

OM Forum

What Is Interesting in Operations Management?

G rard P. Cachon

The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104,
cachon@wharton.upenn.edu

This essay discusses my view of the essential characteristics of interesting research in general and in operations management in particular. It is based on my Manufacturing and Service Operations Management Distinguished Fellows presentation given at the University of Michigan, June 27, 2011.

Key words: operations management; interesting; impact

History: Received: September 5, 2011; accepted December 8, 2011. Published online in *Articles in Advance* February 28, 2012.

What is interesting research? Is it merely in the eye of the beholder? Or is there something more systematic about what is interesting? We all want to write, and editors should only want to publish, interesting research. And we all (hopefully) think our own research is interesting. But what will others find interesting? I claim that there is one key ingredient to “interesting.” Interesting means unexpected—interesting research piques your curiosity, it induces a pause for contemplation, and most importantly, it contradicts how you think about the world. In this essay, I develop this idea further and apply it to the literature in operations management. For other work that identifies “unexpected” as a critical component of “interesting,” applied in the domain of sociology, see Davis (1971).

Unexpected and Expected

Interesting research reveals a new perspective on the familiar. It poses a question that has not been asked before, or it follows an accepted question with a new answer, an answer that is orthogonal to those that preceded it. For example, to the question “What is the shape of the Earth?” Columbus countered with “round,” a sharp departure from the common wisdom of the time.

Confirming what is expected to be true is simply not interesting. For example, do we care or should we care if a paper reports, even if through sophisticated means, that callers will balk more frequently if they have to wait longer? Or that employees work faster if they are paid more? How could it be any other way? One might retort with “But isn’t it important to show that people actually put in more effort when they are

paid more?” I do not think so. It would indeed be interesting to show the opposite because that would violate our presumptions about worker motivation. But demonstrating or proving the expected, no matter how hard it is to do, remains bland, pedestrian, and just flat out uninteresting.

Given that interesting equals unexpected, it is important to be clear about what is meant by “expected.” Expected is what is assumed to be “true,” “known” or “given” at the time the research is developed, or what could be known with minimal effort. For example, if the expectation is that a linear program is hard to solve, then the invention of a method that reduces the necessary search space and is guaranteed to find the optimal solution (e.g., George B. Dantzig’s simplex method) is unexpected. After that, if the expectation is that the solution method must traverse the boundaries of the feasible region (as in the simplex method), then a method that approaches the optimal solution from the interior would be unexpected (as in interior point methods; e.g., Karmarkar 1984).

A Formula for Interesting

How does one create an unexpected, and therefore interesting, result? It has been my experience that the main idea of an interesting paper can be described with the following template:

“What was thought to be *X* is really *Y*.”

Uninteresting papers are unable to offer a short, simple, and precise version of the above template. The “*X*” and “*Y*” are apparent in an interesting paper (and they are apparent due to the careful writing of the authors—interesting papers do not just happen, they

are crafted so that it is clear to the reader why they are interesting). Fortunately, as the following examples demonstrate, there are many to choose from. The list below is by no means exhaustive (nor meant to be).

What Was Thought to Be Exogenous Is Really Endogenous

This is a powerful and often-used approach to developing interesting research. To illustrate, consider demand uncertainty at each level of a supply chain (e.g., manufacturers, wholesalers, and retailers). Where does that uncertainty come from? It generally was assumed to be exogenous, but Lee et al. (1997) provide four reasons why it can be endogenous. Not only endogenous, it amplifies as one moves up the supply chain, a phenomenon labeled the “bullwhip effect.” (What was thought to not exhibit a pattern does have a pattern, i.e., variances increase in a particular way.) Hence, much of the demand uncertainty within a supply chain is actually self-inflicted, caused by the actions of the firms themselves, and therefore can be reduced by prudent management.

Porteus (1985) provides an early example of this approach with the economic lot-sizing problem—choose a set of production quantities to minimize holding and setup costs, where a fixed setup cost is incurred whenever a production lot is started. Rather than assume the setup cost is given (i.e., exogenous), Porteus (1985), motivated by Japanese manufacturing practices, argues that a firm can exert effort to reduce its setup costs (i.e., the setup cost can be endogenous).

