management science, I will also discuss how to improve the communication of important scientific findings.

Why We Need Management Science

We need management science because not everything can be learned from our practical experience. Many things contradict our experience and many are difficult to assess. Here is a list of important and interesting questions in management science that I have studied: Is formal planning useful for firms? Is the Boston Consulting Group matrix a useful technique? Are mission statements useful? Does the pursuit of market share increase long-term profits of a firm? What is the best way to prepare a five-year sales forecast? What is the best way to forecast the outcome of a conflict situation? What is the best way to design a survey? Perhaps you have your own answers.

Management science provides some answers:

- Is formal planning useful for firms? Yes, definitely. There is much research and it shows that firms that use formal planning are more profitable than those that do not (Armstrong, 1982b).
- Do the BCG matrix or other portfolio matrix types help firms make better decisions? No, firms make less profitable decisions as a result of using these matrices (Armstrong & Brodie, 1994).
- Are mission statements useful? Actually I have not studied this, though I would like to. Some people think mission statements are useful, but my hypothesis would be that they

are damaging.

- Does the pursuit of market share improve long-term profits? No, it harms profits. Firms whose objective is to increase market share earn less for their investors (Armstrong & Collopy, 1996).
- What is the best way to prepare a five-year sales forecast? Use a good econometric model (Armstrong, 1985).
- What is the best way to forecast the outcome of a conflict situation? Use role-playing techniques (Armstrong, 1987).
- What is the best way to design a survey? Buy a copy of Dillman (1978) and follow it religiously. You should be able to get a 70% return on mail surveys.

Advancing Management Science

The first task in trying to improve the production and delivery of management science is getting the objectives right. One area of management where there is considerable agreement is the belief that objectives have a big impact on the success of individuals and small groups. The four major factors in defining objectives are that they should be relevant, explicit, measurable, and challenging.

There are also some constraints. Gordon & Marquis (1956) took research reports from different types of institutions; universities, business firms, and not-for-profit organisations. They said two things are needed for useful research: one is to be near the problem so the objectives are clear, and the other is to have the resources to solve the problem. Businesses are near the problem but they do not have the resources for doing the research. Universities have the resources for doing the research but they are usually removed from the problem. Not-for-profit organisations are close to the problem and have the resources. Gordon & Marquis took a set of research studies from these three different types of organisation and disguised who produced them. Then they gave them to a panel of experts and asked them which reports were the most useful. It turned out that the not-for-profit organisations produced the most useful research.

Finally, there is the constraint that you need a motivation to publish. If you were a business firm, why would you want to publish your major findings? Not-for-profit organisations often lack motivation to publish. This explains why over 95% of the research that is published is by academicians. There is a problem because the objectives of academic researchers are vague. One of the consequences of this is that few published studies lead to useful management findings.

Few Studies Lead to Useful Management Findings... and This is Getting Worse

When I started out in the 1960s I would pick up a journal and find many interesting and useful papers. Now I pick up journals and typically do not find a single useful paper. Holob et al (1991) hypothesised that the output of published scientific papers was going up geometrically, but the number of important papers was only going up arithmetically. They did a number of analyses, and found a gradual increase in important, useful papers, but a dramatic increase in the number of publications. The result is that the proportion of important papers keeps dwindling. They called this the Iron Law of Important Papers.

It is difficult to locate useful papers because they are mixed in with all the unimportant

papers. So, how do you find the ones that are useful? Some people might suggest looking in the best journals, since they receive many submissions and the papers in them have been reviewed by two or three people. Unfortunately it doesn't work out that way, because journals seldom use importance as a criterion for publishing papers. They send papers to reviewers who usually check for quality. Most editors use some vote counting scheme - this paper got positive reviews so we'll publish it, this one got negative reviews so we won't (Armstrong & Hubbard, 1991). Consequently, they are not looking at importance; they are looking at quality.

It also turns out reviewers do not even do a good job with quality. Even after a paper has been reviewed by two or three reviewers it typically has serious errors. There have been a number of studies to show this. *The Economist* described one such study: an author had a paper accepted by a journal, then got together with the editor, put in eight errors and sent it out to a few hundred people to review. Few reviewers identified more than two of the eight errors.

