Taking the Dogma out of Econometrics: Structural Modeling and Credible Inference

Aviv Nevo and Michael D. Whinston

I n an influential paper with a catchy title, Leamer (1983) criticized the state of applied econometric practice. In the 25 years or so that have passed since the Leamer article was published, empirical work in economics has changed significantly. Without doubt, one of the major advances has been what Angrist and Pischke in this journal call the "credibility revolution." Applied work today, compared to 25 years ago, is based on more careful design, including both actual and "natural," or "quasi-," experiments, yielding more credible estimates.

Empirical work has also changed in at least two other significant ways since Leamer's (1983) article. First, econometric methods have advanced on many dimensions that allow for more robust inference. For example, nonparametric and semiparametric estimation (Powell, 1994), robust standard errors (White, 1980), and identification based on minimal assumptions (Manski, 2003; Tamer, forthcoming) are methods aimed at improving the credibility and robustness of data analysis.

A second major development, and our main focus here, has been in the improvement and increased use in data analysis of what are commonly called "structural methods"; that is, in the use of models based in economic theory. Structural modeling attempts to use data to identify the parameters of an underlying economic model, based on models of individual choice or aggregate

■ *Aviv Nevo is HSBC Research Professor of Economics and Michael D. Whinston is Robert E. and Emily H. King Professor of Business Institutions, both in the Department of Economics, Northwestern University, Evanston, Illinois. They are both Research Associates at the National Bureau of Economic Research, Cambridge, Massachusetts. Their e-mail addresses are ⟨nevo@northwestern.edu⟩ and ⟨mwhinston@northwestern.edu⟩, respectively.*

relations derived from them.[1] Structural estimation has a long tradition in economics (for example, Marschak, 1953), but better and larger data sets, more powerful computers, improved modeling methods, faster computational techniques, and new econometric methods such as those mentioned above have allowed researchers to make significant improvements. However, this development has been uneven across the various applied fields within economics. For example, structural analysis appears today in a large fraction of (but still far from all) empirical work in industrial organization, but is much less common in some other fields, such as labor economics.

While Angrist and Pischke extol the successes of empirical work that estimates "treatment effects" based on actual or quasi experiments, they are much less sanguine about structural analysis and hold industrial organization (or as they put it, industrial "disorganization") up as an example where "progress is less dramatic." Indeed, reading their article one comes away with the impression that there is only a single way to conduct credible empirical analysis. This seems to us a very narrow and dogmatic approach to empirical work; credible analysis can come in many guises, both structural and nonstructural, and for some questions structural analysis offers important advantages.

In this comment on Angrist and Pischke's article, we address their criticism of structural analysis and its use in industrial organization, and also offer some thoughts on why empirical analysis in industrial organization differs in such striking ways from that in fields such as labor, which have recently emphasized the methods favored by Angrist and Pischke.

Credible Identification and Structural Analysis: Complements, Not Substitutes

We firmly believe in the importance of credible inference, or "credible identification," and applaud the ingenious approaches to generating or identifying exogenous variation that often appear in the work using actual or quasi-experiments. Moreover, we don't think anyone (or, at least, anyone sensible) in more structurally oriented fields, such as industrial organization, would disagree with the importance of credible sources of identification. While authors of structural papers are sometimes more focused on issues such as estimation and modeling methods, this should not be taken to mean that they do not appreciate the need for credible sources of identification. In the industrial organization seminars and conferences we attend, discussions of identification and its credibility play a

---

[1] By a structural model we do not mean the econometric textbook definition (Greene, 2003, Chapter 15), but rather an economic behavioral model that defines the relationship between exogenous and endogenous variables, both observed and unobserved by the researcher. For further discussion, see Heckman (2000).

central role in the presentations and discussions, regardless of whether the paper's approach is "structural" or not.

However, empirical analysis must deal not only with credible inference, but also with what might be called "generalization," "extrapolation," or "external validity" (to use the terminology of Deaton, 2009). This is where structural analysis comes in. Structural analysis is not a substitute for credible inference. Quite to the contrary, in general, structural analysis and credible identification are complements.

When sources of credible identification are available, structural modeling can provide a way to extrapolate observed responses to environmental changes to predict responses to *other not-yet-observed changes*. In an ideal research environment, this would be unnecessary. Whenever we would be called upon to predict the effect of a proposed policy or anticipated change in the economic environment, there would be many prior events where the same change happened exogeneously (whether through actual randomized trials or naturally occurring ones) and we could use these to estimate a treatment effect. That is, as Angrist and Pischke note, past evidence would then be rich enough to provide a "general picture." Unfortunately, the real world is not always so ideal. The change we are interested in may literally never have occurred before, and even if it has, it may have been in different circumstances, so the previously observed effects may not provide a good prediction of the current one. Structural analysis gives us a way to relate observations of responses to changes in the past to predict the responses to *different* changes in the future.

It does so in two basic steps: First, it matches observed past behavior with a theoretical model to recover fundamental parameters such as preferences and technology. Then, the theoretical model is used to predict the responses to possible environmental changes, including those that have never happened before, under the assumption that the parameters are unchanged.