Recent examples of converting “exogenous” into “endogenous” include the highly influential work on strategic consumers: although it was assumed that operating decisions had no direct impact on demand, in fact, consumer demand responds directly to the operational choices the firm makes. With this new perspective we learn that ignoring strategic behavior can lead a firm to carry too little inventory (Dana and Petruzzi 2001), or that optimal prices may rise over a selling season (Su 2007), or that committing to never mark down prices may help a firm (Su and Zhang 2008).

What Was Thought to Be Complex Is Really Simple

Return to the lot-sizing problem but consider a more complex version, the multiechelon version with one warehouse and N retailers. This is a very difficult problem, and the optimal policy is unknown. However, Roundy (1985) provides a different view on the problem—instead of trying to identify the optimal policy, identify a simple policy that has provably good performance. In other words, what was thought to be complex (a problem with an unknown optimal policy) is really simple (a simple policy works quite well).

Although the work of Roundy (1985) exemplifies this approach, it is by no means the only example.

In fact, many papers in operations management follow this path. For instance, Lariviere and Porteus (2001) take a problem (selling to a newsvendor) that is ill behaved in general (i.e., complex) and show that it is well behaved (i.e., simple) for a large class of distribution functions. Gallego and van Ryzin (1994) show that a complex stochastic dynamic problem (what sequence of prices should be chosen to maximize revenue from a finite set of inventory sold over a fixed period of time) can be well approximated by a deterministic counterpart. Their insight led to the unexpected conclusion that simple, fixed price policies may perform nearly as well as a fully dynamic pricing policy.

What Was Thought to Be Simple Is Really Complex

In the newsvendor problem a decision maker chooses a quantity of inventory to order before facing stochastic demand—order too little, and there will be lost sales, but order too much, and there is costly leftover inventory. The optimal policy is known, but how do people actually make newsvendor-type choices? This is a choice task with random outcomes, and there are a number of established theories for how people make choices under uncertainty. For example, they could be risk averse or risk seeking. Unfortunately, these theories do not describe actual decisions in this setting very well (see Schweitzer and Cachon 2000). (I fully recognize that it is presumptuous to assume that my papers are interesting, so I only mention a few.) Another theory is needed. What was thought to be simple (an established theory can provide the answer) is really complex (a new theory is needed).

What Was Expected to Be a Small Effect Is Really a Large Effect

Say a paper demonstrates that what you would think would be a small effect is actually a large effect. That’s interesting. And there are number of good examples of this approach.

What is the impact of traffic congestion on public health? To answer this question, Currie and Walker (2011) exploit the adoption of the E-ZPass toll system on highways in Pennsylvania and New Jersey. E-ZPass is a radio-frequency identification tag that reduces toll plaza congestion by allowing vehicles to pay their toll without stopping. Less congestion leads to less pollution, which can lead to better health. Specifically, does the reduction in congestion in a toll plaza have a measurable impact on pregnancy outcomes of people who live within two kilometers of the toll plaza? One could easily assume (as I would) that this effect, if it exists, would be too small to measure, but they find a substantial effect: premature births decreased by 10.8% and low-weight births decreased by 11.8%.

As a Ph.D. student, my first mentor, Colin Camerer, told me that there are three kinds of papers in this world. The first, and maybe the majority, are the ones that should never have been written—the question is not interesting, and the answer is not compelling. The second are the papers that you are glad somebody else wrote. You wanted to know the answer to the question, but you are content that somebody else was willing to do all of the work to provide it. These can be interesting, just not worth the personal effort to achieve the answer. And the third are the papers that you wish you wrote. There are not many of those. Currie and Walker (2011) is definitely in my “I wish I wrote it” category. The next one is as well.