The Ideal Scientific Paper

What would be the ideal paper in management science? It would say here is an important problem, and this is why it is important. Next, it would describe how people are currently solving this problem. It would propose some alternative procedures for handling the problem, and would report the results of empirical tests designed to find which of these procedures works best. This should lead to the selection of the most effective procedure. Then it would call for further research. Finally, there might be an appendix with details about how many different procedures were used, or perhaps the author would simply ask people to write for the details.

The Typical Paper

What does a typical journal paper look like? It starts with a research review of a set of loosely related studies. Typically, these studies are not used in the research, they are just there to let reviewers know the author read them. In most of the reviewing I do for major journals, I advise that these pages be eliminated. Then there are assumptions, typically well stated, but seldom realistic assumptions, followed by the logical implications of the assumptions. These arguments are complex and hard to understand. Or there are sophisticated mathematics, which is even more impressive. This is followed by an optimum procedure, given the assumptions. The paper ends with a short call for further research. The major part of this typical paper, at least what is left after I have crossed out most of the literature review, consists of items that constitute appendix material for the ideal paper.

Typical Papers Assume that the Problem is Important

Typical papers assume that the problem is important, but this assumption is not justified. This is normal science, which Thomas Kuhn says does much good, but it is only useful if it is directed toward important problems (Kuhn, 1970). In practice, there is a bias against important problems in the social sciences; it is difficult to get work published if it deals with important problems. This has been proven in a number of experimental studies. Mahoney (1977) picked something almost everybody agreed to be true in behavioural psychology, made up data, and then wrote two versions of a paper. One version supported what everybody believed, and the other contradicted it. He sent the paper to over 80 reviewers. The paper recommended for publication was the one that agreed with the many studies that had already been published. Most reviewers of that version said it was based on good methodology. The version with conflicting results was usually rejected, and the reason it was rejected was that

the methodology was poor, yet the methodology was identical in each version.

In my opinion, many management science studies are unimportant. For example, for the past quarter century people have been doing research using Box-Jenkins procedures. We know from validation studies that this method has no value to forecasters, yet researchers keep working in that area. Strategic planning is another area where people are spending much of their time. They ask questions such as: What defines an excellent firm? What are the characteristics that would match our strategy with the firm and the product? I believe this research is doomed to failure. I have never seen a study that has ever had anything sensible to conclude about what strategy a firm should follow based on that sort of research.

Typical Papers Do Not Observe Scientific Standards

Typical papers also fail to follow the procedure I am talking about when I use the term 'scientific method'. These typical papers tend to use the procedure of advocacy; in other words, they propose an idea and then try to convince the reader without any objective evidence. Here is an example. Escalation bias, in simple terms, is "managers throwing good money after bad." The original studies had an experimental group invest money. Some projects did well, and others did poorly. Then, each group would be given alternatives for investing more money. The researchers concluded that people tended to invest more money in a project that was not doing well; they throw good money after bad. However, we did research on advertising and new product problems and found that this bias did not occur (Armstrong, Coviello & Safranek, 1993). Furthermore, there is no evidence that, if the bias does occur, it is irrational. With the information provided, one could not say people were making poor decisions. The issue is what will happen in the future, and, on that, the subjects had no evidence, so there was no correct decision.

I had trouble publishing my paper on escalation; it was turned down by many journals. I wrote a paper about the publication process and included some of the reviews so that people could see what life is like for researchers who challenge existing beliefs (Armstrong, 1996). In my research I found five other papers reporting a failure to replicate escalation bias. None of these papers, including mine, have been cited by the papers that have been published recently on escalation bias. This occurred even though I sent copies of my paper to all the people doing research in the area. Reviewers do not seem to be interested in citing evidence that contradicts their beliefs.

Despite its importance, replication is not required in order for a paper to be published. Worse, once a finding is published, replications are difficult to publish. Ray Hubbard and I (Hubbard & Armstrong, 1992) did a study on this in marketing, and Hubbard & Vetter (1996) extended this to accounting, economics, finance and management. Only a very small percentage of all papers published are replications or extensions.

Obfuscation is rewarded. In one of my studies I found that the more complex you make the writing or the mathematics, the more highly regarded the paper. I ranked academic journals on how difficult they are to read by their "fog index". Then I sent a survey asking people to rank these journals in terms of prestige. The more prestigious the journal, the harder it was to read. I thought this was interesting; you should make your journals harder to read if you are trying to add prestige. My paper was rejected initially because reviewers claimed that top journals deal with more difficult topics and that is why they are harder to read. So I did another study in which I took conclusions sections from a number of different papers and made some of them more difficult to understand and some of them easier to understand. I sent them out to reviewers; some got the difficult version, some the easy version. All were

told: "Here's a sample of what Professor X wrote. Please rate Professor X's competency." He was rated as much more competent by reviewers who received the most complex sample of writing (Armstrong, 1980b).

