Proceedings of the 2008 Winter Simulation Conference
S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler eds.

A PRACTITIONER, A VENDER, AND A RESEARCHER WALK INTO A BAR: TRYING TO EXPLAIN WHAT RESEARCHERS DO

Bruce W. Schmeiser

School of Industrial Engineering
Purdue University
West Lafayette, IN 47907–2021

ABSTRACT

Practitioners, venders, and researchers form three influential groups within the WSC community. The groups depend upon each other, yet not-enough interaction exists. Researchers, in particular, have difficulty explaining their culture and contributions. This paper is an attempt at such an explanation.

1 INTRODUCTION

Since at least 1976, when I first attended the Winter Simulation Conference (WSC), the WSC program has been organized into tracks. Over the years, the number of tracks has increased, but year after year each track draws its own crowd. Practitioners participate in the various application tracks, venders participate in the software tutorials (and, since 1984, the exhibit area), and researchers participate in the Modeling and Analysis Methodology tracks. Of course a bit of mixing occurs among the tracks, but most discussions among practitioners, venders, and researchers are outside the program at user-group meetings, at society business meetings, and in the hall smoozing. Even there, the discussions seldom allow practitioners, venders, and practitioners to understand each other.

If asked, most conference attendees would quickly identify themselves as being a practitioner, a vender, or a researcher. (Those who easily identify with more than one group are rare; the language venders of the early 1980s come to mind.) Most attendees could (and would) quickly describe the other two groups. (The animal-vegetable-mineral analogy might be fun, asking each attendee which matchup to practitioner-vender-researcher would be most appropriate.) My sense is that substantial misunderstanding exists.

Venders probably have the least misunderstanding. Having a product to sell, they need to understand their buyers, who are typically practitioners. Needing to develop a product to sell, they need to understand the state of the art, which arises ("by definition", as I and other researchers like to say) from the researchers. Ideally the state of the art is embedded in the vender's product; the practitioner should not need to be bothered to learn the state of the art.

Practitioners have more misunderstanding. They have a job to do. I hope that the job is to answer a set of questions so that a decision maker (who might be the practitioner) can make better decisions than if the practitioner weren't doing his or her job. To do the job, a simulation practitioner needs deep understanding of the job context, excellent social skills, a good understanding of at least one vender's product, and (ideally) no understanding of the state of the art.

Researchers have the most misunderstanding. I am a researcher. I would like to understand other groups better. For perspective, I occasionally attend other tracks, but my own track (Analysis Methodology) invariably calls. In the selfish hope that others will make the effort to describe their groups to me, I describe here what we researchers do. As must be, other researchers have different views, but to a large extent we are alike and what I say here reflects what is going on in the research community.

For perspective, because I think that my background is similar to many other researchers, in Section I describe how I became a researcher. Since most (almost all?) WSC researchers are professors, Section 3 is about what professors do for a living. Section 4 is about research, the process of stating and defending a problem and its solution(s). Section 5 is about the WSC research community, at least as seen in the Analysis Methodology track. Finally, in Section 6 I try to extract lessons from two of my own research topics.

2 MY BACKGROUND

I was a practitioner once, with Electronic Data Systems (EDS)in the early 1970s. Being a practitioner was difficult, and not because of punch cards and jockeying for computer time in the middle of the night. The difficulty was people and communication skills: interviewing non-technical clients to create models, extracting cost data from accountants who didn't understand why I wanted to exclude overhead costs (and I who didn't understand that the term "marginal cost" would have been useful), and writing weekly status reports. Also, I noticed that the obvious higher-paying career path meant managing other employees, and management appeared to be really difficult.

After about a year of the real world, I began applying to Ph.D. programs. I needed a career path that allowed me to follow my nerdy technical strengths. As with many eventual WSC researchers, my strengths were short on people skills and long on mathematics, statistics, and computer science. Indeed, my undergraduate mathematical-sciences degree was a set of courses combined by the mathematics, statistics, and computer science departments at The University of Iowa.