Jordan and Graves (1995) study the question of production flexibility. Say you have a set of production facilities and a set of products to produce. Demand is uncertain, and each facility has a fixed (and limited) capacity. You can invest to give a plant the ability to make more than one product. But this flexibility is expensive, so how much do you need? An intuitive answer is that a lot of flexibility is worth much more than a little bit. It turns out, we now know, thanks to Jordan and Graves (1995), that what was assumed to be a small effect (a little bit of flexibility provides only a little benefit) is actually a very large effect (a little bit of flexibility, done right, provides essentially the same benefit as complete flexibility).

There are other examples of this approach. Fisher and Raman (1996) show that smartly exploiting a little bit of early season sales can generate a substantial profit increase (whereas it was assumed that a small sample of data would have little value, it can actually have a large value). Next, although it was assumed that there would be no value to operational flexibility when demands across products are perfectly correlated, Van Mieghem (1998) shows that the value in this case can be positive (when the products have different margins).

What Was Thought to Be a Large Effect Is Really a Small Effect

In the early 1990s there was considerable interest in the use of information technology to improve supply chain performance. The key technologies of the time were bar coding and electronic data interchange. (The internet was only starting to gain momentum.) Everyone assumed (including myself), that exploiting the data these systems created and shared across the supply chain would lead to tremendous improvements in inventory performance. But Cachon and Fisher (2000) show that what was expected to be a large effect (the reduction in inventory through the use of shared information across the supply chain) was actually a small effect (it is much more effective to move goods faster than it is to be smart with how you move them).

A corollary could be “What was thought to be common is not so common.” Cachon et al. (2007) take this approach with the bullwhip effect.

What Was Thought to Be a Large Effect Is Really Much Larger

Quick response has been shown to provide substantial value to firms (Iyer and Bergen 1997) because of its ability to better match supply with demand. However, Cachon and Swinney (2009) show that ignoring strategic consumer behavior *underestimates* the value of quick response (what was thought to already be a substantial amount) by as much as 500%.

What Was Thought to Be Easy Is Really Hard

Ask students in teams of four to make inventory decisions in a serial supply chain. Demand is constant, and lead times are constant. This should not be terribly hard, yet, as first observed by Sterman (1989) and has been demonstrated countless of times, this is a very difficult task—subjects make a number of errors that lead to large mistakes and terrible operating performance.

What Was Assumed to Not Be a Problem Is Really a Problem

Nearly the entire literature on inventory management assumes the decision maker knows how much inventory they have. But is that true? In fact, DeHoratius and Raman (2008) show that what was assumed to not be a problem (inventory records are accurate) is really a problem (in fact, errors are common enough to lead to poor decisions).

What Should Improve Performance Really Harms Performance

Consider a large number of commuters who want to drive to work and have several possible routes to choose from in a network of roads. The congestion on a particular route depends on how many drivers choose the route, so one would think that things can only get better when you add capacity to the network (another route) while holding the number of commuters constant. But Braess (1968) shows that this might not be the case—due to self interested behavior, adding capacity may actually increase commuting times for all.

Is “Interesting” Everything?

Are we done? Is “interesting” sufficient, the exclusive goal, everything that we want? Although interesting is desirable, and surely necessary, I contend that it is not sufficient. We want research to also be “important”—important research is useful, either for the creation of more research or, better still, for the utility of society. To be useful, research generally

must be broadly applicable, or at least applicable in a domain of significance (e.g., healthcare rather than just dog grooming). A result is not important if it applies only in a special case or for unlikely or unreasonable parameters. When a result is useful and broadly applicable, it is likely to change behavior, the behavior of other scholars (e.g., they build upon the research or they change the questions they work on) or, ideally, the behavior of those outside academia (managers, policy makers, etc.).

Is that it? We only want interesting and important research? What about being correct? Shouldn't a result also be "true"? Strictly speaking, interesting does not need to be true. Take the famous Hawthorne effect. In the 1920s a series of experiments were conducted at Western Electric's Hawthorne plant to determine whether better lighting improves worker productivity. The surprising finding was that productivity improved whenever there was a change in lighting—lighting did not matter per se, but rather that workers knew that management was paying attention to them. This is indeed an interesting result, maybe even a fascinating result, but it may not be correct (see Levitt and List 2011 for details).

