

The Power and the Limits of Industrial Organization

Severin Borenstein^{1,2} 

Published online: 7 March 2016
© Springer Science+Business Media New York 2016

Abstract In the last few decades, mainstream industrial organization has developed very powerful tools for analyzing firm behavior within a fairly narrow framework. In this essay, I argue that while these tools are important, their development has crowded out other important aspects of firm behavior from the mainstream of IO. I conclude that in order to continue to grow the power of IO to predict firm behavior, the field will need to embrace aspects of firms that have moved away from IO in recent years, including innovation, boundaries of the firm, regulation, managerial incentives and decision biases, and others.

Keywords Boundaries of the firm · Innovation · Managerial bias · Oligopoly · Regulation

I am very honored to be named the 17th distinguished fellow of the Industrial Organization Society.¹ The 16 previous fellows include two Nobel Prize winners—my Berkeley colleague, Oliver Williamson, in 2005 and my colleague in graduate school, Jean Tirole, in 2014—and my dissertation advisors Dick Schmalensee and Paul Joskow, as well as my close friend since grad school, Carl Shapiro. This is daunting company to find myself in.

¹ This essay is based on my Keynote Address at the 13th Annual International Industrial Organization Conference, Boston, MA, April 25, 2015.

✉ Severin Borenstein
borenste@haas.berkeley.edu

¹ Haas School of Business, University of California, Berkeley, 545 Student Services #1900, Berkeley, CA 94720-1900, USA

² National Bureau of Economic Research, Cambridge, MA, USA

Being told of this honor led me to once again reflect on the field of Industrial Organization, and where it sits in economics. About 10 years ago, feeling a bit frustrated with my own research, and inspired by a discussion with my long-time friend and collaborator, Nancy Rose, I started asking myself what are the big questions of Industrial Organization. Then, at a few different IO conferences, I asked others—from junior to very senior faculty, and across the spectrum of theorists and empiricists.

The answers I got suggested that there is a diversity of views on whether our field should be primarily descriptive or prescriptive, but the focus of both types of responses still fell under the broad theme of “how do we model firm behavior?” as an input to “when might government intervention improve outcomes?”

In addressing these questions, I think the tools of what today is mainstream IO are very powerful, yet at the same time humbly incomplete. I emphasize the mainstream tools, because there are talented researchers working in areas that were once considered part of IO—such as innovation, boundaries of the firm, regulation, financing constraints from imperfect capital markets, and managerial incentives and decision biases, among others—but over the last few decades have become less integrated with IO. Instead, the focus in IO has increasingly been on modeling oligopoly interactions under a fairly narrow set of assumptions that exclude other drivers of firm behavior.

My current view of mainstream IO may be best described with an analogy from econometrics: I still think those tools have highly significant coefficients in describing the activities of firms, but over the years the “R-squared” I attribute to them has declined. Such “IO effects” matter, but in all the complexity of intrafirm and interfirm dynamics, focusing only on these effects does not yield great predictions. As a result, I’ve become more restrained in my claims about what mainstream IO can teach us about the world.

Maybe it is part of the normal lifecycle of an academic to realize gradually that one’s own narrow interest isn’t the only thing that makes the world go round, but that has certainly been my experience. While I find that disappointing, recognizing it has pushed me to think more broadly, to be more interdisciplinary, or at least “inter-fieldinary.” That sort of work is riskier, but I also think that it’s where the big gains will come in the future. I make no claim to big gains in my own research, but it has led me to new and interesting areas.

Much of my own work has been in energy, which is where I want to focus to illustrate my points. Energy markets have been so attractive to me in part because they are relatively simple: in many cases just commodity markets, though often with firms that are large enough to behave strategically. That’s not to say there aren’t some complexities, but there is much less of the product differentiation, reputation, advertising, branding, and related issues that come up in many other industries. I would argue that this is where the application of the core IO ideas is most straightforward.

Of course, as in all scientific study, the real test of the tools of IO is their predictive power. Can they tell us what will happen in the future? Can they predict “out-of-sample”?

The deregulation of electricity markets in the late 1990s illustrates both the strengths and the limits of mainstream IO analysis. Work that Jim Bushnell and I did in the couple of years prior to the 1998 opening of the California electricity market used a Cournot–Nash equilibrium model to conclude that even firms with relatively small market share could exercise substantial market power.² That work, and a related paper, Borenstein et al. (1999), showed that traditional concentration measures—including the US Department of Justice merger guidelines on which the Federal Energy Regulatory Commission was basing its policy—were likely to understate substantially the risk of significant market power.

While there is not completely unanimous agreement that market power played a large role in the 2000–2001 California electricity crisis, my own reading of the ex post analyses is that they support our predictions quite compellingly.³ Still, there were many aspects of the California electricity market—both prior to and during the crisis—that did not fall in line with standard IO analysis. And similar issues have arisen in other electricity markets.