Statistical significance is rewarded. Significance testing was not common years ago, but over time it has increased, so now almost all papers with empirical aspects report statistical significance. There are problems with this. One problem is that statistical significance is seldom related to anything important to management. Another problem is that the writers seem to misunderstand what they are doing; for example, they misinterpret type I and type II errors. McCloskey & Ziliak (1996) examined papers in economics journals to learn how economists used statistics. Papers by leading economists contain many errors in interpreting statistical significance.

In psychology, Cohen (1994) in "The Earth is Round (p less than .05)" summarised years of discussion of this problem. After this, a committee of leading researchers from the American Psychological Association was formed. They recommended banning the use of significance tests in published papers. (For a discussion on this issue see Shrout, 1997). *The American Journal of Public Health* had an editor a number of years ago who decided to ban the use of statistical significance. So what do you do? You use confidence intervals. An editor could write to an author and say, "I notice you have reported statistical significance levels. We are happy to publish your paper as it has passed the review process, but you will have to take out all the significance tests. Instead, you can report confidence intervals. Alternatively, you might want to send it to another journal."

The Author's Formula

After reviewing a wide variety of research findings (Armstrong, 1982a), here are my conclusions about getting a paper published. First, do not examine important problems. Second, do not challenge existing beliefs. Third, do not obtain surprising results. If you violate all three of these rules, you have a good chance of having your paper rejected. Also, do not use simple methods, do not provide full disclosure, and do not write clearly. I do not like these conclusions; unfortunately, these are the guidelines for getting your articles into leading journals.

Rejection of New Findings

What do practitioners do with the findings that are published? If the findings agree with what the practitioners already believe, they probably feel happy. If the findings conflict, the managers tend to reject the information. Thus, the papers that have the most value to them are often ignored or rejected.

Steve Dakin and I did a study of New Zealand personnel consultants (Dakin & Armstrong, 1989). We gave them a list of different ways of testing new applicants for a job and asked them which procedures would allow them to see who is going to be the most successful at the job. We had 11 different procedures for testing new employees, based on a meta-analysis involving hundreds of studies collected over half a century. When we ranked these procedures from 1 to 11 according to what the consultants thought were best, there was no correlation between that ranking and the ranking provided by the empirical literature.

The same thing applies to the design of mail surveys. I don't know why people ignore the research in this area because it is easy to apply. I have a rule: "Don't fill out a mail survey unless the people conducting it know the research that has been done." This rule has saved

me much time. I filled in a questionnaire in 1996, but it has been years since I last did that. Usually, I pick up a questionnaire and find that the people who designed it were unaware of research on survey design.

Another study I am currently working on is how to persuade people to do things; much of this in the context of advertising (information on this project is available at the website www-marketing.wharton.upenn.edu/~esap). There is about 70 years of research on this subject, valuable research in my opinion, but it is ignored by people in the advertising business. In fact, they not only ignore it but they get annoyed when you tell them about it. Why is it that practitioners ignore valuable information?

Hope For The Future: Two Positive Trends

What is going to happen to turn all this around? Actually, the problem is going to solve itself to some extent, but we can also do something about it through the development of expert systems and software. We can put new findings into expert systems and new technology into software.

When you have an expert system it solves the timing problem. I teach most of my courses by using projects. Today's topic may not be needed by some students for a few weeks, then maybe when a few weeks are up, they forget what I said. Or they may encounter a problem before it is discussed in class. So I have taken my whole course and put it on an expert system, and students can get advice when they need it.

This solves the timing problem for students when they are studying and when they leave the university. In five years time, when they're working, they will need this information. They are not going to remember much of it, but since they have access to the expert system they can always get help. In fact, they don't even have to take the course; they can just buy the expert system and use it when they need it. Another nice aspect is that we can put the latest research into the expert system, so you have to take action *not* to use the latest research.