I was never quite a mathematician, though, because I was often the student in class who disappointed the professor by asking whether "this stuff" was ever used. Other than for a part-time programming stint with Westinghouse Learning Corporation, where I implemented Fortran programs to compute descriptive statistics for American College Testing's Freshman Profile, I saw no evidence that the useful "stuff" that I was learning was indeed useful. Fortunately, before graduation I was required to meet with an adviser to ensure that I had taken sufficient and appropriate courses for my Bachelor of Arts degree. He asked a life-changing question: What did I enjoy more, pure or applied mathematics? I clearly didn't know, because I asked in reply which had I taken. He suggested that I go across the street to take an introduction to operations research; "across the street" referred to the engineering building. The course was taught (well) by Professor John Liitschwager in industrial engineering. More important than any course concept, however, was that engineering provided the technical problem-solving community that I had been looking for. But graduation and the military draft loomed.

An April 1969 motorcycle accident gave me a one-year draft deferment, so I enrolled in the IE Ph.D. graduate program at Iowa. Professor John Ramberg introduced me to systems simulation, first in a fall (Saturday morning, ouch) course and then in a spring readings course. Based on one of his suggestions during the readings class, I volunteered to investigate how to generalize Tukey's Lambda Distribution, an effort that became my MS thesis and introduced me to the joy of falling asleep and waking up thinking about a single research question for an entire summer. And then I won the first draft lottery (with number 183), removing a primary reason to be a student. I dropped out with an MS degree and suddenly was an operations-research practitioner at EDS. But more than freedom from the draft caused me to drop out: The salary of $11000 per year seemed huge, IE at Iowa required at least one year of work experience to obtain a Ph.D., and I had once again begun to wonder whether the course material was "useful stuff".

My practitioner experience at EDS, brief as it was, made me a believer that the mathematics, statistics, and computer science "stuff" was useful. And that bit of practitioner experience has guided my research career, the point of which has been to try to improve ability of practitioners to do a good job, ideally automatically, without the practitioner being aware, by having the research results embedded within the off-the-shelf software.

3 ON BEING A PROFESSOR AT A RESEARCH UNIVERSITY

As with most WSC researchers, I am an ivory-tower researcher, having been a professor since I graduated from Georgia Tech's School of Industrial and Systems Engineering in 1975. Professors engage in teaching, research, and service, but for most of us the reward system is hugely skewed toward research, with quality of teaching and service ignored until students march in protest or a law suit is threatened. Many of us consult a bit, giving some advice to a practitioner or to a vender, but we seldom face the real-world issues of implementing and selling models or of creating and selling products. Some professors do face an almost-real-world issue—attracting funding—a revenue source that is an ever-increasing share of research-university budgets.

Tenure—the awarding of a life-time employment contract—is at the heart of the life of a professor. Conceptually the argument for tenure is simple: If the position is called *professor*, then the person in the position should *profess*. Professing well is aided when there is no fear of losing one's job. The content of engineering and science teaching usually is not controversial, but management professors more-often deal in opinion. More relevant to most of us than freedom of class-room speech is the freedom to choose research topics, popular or not. That freedom is central to an interesting research career. And maintaining the professor's interest is crucial in a tenured world where lack of productivity is not a criterion for termination.

Only a tiny part of a professor's activities are required. Showing up for class, almost all of the time, is required. And some committee activities. And some formal advising. I'll estimate that required activities comprise twenty percent of what I do; the rest is voluntary. Much of the rest is difficult to measure, both the quantity and the quality. Advising undergraduate and graduate students about course selection, career direction, and thesis adviser. Advice to former students. Writing letters of recommendation. Writing grant and contract proposals. Attending seminars, sometimes to be supportive rather than from a fundamental interest. Informally talking with colleagues about their research. Visiting other institutions to give a seminar, provide curriculum advice, or generally promote interaction. Attending and speaking at conferences. Service to professional societies. All of these activities are desirable; none are required.

And I haven't mentioned research, which is what draws many of us to universities. After tenure is awarded, research is another non-required activity. A fundamental question that I ask myself when evaluating a tenure case is whether I think that the candidate loves doing research. Research effort arising in response to a reward system, rather than for the inherent enjoyment, can quickly fade.