For example, prior to the electricity crisis, many California sellers that appeared to lack market power were bidding into the uniform-price wholesale electricity auction at prices that were substantially above any plausible measure of their marginal costs. The explanation that we heard more than once from generators was that they needed to bid above marginal cost in order to recover their fixed costs. It is tempting to dismiss statements so at odds with economic incentives as disingenuous; but for many firms and in many of the bidding hours, there does not seem to be any explanation that is well grounded in economic optimization.

A very similar puzzle has been presented and analyzed rigorously in a different electricity market setting by Hortaçsu and Puller (2008). They show that many electricity suppliers in the Texas market leave substantial money on the table through departures from optimal bidding. Some have tried to defend the virtue of IO analysis by pointing out that Hortaçsu and Puller find fewer large firms making egregious errors. Still, some of the large firms seem to be far from efficient bidding, and there are numerous smaller firms missing out on more than \$1000 per hour, appearing to make it well worth learning more economics. Are these firms just confused, or are they optimizing in an environment that is more complex than the simple oligopoly IO model would suggest?

When firm behavior departs from simple IO-model predictions, a common response is that time will weed out the less capable firms, so model predictions will become more accurate. While there is surely something to this argument, I don't think that the body of empirical work suggests that IO analysis explains the behavior of even long-lived firms particularly well. There are plenty of examples of anomalous behavior among firms and industries that have been around for decades. Again, that's not necessarily because IO analysis is incorrect, but because it is focusing on a rather narrow aspect of what drives firm behavior.

² The paper, Borenstein and Bushnell (1999), was first presented at a conference in March 1997, a year before the deregulated California market opened. An earlier paper, Borenstein et al. (1996), laid out the theoretical application of a Cournot equilibrium analysis to electricity markets.

³ See Borenstein et al. (2002), Joskow and Kahn (2002) and Puller (2007). For opposing views, see Harvey and Hogan (2001) and Sweeney (2002).

The low R-squared of mainstream IO is in part just a result of the fact that firms are complex organisms with much more going on than competitive strategy, or even economics. Anyone who has ever worked inside a large organization, including a university,⁴ knows that economic incentives often are not the only, or even the best, lens for understanding the actions of the organization as a whole.

But it is also, I believe, a result of the recent tilt in IO to model something—anything—rigorously and precisely with the tools of mainstream IO, even if it is only a small slice of the picture. That approach can no doubt contribute real insights; but many aspects of firm behavior are likely to be overlooked. As IO has increasingly ceded analysis of the boundaries of the firm, managerial incentives and decision biases, innovation and the rest to other fields, both IO and the analysis of firms have suffered.

That is not to say that the standard tools of IO (and price theory) aren't the ones that we should reach for first in many circumstances. Besides the analysis of market power in electricity markets, other examples from energy markets that the standard tools of IO help explain include:

- why gasoline price ceilings in Georgia after Hurricane Katrina and in New Jersey after Hurricane Sandy led to shortages and queuing;
- why subsidized renewable electricity generation from zero-marginal-cost sources like wind and solar will reduce average wholesale electricity prices, which will also reduce investment in conventional generation that is needed to balance the electrical grid every second;
- why capacity-constrained pipelines that carry crude oil from the Midwest to the Gulf Coast starting in 2011 led to lower crude prices in the Midwest, but not lower gasoline prices, because pipelines that carry gasoline between the Midwest and Gulf Coast remained unconstrained;⁵
- why Saudi Arabia's decision to continue aggressive crude oil production as prices fell in 2014 may have worked out well for them if the supply curve from fracking and other unconventional—and price-taking—sources were highly elastic around \$60–\$70 per barrel as many observers thought, but instead helped create a crash when the unconventional producers were able to continue to innovate and drive their costs much lower. (Mainstream IO tools, however, are not much help in explaining why fracking innovation seems to have accelerated as the price of crude fell.⁶)

Similar examples certainly exist in other industries, though perhaps energy presents some of the cleanest cases, because of the relative homogeneity of products. While these sorts of applications are heartwarming to the IO economist, one can draw nearly daily examples of puzzling behavior from the *Wall Street Journal*, *Financial*

⁴ And any economist who has become a Dean.

⁵ See Borenstein and Kellogg (2014).

⁶ IO reasoning does suggest the reverse causality explanation: that cost reductions in fracking during 2015 were the cause of continued crude oil price declines. But those same tools suggest looking at production quantity data, which don't really support that hypothesis.

Times, *Bloomberg*, and other business publications, not to mention the trade press in major industries, where the economic thinking on which we rely seems to play a secondary role to regulatory strategy, pressures from financial markets, and managerial over-confidence. Just read the ex ante justifications for many mergers, as well as the post mortem on why they have not worked out as well as planned.

