

OM Forum

Three Rs of Operations Management: Research, Relevance, and Rewards

Jan A. Van Mieghem

Kellogg School of Management, Northwestern University, Evanston, Illinois 60208, vanmieghem@kellogg.northwestern.edu

This article presents three personal experiences and reflections on the process and objective of research in operations management. This invited OM Forum contribution is based on my Manufacturing and Service Operations Management Distinguished Fellow inaugural lecture given at Columbia University on June 18, 2012.

Key words: operations management; research; relevance; rewards; research relevance matrix

History: Received: June 19, 2012; accepted: November 3, 2012. Published online in *Articles in Advance* December 14, 2012.

There are those days in the life of a scholar that follow what we call a “Poisson process”: they are rare and unexpected. February 16, 2012, was one such day for me. That was the day Manufacturing and Service Operations Management Society President Steve Gilbert informed me that I had been elected a Distinguished Fellow. I was, and still am, touched and humbled by this recognition of my work on the generalized $c\mu$ rule and on the role of flexibility and capacity in networks. Rather than describing that work here, this essay instead presents some reflections on the objectives and rewards of research in operations management. To give some perspective on the origins of these reflections, I share here three personal stories that illustrate the “three Rs of operations management (OM).” And during this process I will take the opportunity to thank several individuals and institutions that were instrumental in my professional development.

1. Research

The first story is about what research means to me personally, and the convoluted way in which I ended up pursuing it professionally. During my study years in electrical engineering at the University of Leuven in Belgium, I wore bunny suits in clean rooms while studying semiconductors and microprocessor design, and a lab coat in chemistry. But that popular notion of research that centers around lab experiments did not appeal to me. As I continued my studies, however, I came to understand that there is another notion of research that is broader and extends beyond the lab. Indeed, etymology is illuminating, and the word *research* is derived from the French *recherche*—to seek

out. In my native language of Dutch, *onderzoek* goes even further: to search under. Research is searching for answers, but how and where?

In 1989, a fellowship from the Belgian American Education Foundation brought me to Stanford University. My original aim was to study cryptography, but the leading professor in that field was pursuing other activities and so I moved into quantum mechanics and engineering administration. This dual focus reflected my admiration of beauty in mathematical theory coupled with an aspiration for relevance.

The Greek philosophers discussed the connection among beauty, the “ideal,” and mathematics. Classic beauty possesses elegance, symmetry, and proportionality. Now add parsimony, logic, and generality, and we arrive at the beauty in mathematics and quantum mechanics, for which I was a teaching assistant at Stanford. Rationalists and romanticists made an even further connection: “Beauty is truth, truth beauty,—that is all,” according to Keats (1820). If research is about searching for the truth, then isn’t it therefore also about beauty?

But how did I move into the type of research I have pursued for the past 20 years? Although in line with the dual focus described earlier, the specific choice of the field was rather incidental: My friend and fellow Belgian, Christoph Crombez, who was pursuing a Ph.D. in Business at Stanford at the time, told me about the Ph.D. program track in operations. (As an electrical engineering student, I had never heard of such a program before. Even today, our Ph.D. programs remain hidden gems, and we should think of how we can straighten and illuminate the pathways to our field to attract more talent.) I still recall the doctoral brochure describing the program as “Applying

mathematical models to improve decision making.” While the title “doctor of philosophy in business” sounded somewhat contradictory, the program itself appealed to me. So, after completing my masters degree in electrical engineering I moved into research at the business school with J. Michael Harrison as my advisor and mentor. And hence my quest for truth, and therefore beauty, began.

Years of experience have shaped my current view on what is “desirable” research. Desirable research seeks to discover truth by creatively building a thing of beauty. I have a deep passion for building, for beauty, and for creativity. Simply stated, the research process starts with a question to which you build an answer by creatively putting pieces together. Sometimes you may intuitively conjecture the answer and the task is in constructing the argument and validating the conjecture. Other times you may not know the answer but teach yourself by trial and error in putting pieces together. There is wonder in that process of building knowledge almost out of nothing. And there must be beauty in that process, and it must result in a deliverable. (The rationalization of “I also thought about that but never wrote it up” doesn’t count.)