Even if we do nothing, the Internet is going to have immense positive impacts on the transmission of scientific knowledge. You can take typical papers and put them directly on the Internet, so people can get any information they want. In the future, journals could just report findings. They would say, "Here are the findings and here is the evidence; for further details, check the web site". The published papers will be more readable, more papers will be published, and it will improve efficiency. Another nice thing is that we will have open peer review and, instead of having two people review a paper for a journal, anybody who reads it can write a review.

Bernard Hibbitts published a paper in a law review (Hibbitts, 1996). I imagine ten people might read that. He also put it on a web site, and in 18 months this site was accessed 6,000 times. People had a lot to say about that paper, and you can see what their comments were. Thus, you can have open peer review of papers and find out whether people use the research and what they think of it.

Management Science Often Ignored Due to Inefficient Delivery

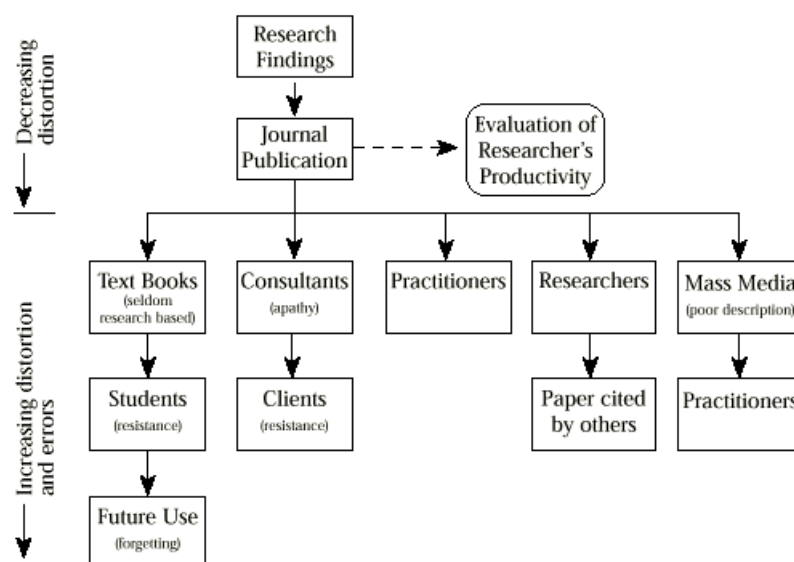
I want to give you a sense of how difficult it is to get management science information out. The problems are illustrated in Figure 1. As you will see, the current diffusion system is inefficient.

For a start, there is an 80% to 90% chance of having a paper rejected by a leading journal in

the social sciences. I have a whole file drawer of rejections, but because I am persistent nearly all my papers are published eventually. It helps that I know many editors and I am a contributing editor. But my guess is that perhaps half of the research findings are never published.

If a paper gets into a journal, it has convinced two or three people that it is good research. This reviewing and rewriting process is helpful and it improves papers, but the key is whether the findings are used. A first step might be to ask how many of these journal findings make it into textbooks. We (Armstrong & Schultz, 1993) did a study of marketing principles books and found the answer was zero. I am sure it is higher in other areas and I know Rossiter's textbook on advertising, for example, includes many findings, as does Nagel's textbook on pricing, so there are textbooks that do contain research findings. Then you have to convince students to accept the findings. But even if they do, they are likely to forget them five years later. So the loss of research from journal publication to use by students after they graduate is immense.

Figure 1. Current diffusion system for management science



One possible route for dissemination of journal findings is through consultants. Some consultants read journals, but most do not. Typically, they have no interest in research. Perhaps they feel their clients won't be interested in research.

One might expect that the most direct route is for managers to read academic journals. I find it puzzling that they seldom read journals. However, given the low likelihood of finding something that is useful and intelligible, managers have some justification for ignoring the literature. Some findings are adopted by practitioners through the trade press, and that is probably a useful way of delivering research findings.

Researchers disseminate their findings to one another. Researchers are pretty good at keeping up with the literature, as long as it is in their field. It turns out, though, that most papers in management science are never cited. How can we change this?

Management Science Possibilities

I suggest we take research findings and publish all of them on the Internet. This is how it would work. All papers submitted to journals would be reviewed. Some papers will receive bad reviews. Instead of rejecting them, the editor would tell the author, "You can withdraw your paper or we will publish it on the Internet with the reviews." Because there is a low cost, journals could publish all the papers submitted, along with all the reviews.