So what can IO economists do to increase the field's R-squared? I think the big gains in the next decade will come much more from broadening than from deepening: from combining an IO approach with thinking about firm behavior that is outside the narrow IO box. Of course, some economists are already doing this, but I think it is especially important now to expand the IO tent. Two directions that I think are particularly important—and, not coincidentally, particularly likely to push IO economists outside their comfort zone—are behavioral economics and innovation.

The responses of firms to the behavioral biases of consumers is a growing topic within IO,⁷ but when we talk about large companies it is still common to hear behavioral economics dismissed. Yet, when I talk to my former MBA students about managing in the constantly changing environments they face, it's clear that there is a lot of rule-of-thumb and go-with-the-gut decision making. And then add in managerial overconfidence and incentive-distorting compensation, as documented in the corporate finance literature.⁸ Without taking on the important roles that imperfect decision-making processes play in firms—even if they are just random errors and especially if they are systematic biases—I think IO hits a low ceiling on its R-squared.

Innovation is what makes companies successful. I don't just mean inventing new products but also inventing new processes, pricing, and business models, including innovations in firm boundaries. Even price-taking firms can earn extraordinary profits through innovation, as the fracking revolution in hydrocarbons has made so clear over the last decade. Also in the last decade, many Internet firms have built entire businesses around “customer acquisition,” essentially figuring out how a customer can be contacted without incurring a digital door slam. Then they have had to figure out how to “monetize” those acquired customers. And, of course, online sellers have led the innovative use of “big data” to improve all dimensions of their operations. Eventually, such increased sophistication might make mainstream IO tools more valuable again, but it is hard to imagine that IO's R-squared will even be able to hold steady over the next decade if its boundaries aren't expanded to give innovation a more central role.

Broadening IO will be more difficult and frustrating than will deepening it. Broadening will require a tolerance for less precise presentation of some ideas, at least at first, while we work out how they can be rigorously modeled and tested. But broadening is also what will continue to grow the IO toolbox in ways that increase its R-squared and thereby its value to society.

⁷ See Grubb (2015) and the other articles in that issue of the RIO.

⁸ See Malmendier and Tate (2015) on overconfidence, Murphy (2013) on incentive-distorting compensation.

Acknowledgments I thank innumerable colleagues who have shaped my thinking about industrial organization without suggesting that any of them agree with the opinions I have expressed in this essay. I'm especially grateful for hours of valuable discussions with Jim Bushnell, Joe Farrell, Ryan Kellogg, Chris Knittel, Nancy Rose, and Carl Shapiro. I have no conflicts of interest that relate to this research.

References

- Borenstein, S., & Bushnell, J. B. (1999). An empirical analysis of market power in a deregulated California electricity market. *Journal of Industrial Economics*, 47(3), 285–323.
- Borenstein, S., Bushnell, J. B., Kahn, E., & Stoft, S. (1996). Market power in California electric markets. *Utilities Policy*, 5(3/4), 219–236.
- Borenstein, S., Bushnell, J. B., & Knittel, C. R. (1999). Market power in electricity markets: Beyond concentration measures. *Energy Journal*, 20(4), 65–88.
- Borenstein, S., Bushnell, J. B., & Wolak, F. A. (2002). Measuring market inefficiencies in California's restructured wholesale electricity market. *American Economic Review*, 92(5), 1376–1405.
- Borenstein, S., & Kellogg, R. (2014). The incidence of an oil glut: Who benefits from cheap crude oil in the midwest? *Energy Journal*, 35(1), 15–33.
- Grubb, M. D. (2015). Behavioral consumers in industrial organization: An overview. *Review of Industrial Organization*, 47(3), 247–258.
- Harvey, S. M., & Hogan, W. W. (2001). Further analysis of the exercise of market power in the California electricity market. Retrieved February 21, 2016 from Harvard Kennedy School of Government web site: <http://www.hks.harvard.edu/fs/whogan/>.
- Hortaçsu, A., & Puller, S. (2008). Understanding strategic bidding in multi-unit auctions: A case study of the Texas electricity spot market. *RAND Journal of Economics*, 39(1), 86–114.
- Joskow, P. L., & Kahn, E. (2002). A quantitative analysis of pricing behavior in California's wholesale electricity market during summer 2000. *Energy Journal*, 23(4), 1–35.
- Malmendier, U., & Tate, G. (2015). Behavioral CEOs: The role of managerial overconfidence. *Journal of Economic Perspectives*, 29(4), 37–60.
- Murphy, K. J. (2013). Executive compensation: Where we are, and how we got there. In G. Constantinides, M. Harris, & R. Stulz (Eds.), *Handbook of the economics of finance* (pp. 211–356). Oxford: Elsevier Science.
- Puller, S. L. (2007). Pricing and firm conduct in California's deregulated electricity market. *Review of Economics and Statistics*, 89(1), 75–87.
- Sweeney, J. L. (2002). *The California electricity crisis*. Stanford, CA: Hoover Institution Press.