The good news, at least to those of us who are night owls, is the freedom to work when we please. For me, that means that afternoons can focus on advising, discussions, seminars, and teaching, with the nights free for research. And that others can find their own schedule. Alan Pritsker and I occasionally would overlap our day in Grissom Hall at about 3:30am, he arriving to write new material until 8a.m. For research productivity, we all find a few uninterrupted hours every day, and the university environment provides few constraints.

Having few constraints, however, provides the argument against the tenure system. Any tenured professor can choose to ignore the non-required activities, and not much can be done about it. The positive solution is allow the professor to maintain interest. And that is aided greatly by the free choice of research topics.

A researcher's freedom to work on any topic can be scary. Having a boss to assign a task is easier, much like an entry-level position is easier than a managerial position. Most students, before attempting a thesis, have spent their lives taking courses, with the instructor providing assignments and the students providing solutions. Many excellent students, as measured by grade-point average, struggle when the assignment is to decide what to do. For several years, in my IE680 research-level simulation course, I have asked the students to prepare for class, but to submit only whatever they choose during the semester. Very few students like the freedom to decide what to do. Even after I make some suggestions (work specific examples, or write computer code, or write a paper summary, or make a class presentation, or engage me in discussion), the uncertainty about what *I* want dominates. Seldom does a student naturally enjoy the freedom to spend time on what he or she wants when a course grade is to be assigned.

Similarly, most students prefer to have a thesis topic and problem posed for them. At the thesis defense, a classic question is why the student chose the particular topic and problem; the classic bad answer is that the adviser suggested them.

4 THOUGHTS ABOUT PURSUING RESEARCH

The following subsections contain random thoughts about research.

4.1 Stating the Problem

Stating a research problem is difficult. Very often when I ask someone about his or her current research, the reply is about the solution. For example, "I am using neural networks to choose input models". Other times, the problem is only vaguely stated. For example, "I am working on the initial-transient problem." Despite being difficult, I find the exercise of stating the research problem, with no mention of solution approach, useful.

In particular, I like to state the problem in two ways. First, from the point of view of a practitioner. I want to know what the practitioner is assumed to know, what is to be found, and how the quality of the solution will be measured. For example, given a vector of n steady-state data from an unknown process, estimate the marginal standard deviation to minimize the mean squared error. Problem statements usually need to be clarified. For example, what computational effort is allowed? In many simulation contexts, computation needs to be O(n) because otherwise the time spent computing the point estimator could be used to obtain more simulation data. But many reasonable standard-error estimators are quadratic forms, so knowing whether $O(n^2)$ computing is acceptable should be part of the problem statement.

The second way to state the problem is from the point of view of the researcher. In my world, the researcher's problem is to create a method, usually implemented in computer code, to solve the practitioner's problem. Ideally, a design criterion is that the practitioner need not specify any "magic" parameters. For example, a practitioner should not be asked for an initial step size, a number of batches, or an amount of initial data to be deleted. The practitioner should need to provide only the problem context and the quality of the solution desired.

When an adequate problem statement is available, researchers should be able to determine which of two solutions is preferable. A good exercise is to have someone else read the problem statement to check whether the problem statement can be broken. For example, a few years ago I was on the committee of a Ph.D. student who proposed a scheduling algorithm designed to minimize total tardy time. First, of course, I corrected the statement to "expected total tardy time". Then I asked the meaning of "total tardy time", which turned out to be the sum of tardy times of all customers who departed the facility. I proposed the rule that no tardy customer be allowed to leave the facility, which always would return zero total tardy time. The student protested that of course he wasn't interested in such a rule. But if such a silly rule satisfies the problem statement, then certainly there are other rules that might also be unacceptable even while seeming to be good. As it turned out, the student's simulation experiments didn't clear the queue, causing a biased comparison: poorer scheduling rules usually ended with unseen tardy customers still waiting for service while better rules usually cleared all customers from the system.

Defending the problem is as much part of research as defending the solution. Typically thesis defenses pursue these three questions and answers: (1) Why is the problem of interest? Because it is important. (2) If it is important, why hasn't it been solved? Because it is difficult. (3) If it is difficult, why do you think that you can solve it? And then we come to the *hook*, the insight into the problem that might allow the researcher to solve an important and difficult problem.