Another closely related use of structural modeling is to conduct welfare calculations. In some cases, for example, we might be able to predict price changes due to a proposed policy, but without an economic model we could not compute the welfare implications of these changes. We may want to know the overall effect on consumers if some prices go up and others down, or we may want to compare effects on consumers with the changes in firms' profits. When changes in consumers' well-being and firms' true economic profits are unobserved, estimation of treatment effects is not possible, but inferences about underlying preference or cost parameters drawn from observed behavior can allow us to predict these welfare changes. In fact, this use of structural models can again be seen as an example of extrapolation. If we could see previous examples of consumers and firms choosing between the "before" and "after" outcomes we would not need a model. Rather, we could infer welfare changes based on which outcome they chose. But this is usually impossible, so instead a model is used to extrapolate from observation of *other* choices by consumers and firms to predict whether they would prefer the before or the after outcome.

To illustrate these points, we focus on the example highlighted by Angrist and Pischke from industrial organization: the analysis of mergers.

The Analysis of Mergers: Who Can You Trust When It Comes to Antitrust?

As an example of a field that does not fit their mold, Angrist and Pischke offer industrial organization. In particular they discuss merger analysis and conclude that industrial organization has got it wrong. The merger example is a good one, but it demonstrates not the "disorganization" of industrial organization, but rather the limitations of Angrist and Pischke's approach.

Angrist and Pischke contrast two possible approaches to merger analysis: one that they describe as the "transparent analysis of past experience" (that is, quasi-experimental analysis of treatment effects) and the other as the "complex, simulation-based estimates coming out of the new empirical industrial organization paradigm." To them, it is hard to see why one might favor the latter over the former.

Consider the problem faced by an antitrust agency or a court confronted with a proposed merger between two firms and charged with protecting consumer welfare.[2] Should the merger be allowed? For simplicity let's assume that the firms both produce substitute varieties of the same differentiated product (that is, this is a "horizontal" merger). Economic theory gives the basic tradeoffs. The merger will cause the two firms to make pricing decisions jointly, internalizing the effect of their price choices on each other's profits. This increase in "market power" will tend to raise prices, although the precise amount depends on factors such as demand substitution between the products of the two firms, the diversion of consumers to rivals caused by a price increase, and the structure of costs. On the other hand, the merger might result in some reductions in marginal cost that would offset the incentive to increase prices. Indeed, with large enough efficiency gains, prices might decrease as a result of the merger.

The key question facing the antitrust agency or court is which of these two effects dominates. If prices go up, consumers will be harmed and the merger should be blocked; if they go down, consumers will be better off and the merger should be allowed.[3] (When some prices rise and some fall, the overall impact on consumer welfare would need to be assessed.)

---

[2] This is basically the situation in the United States and many other countries. If the goal of the agency or the court is different (for example, to maximize total welfare) then the details of the discussion that follows would differ, but our basic points about the Angrist and Pischke thesis would be unchanged.

[3] In fact, this prescription ignores a number of potentially complicating factors in the determination of an optimal merger policy. For example, if market conditions may change in the future, a merger that is good for consumers today might become bad for them in the future, and vice versa. Moreover, today's decision might alter the set of future merger proposals, introducing another effect on consumer welfare (Nocke and Whinston, 2008). In addition, even absent these dynamic effects, it may be optimal to block certain types of mergers to encourage firms to propose other ones that would be more beneficial for consumers (Lyons, 2002; Armstrong and Vickers, forthcoming; Nocke and Whinston, 2010). We simplify here to remain focused on the issues that Angrist and Pischke raise.

Extrapolation of Merger Treatment Effects

How do Angrist and Pischke propose to address this tradeoff? They propose to look at outcomes in past mergers. Of course, simply looking at the average effect of *all* previously consummated mergers is unlikely to provide a very useful prediction. Angrist and Pischke never provide details, but apparently what they have in mind when they suggest the use of "direct" evidence is some sort of predictive model that averages over the outcomes in "similar" past mergers to predict the effects of a current merger.

There are several problems with this approach. The most important problem, in our view, is how to define "similar" mergers. Clearly, we would not want to predict the effects of a merger, say, in the retail gasoline industry based on what happened after, for example, a merger in the cereal industry. But it is also unclear whether we would want to use a past merger in the gasoline industry to predict the effects of a current proposed merger in the same industry. The circumstances of the industry could have changed or the characteristics of the merging firms may differ from those in the previous merger, and therefore the previous merger might not provide a good prediction of what will happen.

As an example of a way to "trace a shorter route from facts to findings," Angrist and Pischke offer the analysis by Hastings (2004). Hastings analyzes the price effects of the acquisition of Thrifty, a California gasoline retail chain selling unbranded gasoline, by ARCO, a national branded and vertically integrated gasoline chain. After the merger, ARCO re-branded the Thrifty stations with the ARCO name and colors. Hastings studies how rivals' prices changed as a result of the merger. To do so, she compares the differences in price change, before and after the merger, between gas stations that were near a Thrifty station (the treatment group) and those that were not (the control group). The circumstances of the acquisition provide a reasonable basis to think that the merger can be considered as exogenous to the local market; that is, it seems unlikely to be correlated with any unobserved factors that would have changed prices in markets containing Thrifty stations differently from prices in markets without them. She finds that gas stations that were near a Thrifty station raised their prices after the merger more than those that were not, indicating that the merger caused prices to increase.