Beauty is in the eye of the beholder, and here is one example of beauty in the research process and deliverable: Harrison (1988) introduced powerful theory to apply Brownian motion to operations and economics. Application of that theory typically requires the stationary distribution of reflected Brownian motion (RBM). Given that its analytic specification often is intractable, numerical methods were developed to compute the stationary distribution of RBM. To validate the accuracy of these numerical methods, it is useful to have exact analytical test cases. My advisor told Liz Schwerer and me that perhaps the Brownian model of a closed, balanced three-station queuing network could be calculated exactly using conformal maps. Liz and I pursued this challenge as a summer project and constructed a Schwarz–Christoffel transform to map the complex upper half plane onto the interior of a triangle (Figure 1). The queue length process of a closed, balanced network can be approximated by driftless RBM in a triangle. Inverting

that process allowed us to develop explicit formulas for the stationary distribution of RBM in a general triangle and for various performance metrics of a closed, balanced three-station queuing network.

In my eyes, there was beauty in this process and in its product (Schwerer and Van Mieghem 1994). Yet, I did not include that paper in my doctoral dissertation, and almost 20 years later I still believe that was the right decision: Although this paper was the first to “solve” a Brownian model of a non-product-form closed queuing network, it has only garnered four citations in Google Scholar. In operations management, desirable research often needs to go beyond discovering truth by creatively building a thing of beauty. This brings me to my second, much shorter, story about relevance.

2. Relevance

Something is relevant if it has bearing on or a connection with the subject at hand. Given that I was in a business school, it seemed logical that my dissertation connect to dollar signs, which brought me to the study of investment in flexibility.

Often, relevance implies importance. As such, relevance is a relative concept, and one must ask, relevant to whom? Research must be relevant to the researcher for knowledge to advance. Yet solely aspiring internal relevance carries the risk that research becomes an intellectual exercise in self gratification, i.e., the quintessential “ivory tower” syndrome. To survive and prosper, research must be externally relevant, which in our field means relevant to students and practitioners. To meet that challenge, we can adopt at least four research styles that each aspire to relevance in their own way. The dominant paradigm has changed over time, as the following quick historical review shows.

Originally, our field oozed with relevance to the military to help find better logistics and training schedules during World War II. Paraphrasing Frederick Winslow Taylor, this research style aims to find the best way of managing an operation and is grounded in applications with an abundance of data.

Figure 1 Beauty in Research Is in the Eye of the Beholder


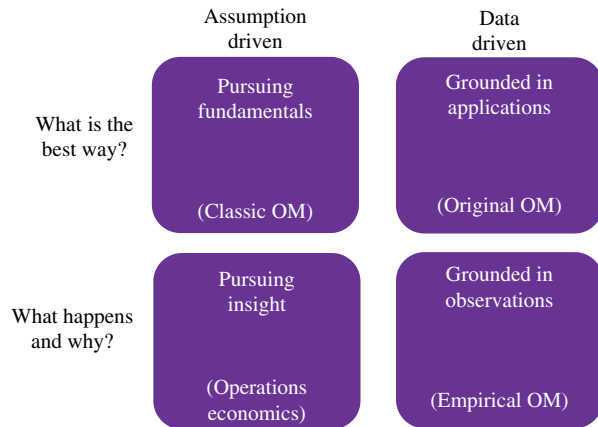


Figure 2 A Research Relevance Matrix for Operations Management

This research style occupies the upper right of the research relevance matrix in Figure 2.

After the war, these same mathematical optimization techniques were used to analyze and optimize complex problems in industry. As the complexity grew, more sophisticated techniques were needed and developed, and more fundamental questions were posed, such as, What is the best way to manage a multiechelon inventory system? To answer that question precisely, Clark and Scarf (1960) defined and analyzed a model, a carefully controlled setting that is specified via a set of assumptions. This has been the classic research paradigm in operations for decades, occupying the upper left of the matrix in Figure 2. This paradigm is powerful due to its generality, and its level of abstraction makes it the most likely research style to create a thing of beauty. However, pursuing generality and abstraction for its own sake runs the risk of losing external relevance. As a consequence, the 1980s saw several researchers leaving the field and porting the operations toolkit to make influential contributions in marketing, finance, and economics.