4.2 Research is Opportunistic

The pattern of choosing a topic, focusing on a problem, and then finding a hook is only sometimes sequential. Often an adviser suggests a thesis topic and the student struggles to define a problem that has a hook. In contrast, experienced researchers live in a world of several topics—often a set that increases with experience—and wander around reading and listening to many ideas, finding synergies among disciplines that provide hooks to problems.

For many of us, off-topic seminars are useful because the speakers' comments can stimulate useful thoughts about our set of topics, either about new problems or about new solutions. My guess is that the process is common in other fields. Opportunistically, just now I remember a short internet article (Ker Than 2005) that contains the quote from Bradley Folley, a graduate student in clinical psychology at Vanderbilt University. "Creativity at its base is...taking things that you might see and pass by everyday and using them in a novel way to solve a new problem." Years ago, Lee Schruben told me that, at the beginning of his professorial career, he tried to sit-in a course every semester. Beyond the continued learning, my impression is that attending courses contributed to his impressive creativity. How else does someone think of solving the warm-up problem by reversing the output time-series data and then using quality-control concepts to ask when the reversed time series is out of control? How else to think of using a Turing test (originally proposed in 1950 to determine whether machines can think) for model validation? How else to relate graph theory to next-event modeling?

Being opportunistic is at the heart of a distinction between research and development. Like research, development can be difficult to do well. But development can be done on a schedule, in contrast to research, which requires the hook, which requires something unexpected. For true research, sitting down to "do" research is difficult. Progress is not linear. Sometimes another week's passing means only that the researcher now understands why a current idea is not worthwhile to pursue. The boundary between research and development is not pure, especially since after finding the hook much of research is development.

4.3 A World View of Simulation Experiments

From the research point of view, simulation is one of three methods to analyze a given probability model. The other two are to seek closed-form solutions, such as taught in classical probability courses, and numerical methods, such as taught in numerical-analysis courses. Probability models have two parts: the *input model* and the *logical model*. The input model describes the random behavior of individual components; for example, that customers arrive according to a Poisson process with rate six per hour. The logical model describes how the random components interact; for example, that arriving customers begin service immediately if a server is available and otherwise enter the queue. Much like the distinction between data and code in computer science, the boundary between the input model and the logic model is sometimes vague; for example, when customers arrive deterministically every ten minutes. Given an input model and a logic model, analyzing a probability model means to determine the numerical values of one or more *performance measures*; for example, the mean and 0.9 quantile of customer time in the system.

The world view of input model, logical model, and performance measures, which allows every application to fit within the same structure, allows researchers to think abstractly so that results are general. Venders are always seeking common structure, because they are like researchers in seeking generality. Practitioners who use a common structure for all applications know what to look for, much as in optimization modeling one knows to look for decision variables, an objective function, and constraints.

Researchers don't move much beyond the world view comprising input model, logic model, and performance measures. Simulation used for training, for example, seldom is discussed in the Analysis Methodology track. Stochastic root finding and stochastic optimization augment the world view by defining some parameters of the input model and logic model to be decision variables. Researchers often think about how to improve the quality of input models, but usually consider the logic model to be given.

Arising from Barry Nelson's Ph.D. thesis (Nelson 1983), the world view of simulation experiments can be thought of as five sequential boxes, which I usually draw horizontally. On the left is the random-number generator (including initial seed), code that makes a deterministic computer seem to create random numbers, the contents of the second box. The random numbers are assumed to be uniformly and independently distributed over the unit interval. Practitioners hope that the vender has taken care in the choice of generator(s), venders often do take care, and researchers have the luxury of simply stating an assumption. The random numbers from the second box are used to generate random variates, which lie in the third box. These random variates are observations from the input model. The random variates are fed into the logic model, to produce output data (e.g., the customer times in the system), which lie in the fourth box. The output data are combined using a point estimator of the performance measures; the fifth box contains the point estimates, which are then reported to the practitioner. This world view is implicit in researcher discussions, but most textbooks don't discuss it explicitly; Leemis and Park (2006) is an exception.

5 ABOUT THE WSC RESEARCH COMMUNITY

The following subsections are my attempt to explain the culture of the WSC Analysis Methodology track.