You do not even have to wait for journals; you can just set up your own website and print your paper there. Varian (1997) discusses the issues involved with electronic publishing. Of course, people might say, "How do we know this is good research if it doesn't have the stamp of approval from a journal?" The answer is to have continuing peer review. You take all the reviews for a paper and put them on the Internet, thus providing continuing peer review. You could look back ten years later and see how much attention people paid to a paper and whether it was read by practitioners. In this way we would have a better evaluation of a researcher's impact.

I also mentioned earlier the value of software. Practitioners can purchase software packages with the latest research findings. There is no need to wait for the research to be transmitted through the traditional channels.

This process that I am advocating is summarised in Figure 2.

New technology will lead to a much more efficient system for the diffusion of findings from management science. The Internet reduces the costs of storing and transmitting information. It also makes it easier for the client to find relevant papers. Figure 2 summarises the direction that we seem to be headed. I expect that the evaluation of a researcher's impact will also be improved (see the dotted connections).

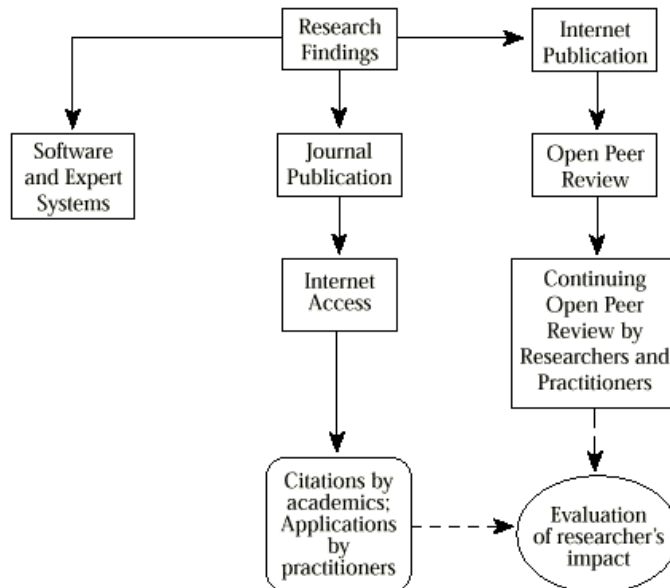
Management Science Impact Sites

What we want to do is to focus on objectives that are relevant, explicit, measurable and challenging. How do we do that? I propose the development of management science impact sites. In these sites, researchers would be asked to describe, in a way that normal people can understand, their most important findings, the evidence for these findings, and why they are important. We can then have these aspects rated by an independent expert panel. They would not know who did the research; they would just see a summary prepared by the researcher. They would rate their importance.

We could summarise these findings, so you could look up a researcher and see how much important research that person has done, or you could see how much a department of a university has done. You could compare universities and replace the current popularity contest among business schools by comparing people on their contributions to management findings. The annual impact ratings could be publicised in the mass media.

You might ask why researchers would participate? After all, it would take much effort to do these summaries. I think they would participate because there would be a cost in not doing so. What if one school comes up with a list of important findings and your school does not? That could be embarrassing. I have had a favourable response from business school deans in the United States.

**Figure 2. Future diffusion system for management science:
Management science possibilities**



Deans

I have some advice for deans. First, a couple of findings. One is that students tend to earn more if they go to a university with a high research impact (Armstrong, 1995). Second, academic prestige does not depend on keeping customers happy; it depends on the school's academic output, the impact of its research (Armstrong & Sperry, 1994). So deans could ask faculty members to report on the impact of their work. This ties in with the management impact site. It would be nice if you could also provide empirical evidence of research impact such as citation rates, beneficial use by practitioners, or web site visits.

Another thing deans can try is the "Three Best" rule, used at a number of major business schools. When you look at someone for possible promotion, ask them to provide their best three papers. This will provide a focus on quality, importance, and impact.

Deans can also try to select productive and creative researchers. A study by Fox (1983) came up with the square root law. She said that the square root of the number of researchers do half the work. So, if you have a thousand researchers, about 30 of them would be doing half the work. But if you consider important research, perhaps you should take the cube root. In other words, perhaps ten researchers out of a thousand do half the important research.