In the 1990s, our field enjoyed a renaissance by embracing the positive research approach from the social sciences. In the early 1990s (during my doctoral education), the field started recognizing that operations are managed and executed by several individuals, each having different motives and information sets. The research question moved beyond prescriptive or normative optimality to the descriptive questions of what happens in multiagent operations and why. Initial attempts to address these questions were either clumsy and not beautiful, or pedestrian and uninteresting—and hence not successful; but this changed in the second half of the 1990s under the leadership of Cachon and Lariviere (1999) and other contemporaries. I believe this movement into what I call operations economics (the lower left of the matrix in Figure 2) has been extremely successful, as

reflected by the enormous growth of operations faculty in business schools.

With every success, however, comes a risk: Will the field eventually be subsumed by the imperial science of economics altogether? Wikipedia defines economics as the social science that analyzes the production, distribution, and consumption of goods and services. So is operations nanoeconomics, or does it have a different focus? I would submit that the answer is somewhere in the middle, but it does raise the need to articulate our focus, differentiation, and relevance. One could argue that our fields' focus and differentiation lies in designing, improving, and understanding processes. One should also ask whether our Ph.D. students can excel at both operations and economics, or whether we shall see more specialization or a fracture in the field? Time will tell.

The last decade, however, has been promising in reemphasizing external relevance with a dramatic and welcome increase in data-driven research. Rigorous empirical research (the lower right of the matrix in Figure 2) is necessary to apply the scientific method to our field. It can estimate and test our theory like, for example, the structural estimation of the news vendor model by Olivares et al. (2008), or the search of the bullwhip effect by Cachon et al. (2007). Empirical research also can stimulate new theory like, for example, DeHoratius and Raman's (2008) work on inventory inaccuracies. Finally, we have closed the loop by a return to, and an increase in, the original mode of research (upper right in the matrix) that is grounded in applications yet generalizable and implemented in other practice settings and in widely used software products. For example, the work by Graves and Willems (2000) at Eastman Kodak, Rao et al. (2000) at Caterpillar, and Gallien and Caro (2010) at Zara evaporate any worries about relevance.

There is more good news: In addition to the multitude of research styles in the research relevance matrix, we also have more channels at our disposal that facilitate relevance. The traditional media of journals, textbooks, and cases remain, but the Internet has made journals and cases globally accessible by anyone in real time. Although traditional publishers resist, electronic textbooks are poised to follow and present a wonderful opportunity for us, the content creators and owners.

The Internet is also creating new channels that facilitate relevance: There now are various blogs by operations professors that have growing readership, both among academics, practitioners, and students. These blogs also can be a wonderful source for research ideas. New blogs are announced on Twitter and Facebook, which also allow you to follow updates and opinions of those colleagues you are

interested in. Many institutions now feature a “popularized papers” website that summarizes academic research in an accessible manner. Third-party organizations like the Henry Stuart Talks now publish “animated audio visual presentations by leading world experts—advanced content in a user friendly format” (<http://hstalks.com/>, accessed November 20, 2012).

How can we use these channels to turn research and relevance into a rewarding process? This leads to my last story.

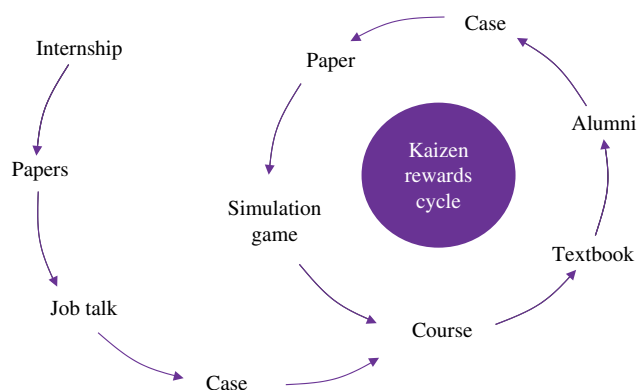
3. Rewards

In operations strategy we teach the benefits of alignment between the goal and the means. Practicing what we preach and aligning research and teaching can lead to rewarding experiences. Picking up my personal story: During my doctoral study, I was an intern at Seagate Technologies working on a global capacity strategy for introducing new disk drives. Disk-drive manufacturing has a steep learning curve, and the question was how that learning curve can be integrated with capacity planning of the global network.