5.1 WSC Researchers

I am pleased that the WSC community of researchers has both high standards and fine collegiality. Recent examples abound. Three days ago, Shane Henderson sent me (and co-authors) an email note pointing out that very similar work was published about six years ago. How much better to have the reference now than to have published our work only to find out later about the earlier work. Similarly, I remember David Kelton pulling an accepted paper before it was published because a conceptual error needed attention. This summer Mike Taaffe has been rewriting "Concise Notes", which I compiled for an introductory probability course; how pleasant when arguing about authorship that everyone is being generous. In contrast to colleagues in other sub-disciplines, I have never had an argument about author order.

5.2 Background Disciplines

Simulation research, at least in the Analysis Methodology track, lies at the intersection of multiple academic disciplines. I've already mentioned that the departments of mathematics, statistics, and computer science defined my undergraduate degree. But every faculty member in each of those departments would say that his or her department includes several disciplines.

For the Analysis Methodology track, the first year of an engineering probability and statistics course allows a listener to understand a bit of what is happening. Frequentist statistics dominates the discussions, with Steve Chick reminding us of the advantages of Bayesian statistics. Courses in queueing, design of experiments, response-surface methods, time-series analysis and survey sampling are relevant. Real analysis arises repeatedly, typically to prove various-forms of algorithm convergence (that is, if the algorithm runs forever it will return the correct answer). Random-number generation uses number theory and depends upon computer architecture, topics that seldom arise otherwise. Numerical analysis is useful, often because improving a simulation experiment often involves combining elements of numerical methods, Monte Carlo, and closed-form analysis.

With the influence of so many disciplines, defining the state of the art is sometimes difficult. Is using another discipline's established idea to solve a problem from our simulation world view a change in the state of the art? The answer depends upon a person's background, with the same presentation often seeming elementary to one listener and insightful to another listener.

5.3 Real Numbers and Computer Numbers

At our best, we blend the various disciplines, but sometimes we ignore the possibilities. An example that I have discussed briefly, at different times, with Jim Wilson, Wheyming Song, Fen Chen, and Raghu Pasupathy is why we use real analysis to prove algorithm convergence. By definition, real analysis is about real numbers. Simulation algorithms, in contrast, use computer arithmetic, which can represent only some rational numbers, a small subset of the real numbers, with the subset depending upon the particular computer. Therefore, we sometimes use real analysis to prove that an algorithm converges to some point θ, even when the computer can't represent θ.

The difference between real and computer arithmetic also arises when we assume that random numbers are continuously distributed on the unit interval. On 32-bit computers using a random-number generator with period $2^{31}-1$, the largest 2000 (approximately) random numbers are represented as one, an effect that doesn't happen to the 2000 smallest numbers. Programmers quickly notice the effect when, for example, they generate an exponential random variate using $x=-\ln(1-u)$. The real-world jolt of having x be undefined keeps us honest in random-variate generation, which is quite different from algorithm convergence where we never notice that θ is undefined (because no algorithm runs forever).

5.4 Engineers Versus Mathematicians

The research community, as a whole and often individually, bounces between the engineering culture and the mathematics culture in two ways. First is the battle between being useful and being elegant. *Useful* means that a practitioner or vender uses the problem solution. *Elegant* means that the problem solution is insightful and has a theoretical base; typically stating conditions under which infinite computing will provide the stated solution. An algorithm that explicitly searches a finite space of solutions is not elegant, and only sometimes useful.

Many of us hope for both useful and elegant. In the world of algorithm design, we often plot expected squared error versus computing effort. If the curve goes to zero, then it converges. If Algorithm A has a curve that is lower than other algorithms' curves, then Algorithm A is useful in that it is a natural choice for a practitioner or vender. Creating the algorithm with the lowest curve is good.

The second way that the WSC researchers bounce between engineering and mathematics is the underlying research goal. A useful cliche is that science seeks to find what exists and engineering seeks to create what didn't exist. In that sense, creative design of algorithms is engineering and the theory upon which algorithms (one hopes) are designed is science. Many of us try to do both, but most know where our interests and strengths lie. And for most of us, despite our symbol-filled papers, we are engineers, designing solutions. The heavy hitters, many of them Peter Glynn's students, try to keep us on a solid foundation.