But if a result need not be true, doesn't that open up the possibility that every cockamamie idea can be labeled "interesting"? No. Although an interesting idea does not need to be true, it does need to be plausible. The Hawthorne effect gained notoriety because it was both unexpected and plausible—it had some prima facie validity. Most importantly, it had data (albeit, poorly collected and analyzed data). So, by any yardstick, the Hawthorne effect was interesting and important—it surely changed behavior of scholars and practitioners.

Conclusion

Interesting research raises more questions than it answers. It is controversial. It invokes responses like "that can't be true" or "this is obviously incomplete." Interesting research should initially leave the reader a little discontent, unnerved, or motivated to prove it wrong or at least incomplete. This is why it can be hard to publish interesting research, and really interesting research is rarely accepted immediately.

Interesting research that is also "important" changes behavior, i.e., it yields the ever so desirable "impact":

$$\text{Interesting} \times \text{Important} = \text{Impact.}$$

References

- Braess, D. 1968. Über ein paradoxen der verkehrsplanung. *Unternehmensforschung* 12 258–268.
- Cachon, G. P., M. Fisher. 2000. Supply chain inventory management and the value of shared information. *Management Sci.* 46(8) 1032–1048.
- Cachon, G. P., R. Swinney. 2009. Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Sci.* 55(3) 497–511.
- Cachon, G. P., T. Randall, G. M. Schmidt. 2007. In search of the bullwhip effect. *Manufacturing Service Oper. Management* 9(4) 457–479.
- Currie, J., R. Walker. 2011. Traffic congestion and infant health: Evidence from E-ZPass. *Amer. Econom. J.: Appl. Econom.* 3(1) 65–90.
- Dana, J., N. Petruzi. 2001. Note: The newsvendor model with endogenous demand. *Management Sci.* 47(11) 1488–1497.
- Davis, M. S. 1971. That's interesting! Towards a phenomenology of sociology and a sociology of phenomenology. *Philos. Soc. Sci.* 1(2) 309–344.
- DeHoratius, N., A. Raman. 2008. Inventory record inaccuracy: An empirical analysis. *Management Sci.* 54(4) 627–641.
- Fisher, M., A. Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Oper. Res.* 44(1) 87–99.
- Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* 40(8) 999–1020.
- Iyer, A. V., M. E. Bergen. 1997. Quick response in manufacturer-retailer channels. *Management Sci.* 43(4) 559–570.
- Jordan, W. C., S. C. Graves. 1995. Principles on the benefits of manufacturing process flexibility. *Management Sci.* 41(4) 577–594.
- Karmarkar, N. 1984. A new polynomial-time algorithm for linear programming. *Combinatorica* 4(4) 373–395.
- Lariviere, M. A., E. L. Porteus. 2001. Selling to the newsvendor: An analysis of price-only contracts. *Manufacturing Service Oper. Management* 3(4) 293–305.
- Lee, H. L., V. Padmanabhan, S. Whang. 1997. Information distortion in the supply chain: The bullwhip effect. *Management Sci.* 43(4) 546–558.
- Levitt, S. D., J. A. List. 2011. Was there really a Hawthorne effect at the Hawthorne plant? An analysis of the original illumination experiments. *Amer. Econom. J.: Appl. Econom.* 3(1) 224–238.
- Porteus, E. L. 1985. Investing in reduced setups in the EOQ model. *Management Sci.* 31(8) 998–1010.
- Roundy, R. 1985. 98%-Effective integer-ratio lot-sizing for one-warehouse multi-retailer systems. *Management Sci.* 31(11) 1416–1430.
- Schweitzer, M. E., G. P. Cachon. 2000. Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Sci.* 46(3) 404–420.
- Sterman, J. D. 1989. Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Sci.* 35(3) 321–339.
- Su, X. 2007. Intertemporal pricing with strategic customer behavior. *Management Sci.* 53(5) 726–741.
- Su, X., F. Zhang. 2008. Strategic customer behavior, commitment, and supply chain performance. *Management Sci.* 54(10) 1759–1773.
- Van Mieghem, J. A. 1998. Investment strategies for flexible resources. *Management Sci.* 44(8) 1071–1078.