As with many projects, however, the starting question need not be the ultimate one. My observations at Seagate led to my doctoral dissertation and research papers on the role of flexibility and capacity in networks. While presenting my findings at a job talk at INSEAD, Christoph Loch suggested I turn the paper into a simple case exercise, which became a cornerstone of the operations strategy MBA course that I developed at the Kellogg School of Management of Northwestern University. At Kellogg, the institution that has graced my business card for 17 years, I have had the good fortune of working with my colleague and friend Sunil Chopra to build a fine research group. This growth strategy required a growth in elective courses, and I started in 1998 to develop a full-quarter operations strategy MBA course.

Ten years later I summarized the course into a textbook (Van Mieghem 2008) where each chapter was accompanied by a mini-case. One of those cases, titled “Mexico–China,” resulted from collaboration with Cort Jacobi, an alumnus and consultant. After assigning the case to the students, Cort discussed in class the consultants’ analysis and recommendation to their client. The analysis showed the simulated total landed cost as a function of the percentage sourced from Mexico and from China. The curves were so nicely convex that they screamed out for analysis, which resulted in the research paper Allon and Van Mieghem (2010a). To teach our findings, we designed an Internet simulation game, described in Allon and Van Mieghem (2010b), that has become a standard module in courses at Kellogg and at other

Figure 3 A Process of Research, Relevance, and Rewards



institutions. These changes must now be reflected in a new edition of the textbook, and may lead to a new kaizen rewards cycle, as shown in Figure 3.

In summary, these three stories illustrate the aspiration of merging desirable research, which discovers truth by creatively building a thing of beauty, with relevance to operations management. With proper foresight and some luck, this can result in a virtuous kaizen rewards cycle and the best job in the world.

References

- Allon G, Van Mieghem JA (2010a) Global dual sourcing: Tailored base-surge allocation to near- and offshore production. *Management Sci.* 56(1):110–124.
- Allon G, Van Mieghem JA (2010b) The Mexico–China sourcing game: Teaching global dual sourcing. *INFORMS Trans. Ed.* 10(3):105–112.
- Cachon GP, Lariviere MA (1999) Capacity choice and allocation: Strategic behavior and supply chain performance. *Management Sci.* 45(8):1091–1108.
- Cachon GP, Randall T, Schmidt GM (2007) In search of the bullwhip effect. *Manufacturing Service Oper. Management* 9(4):457–479.
- Clark AJ, Scarf H (1960) Optimal policies for a multi-echelon inventory problem. *Management Sci.* 6(4):475–490.
- DeHoratius N, Raman A (2008) Inventory record inaccuracy: An empirical analysis. *Management Sci.* 54(4):627–641.
- Gallien J, Caro F (2010) Inventory management of a fast-fashion retail network. *Oper. Res.* 58(2):257–273.
- Graves SC, Willems SP (2000) Optimizing strategic safety stock placement in supply chains. *Manufacturing Service Oper. Management* 2(1):68–83.
- Harrison JM (1988) Brownian models of queueing networks with heterogeneous customer populations. Fleming W, Lions PL, eds. *Stochastic Differential Systems, Stochastic Control Theory and Applications* (Springer-Verlag, New York), 147–186.
- Keats J (1820) “Ode on a Grecian Urn.” Elmes J, ed. *Annals of the Fine Arts*, Vol. iv (Sherwood, Neely, and Jones, London), 638–639.
- Olivares M, Terwiesch C, Cassorla L (2008) Structural estimation of the newsvendor model: An application to reserving operating room time. *Management Sci.* 54(1):41–55.
- Rao U, Scheller-Wolf A, Tayur S (2000) Development of a rapid-response supply chain at caterpillar. *Oper. Res.* 48(2):189–204.
- Schwerer E, Van Mieghem JA (1994) Brownian models of closed queueing networks: Explicit solutions for balanced three-station systems. *Ann. Appl. Prob.* 4(2):448–477.
- Van Mieghem JA (2008) *Operations Strategy: Principles and Practice* (Dynamic Ideas, Belmont, MA).