6 ON MY RESEARCH

My own topics, in the order that I acquired them, are all related to Monte Carlo simulation experiments: univariate input modeling, random variate generation, multivariate input modeling, variance reduction, output analysis, stochastic root finding, and stochastic optimization. Appending U(0,1) random-number generation and the initial-transient problem to the list would encompass most presentations within the WSC Analysis Methodology track.

Below I discuss two topics, with the hope that the research discussion is the focus more than the topics.

6.1 Input Modeling

Based on my short practitioner career, my first independent research focused on the following practitioner context. I was trying to predict the effect of automating Blue Cross / Blue Shield claims processing. For the input model, I needed service-time distributions for various human operations, routing probabilities, travel-time distributions, mail-room arrival volumes, and mail-room arrival times. As was not atypical, I had arrived midweek and needed results for Monday's big presentation. The service-time distributions were difficult because they obviously would affect the conclusion, there were many of them, and because there were no data sets.

My only academic training was the classical approach of hypothesizing a family of distributions, fitting parameters, and testing goodness of fit. With many distributions, no data sets, and not enough time, I needed a quick-and-dirty solution. Like many practitioners before me, I scurried from operation to operation, asking for minimum, maximum, and most-likely service times, and then assumed a triangular density function. Some sensitivity analysis was possible, because over the weekend I had the main-frame computer to myself. Not surprisingly, the results were sensitive to the maximum. Worse, even I didn't know what I meant by the maximum time needed to check a claim form for errors.

My first research funding, from the Office of Naval Research, was to create four- and five-parameter families of distributions. I had only a vague problem statement. I wanted to obtain a distribution family that was so general that it would be adequate for most quick-and-dirty applications. Ideally, it would cover the entire plane of third and fourth standardized moments. Ideally, it would have parameters that have meaning, like the three parameters of the triangular family. Ideally, it would have a closed-form inverse transformation for easy random-variate generation and to support correlation induction. Ideally, it would look good, in the sense that a practitioner tended to think that the fitted shape was reasonable.

The best of several attempts, written with Ram Lal, was a five-parameter generalization of the triangular family; the fourth parameter changed the shape of the left side and the fifth the right side. A referee said that we had set statistics back fifty years, leading me to shove the manuscript in a lower drawer forever. For several years I did teach the distribution in simulation classes, because I found the five-parameter approach useful. Then fitting software, such as Averill Law's UniFit and later Mary Ann Wagner's Bezier curves, arrived, and I thought that need for a five-parameter family was gone. Recently, however, Sam Kotz and Rene van Dorp (2005) recently proposed the same family in their excellent reference *Beyond Beta*.

What is the lesson? A researcher needs to develop a thick skin. Don't take referees, who are usually anonymous, too seriously. Read and think about the comments. Don't make changes only because a referee suggested them. Understand that the same idea better presented is usually publishable; any manuscript that is mathematically correct can find a home, because many journals struggle to fill their pages. But please, if an idea is bad or not new, let the manuscript die.

6.2 On Choosing a Confidence-Interval Procedure

Since the 1950s a classic simulation topic has been to estimate the standard error of the sample mean of a steady-state vector of n data points.

The topic continues to develop today, taking several directions. Solution methods have evolved, from explicit estimation of autocorrelations, batching (spaced, contiguous, partial overlapping, overlapping), ARMA time-series modeling, regenerative processes, standardized time series, and quadratic forms. Criteria have changed, from the probability of covering the mean to that and the expected half width to those two and the variance of the half width. The problem has been generalized to the estimation of any distribution property, rather than only the mean. The problem has also been generalized from the original, which assumed that a vector of data are given, to sequential methods, which assumed that each data point is available, used to update cumulative statistics, and discarded. And stopping rules yield many other variations.

The half century of no-end-in-sight research is not due to lack of competence. Most papers are correct, new, and well written. The reason is, I think, the use of two or three criteria (coverage probability, expected half width, and sometimes variance of the half width). Given the multiple criteria, most reasonable ideas are not dominated. Paper after paper includes tables of Monte Carlo results showing good performance.