Another interesting finding is that, once a degree is obtained, measured ability or cognitive skills don't seem to be related to whether people do important or useful research. Despite this, we spend much time finding out how smart people are. I think that the best measure of whether somebody is going to be creative and productive is not how smart they are but what

have they done so far. Base your selection of faculty on their record in producing useful research.

It turns out that teaching might act as a constraint on the production of management science. So, how can we bring teaching for management science into line with this effort? A recent Provost at the University of Pennsylvania made what I thought was an astounding (and useful) statement: that the primary purpose of students at the University of Pennsylvania is to be engaged in the research process. Classes should discuss the research process and research findings, and students should be involved in that process.

Teacher evaluations are detrimental to research and to learning. Since teacher evaluations have become much more important, the use of serious papers has disappeared from many courses. Business schools use readings from *Harvard Business Review*, *Wall Street Journal*, *Businessweek*, because students like them. Why not replace teacher evaluations with learner evaluations? Ask students to describe what they learned and what they have been able to use. If you cannot eliminate teacher evaluations, use medians instead of averages and place small weight on teacher ratings when making decisions related to promotions.

Reviewers

What do you do if you are a reviewer? After completing an examination of the empirical evidence on journal reviewing (Armstrong, 1997), I reached some conclusions about the best way to review papers.

Don't worry whether a paper should be published or not. Put all your energy into telling authors how to improve their papers. I also like the idea of signing your name. People tend to act a bit more kindly when they sign their name, and perhaps they feel a bit more responsible for doing a good job if they sign their name. I sign my reviews unless the paper is really poor. Last, and most importantly, do not give the editor a recommendation as to whether to publish the paper or not, because he is likely just to count votes. The editor should decide which of the submitted papers are most important.

Authors

The first thing authors should do is to generate a long list of problems and ask others to rate them for potential importance. By getting this list and narrowing it down in advance, you increase your research efficiency. I look at some problems and think, "Obviously that's not an important problem. Why didn't they know that before they did two year's work on it?"

Do not clutter journals with typical papers. Instead, follow the ideal paper format. Then, when your paper is rejected (and it will be), read the article by Gans & Shepherd (1994). I describe some of my own experiences with publishing new findings in Armstrong (1996). Get used to rejection and keep resubmitting; eventually you will make it, because the reviewing system is unreliable (Marsh & Ball, 1989). Or, you can start your own journal, or publish on the web.

Practitioners

Thanks to the Internet, it is now much easier for practitioners to find academic papers that might be relevant. Examine a paper to see if it follows the ideal format. If it does, maybe you can use it. Do not read typical papers, which can often be identified by their titles. Consider these titles from a recent issue of *Management Science*: "A Stochastic Version of a Stackelberg-Nash-Cournot Equilibrium Model" and "Tabu Search and Ejection Chains -

Application to a Node Weighted Version of the Cardinality-constrained TSP."

Focus on papers having findings that you disagree with. That provides your best chance of learning from journals.

Conclusions

The advancement of science has been undermined by the advancement of scientists. Much effort is currently expended on publishing as a way to get promoted, but little attention is given to assessing whether published papers deal with useful findings about important issues. As a result, the number of important papers is decreasing as a percentage of papers published. Thus, the current system does not appear to be efficient in the discovery and use of important research findings.

However, the Internet will help solve some of these problems. It will make it easier to find relevant papers, and, because the cost of publication is low, researchers will not have to rely on journals to disseminate their findings. Furthermore, continuing peer review can inform readers about the value of papers and become part of the evaluation process for academics.

The implementation of management science will continue to benefit from software programs and from expert systems. Software can incorporate new procedures that have been developed by management scientists, and expert systems can incorporate new findings.

The establishment of management science impact sites would serve to motivate researchers to work on important topics. It would also inform consultants and practitioners about the more valuable research findings.

Thus, I believe that technology will be the catalyst for improving management science and the communication of important scientific findings. If I am right, the potential for increasing management efficiency through the application of science is considerable.