Schmeiser and Yeh (2002) propose a single criterion to replace the classical multiple criteria. (The single criterion is the expected squared difference of the p value of an ideal procedure and the p value of a new proposed procedure, but that is a detail here.) Our hope is that the single criterion, or an alternative, will allow the research community to rank procedures unambiguously. Now, six years later, my impression is that no other author has adopted our single criterion. Also no author, even in a hall discussion, has attacked the single criterion.

What is the lesson? Maybe everyone is being kind by not attacking the single criterion. More likely, though, is that research communities fall into comfortable patterns. Thesis advisers know how to point students to publishable research. Students can examine earlier theses for good analogies. Editors and referees know what to expect from a submitted manuscript with the problem statement and criteria are consistent with the existing literature. Given the university reward system, for which publishing papers is important, there is little reason for a community to change.

As the years pass, however, I do find a reason to change. Adding yet another paper to a well-developed literature is less fun. Stating (and defending) new problems is more fun. And my belief that the heart of good research is creating the problem statement continues to grow stronger.

7 IN CONCLUSION: MARKETING

So what happened to the bar that the practitioner, the vender, and the researcher walked into? The purpose of the bar was only to attract you, the reader. Which leads us to the final point, that if we researchers don't work a bit harder to attract attention to our work, much of it will be ignored. And most researchers are not good at attracting attention.

Venders have always been good with catchy pronounceable names. Alan Pritsker was the king of well-known acronyms: GERT, GASP, SLAM, AWESIM. Lee Schruben's SIGMA and Steve Roberts's INSIGHT are perfect names, although both have the researchers' aversion to marketing. Dennis Pegden, Jim Henriksen, and all other serious venders have a sense of the catchy name and marketing. Would OPTQUEST have been so easily adopted by multiple venders if the code had a different name?

As a practitioner at EDS, I tried. Based on GASP II, I wrote a simulation package for insurance-claims processing. I called it SNAP: Simulation of a Network of Activities Program. The acronym was ok, but the words lacked cadence.

Researchers are getting better at names. For years I talked about "the three-step method", a 1970s idea to transform multivariate-normal random vectors to vectors having uniform (0,1) marginal distributions to vectors having arbitrary marginal distributions with a specified correlation matrix. Barry Nelson, in automating the idea, began saying NORTA, for "normal to anything". More recently he and Jeff Hong deftly named their search method "COMPASS" and a later variation "Industrial-Strength COMPASS".

Good names are a step toward better communication among practitioners, venders, and researchers. And that better communication would be good for all of us.

ACKNOWLEDGMENTS

I thank John Fowler, the WSC Program Chair, for inviting me to speak. I also thank my former and current Ph.D. students, without whom my research career would have been far less productive and enjoyable. Honggang Wang helped me with distance-based word-processing problems.

REFERENCES

Ker T. 2005. Research summary. Available via <www.livescience.com/health/050907_schizotype_creative.html> [accessed July 13, 2008].

Kotz, S., and J. R. van Dorp. 2005. *Beyond beta: Other continuous families of distributions with bounded support and applications*, World Scientific Publishing.

Leemis, L. M., and S. K. Park. 2006. *Discrete-event simulation: A first course*, Prentice Hall.

Nelson, B. L. 1983. *Variance reduction in simulation experiments: A mathematical-statistical framework*, Ph.D. dissertation, Purdue University.

Schmeiser, B. and Y. Yeh. 2002. On choosing a single criterion for confidence-interval procedures. *Proceedings of the Winter Simulation Conference*, eds. E. Yucesan, C.-H. Chen, J.L. Snowdon, and J.M. Charnes, 345–352. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHY

BRUCE W. SCHMEISER is a Professor of Industrial Engineering in the College of Engineering at Purdue University. He was WSC Program Chair in 1983 and WSC Board Chair in 1988–1990. He is an INFORMS and IIE fellow. He is the recipient of IIE's 2004 David F. Baker Distinguished Research Award, the INFORMS College on Simulation's

1997 Distinguished Service Award, and various teaching awards.