References

- Armstrong JS (1997). Peer review for journals: Evidence on quality control, fairness, and innovation. *Science and Engineering Ethics*, 3, 63-84. (Full text available on the website hops.wharton.upenn.edu/faculty/armstrng.html).
- Armstrong JS (1996). Management folklore and management science: On portfolio planning, escalation bias, and such (with commentaries). *Interfaces*, 26, 4, 25-55.
- Armstrong JS (1995). The Devil's advocate responds to an MBA student's claim that research harms learning. *Journal of Marketing*, 59 (July), 101-106.
- Armstrong JS (1987). Forecasting methods for conflict situations. In G. Wright and P. Ayton (eds). *Judgmental Forecasting*. Chichester-John Wiley, 157-176.
- Armstrong JS (1985). *Long-Range Forecasting: From Crystal Ball to Computer*. New York: Wiley-Interscience, 2nd edition.
- Armstrong JS (1982a). Barriers to scientific contributions: The author's formula. *Behavioral and Brain Sciences*, 5 (June), 197-199.
- Armstrong JS (1982b). The value of formal planning for strategic decisions: Review of empirical research. *Strategic Management Journal*, 3, 197-211.

- Armstrong JS (1980a). Advocacy as a scientific strategy: The Mitroff myth. *Academy of Management Review*, 5, 509-511.
- Armstrong JS (1980b). Unintelligible management research and academic prestige. *Interfaces*, 10 (April), 80-86.
- Armstrong JS & Brodie R (1994). Effects of portfolio planning methods on decision making: Empirical research. *International Journal of Research in Marketing*, 11, 73-84, (commentary follows).
- Armstrong JS; Coviello N & Safranek B (1993). Escalation Bias: Does it extend to marketing decisions. *Journal of the Academy of Marketing Science*, 21, 247-253.
- Armstrong JS & Collopy F (1996). Competitor orientation: Effects of objectives and information on managerial decisions and profitability. *Journal of Marketing Research*, 33, 188-199.
- Armstrong JS & Hubbard R (1991). Does the need for agreement among reviewers inhibit the publication of controversial findings? *Behavioral and Brain Sciences*, 14 (March), 136-137.
- Armstrong JS & Shultz R (1993). Principles involving marketing policies: An empirical assessment. *Marketing Letters*, 4, 253-265.
- Armstrong JS & Sperry T (1994). Business school prestige: Research versus teaching (with commentary). *Interfaces*, 24 (March-April), 13-2.
- Cohen J (1994). The Earth is round ($p < .05$). *American Psychologist*, 49, 997-1003.
- Dakin S & Armstrong JS (1989). Predicting job performance: A comparison of expert opinion and research findings. *International Journal of Forecasting*, 5, 187-194.
- Dillman DA (1978). *Mail and Telephone Surveys*. New York: Wiley-Interscience.
- Fox MF (1983). Publication productivity among scientists. *Social Studies of Science*, 13, 285-305.
- Gans JS & Shepherd GB (1994). How are the mighty fallen: Rejected classic articles by leading economists. *Journal of Economic Perspectives*, 8 (1), 165-179.
- Gordon G & Marquis S (1956). Freedom, visibility of consequences and scientific innovation. *American Journal of Sociology*, 72, 195-202.
- Hibbitts BJ (1996). Last writes? Re-assessing the Law Review in the age of cyberspace. <http://www.law.pitt.edu/hibbitts/lastrev.htm>; version 1.1, June 4, 1996; *New York University Law Review*, 17, 615-688.
- Holob HW; Tappeiner G & Eberharter V (1991). The Iron Law of Important Papers. *Southern Economic Journal*, 58, 317-328.
- Hubbard R & Armstrong JS (1992). Are null results becoming an endangered species in marketing? *Marketing Letters*, 3, 127-136.
- Hubbard R & Vetter DE (1996). An empirical comparison of published replication research in accounting, economics, finance, management and marketing. *Journal of Business*

Research, 35, 153-164.

Kuhn TS (1970). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Mahoney M (1977). Publication prejudices: An experimental study of confirmatory bias in the peer review system. *Cognitive Therapy and Research*, 1, 161-175.

Marsh HW & Ball S (1989). The peer review process used to evaluate manuscripts submitted to academic journals: Interjudgmental reliability. *Journal of Experimental Education*, 57, 2, 151-169.

McCloskey DN & Ziliak ST (1996). The Standard Error of regressions. *Journal of Economic Literature*, 34 (March), 97-114.

Shrout PE (1997). Should significance tests be banned? *Psychological Science*, 8, 1, 1-20.

Varian HR (1997). The AEA's electronic publishing plans: A progress report. *Journal of Economic Perspectives*, 11, 95-104.

J. Scott Armstrong is Professor of Marketing, at the Wharton School, University of Pennsylvania, Philadelphia.