Hastings's (2004) analysis is based on clever and careful design and sheds light on an interesting question in an important industry. But does it allow us to predict the effect of other possible mergers in this industry? What if two of the largest branded firms in this market wanted to merge? Or if ARCO wanted to acquire a small but branded gasoline retailer in this market (such as Citgo)? Or if ARCO proposed doing this merger without rebranding the Thrifty stations? What if a merger was proposed with convincing evidence of greater cost efficiencies than the ARCO/Thrifty merger? What about a merger in a different part of the country? And what if the acquiring firm in the merger was not vertically integrated?

Of course, if we had previous experiences with all possible types of mergers (and could distinguish them), we could answer all these questions by looking at past outcomes. But given the many possible circumstances of a merger, it seems

inevitable that many possible proposed mergers will not have been seen and studied before. In that case, to use past mergers to predict future outcomes, one needs a model. This model can be a statistical model or it can be an economic model. A statistical model, Angrist and Pischke's preferred approach, would seek to predict the outcome of a merger using either a group of not-too-dissimilar mergers (perhaps all mergers in similarly concentrated industries resulting in similar increases in concentration), or more generally fitting some prediction function based on a set of observable merger attributes. The ARCO/Thrifty example makes clear that this will often be a difficult task to do in a convincing manner, even when some mergers have previously been observed in an industry.

There are other concerns with this approach, beyond the extrapolation issue. One is the difficulty of defining a reasonable benchmark by which to judge the outcomes of mergers.[4] A naive approach would compare outcomes of the impacted firms—the merging parties and their competitors—to unaffected firms. But it is not obvious how to find firms that are good comparisons yet at the same time are not affected by the merger. In Hastings (2004), for example, the use of the control group relies on the assumption that stations further than one mile from a Thrifty station will be unaffected by the merger. If consumers search for stations beyond this distance, this assumption could fail, most likely leading to an underestimate of the merger's effect. Fortunately, Hastings does examine the use of different distances and finds no change in her results, and also documents that the control and treatment group prices moved in parallel prior to the merger.

Finding such a control group is likely to be harder, however, in many other industries. For example, Angrist and Pischke offer Ashenfelter and Hosken (2008) as another example of direct evidence of mergers' effects. Ashenfelter and Hosken examine the price effects of five national branded consumer product mergers and use private-label products as a control group for the products of the merging firms. However, retail prices of private-label products can be affected by a merger of branded manufacturers if marginal costs are not constant, if private-label producers are not perfectly competitive, or if retailers adjust retail margins of private-label products in response to wholesale price changes.

A second difficulty is that the treatment effect approach requires that the mergers effectively be exogenous events. But mergers are an endogenous choice of firms that may be motivated, in part, by past, current, or anticipated future changes in unobservable (to the researcher) market conditions. While we find Hastings' (2004) argument for exogeneity reasonably convincing, we are more troubled by Ashenfelter and Hosken's (2008) exogeneity assumption, which they adopt with little discussion or justification. For example, one of the acquisitions they study is the purchase of the Chex brand by General Mills. Ralston, which sold Chex to General Mills, produces many private-label products and according to reports in

---

[4] Absent such a benchmark, it would be necessary to include as explanatory variables all of the factors that would explain prices absent a merger and which are correlated with the occurrence of the merger (Ashenfelter, Hosken, and Weinberg, 2009).

the press was selling Chex to focus on its private-label business. Therefore, it seems likely that this event could be related to unobserved changes in the demand for private-label products.

Add to these concerns the fact that the merger treatment effect approach cannot produce measures of welfare change and it becomes clear that it is far from the simple solution to predicting merger effects that Angrist and Pischke make it out to be.

Extrapolation Using an Economic Model

An alternative approach to predicting a merger's effect instead consists of using economic theory to simulate what the effect of the merger is likely to be.

The basic idea is simple. Historical data are used to recover the structure of an economic model that consists of demand, supply, and competition. Identification of the fundamental parameters of this structure follows instrumental variable procedures similar to those in classical demand and supply estimation. Using the model, one can then simulate the effect of the merger under a variety of assumptions. The assumptions can include different models of post-merger competition and changes in marginal cost. For example, one could ask what level of cost efficiencies are needed, under the model, to assure that prices will not increase. This can lead to some assessment of the likelihood that these efficiencies will be realized. The typical exercise does not offer a single number, as Angrist and Pischke suggest, but rather a range of numbers under different assumptions.

The data used to estimate the model also do not need to consist of past mergers (although they could when mergers have occurred that can be considered exogenous), which can be very helpful in industries where there have been no past mergers. Moreover, because of this feature, a researcher is more able to use careful design and credible inference to shed light on the likely effect of the merger.

In addition, the method makes calculation of welfare effects straightforward.

Just as before, a model is used to extrapolate from the past to infer the effect of the merger. But while before it was a statistical model, now it is an economic model. So we repeat the question that Angrist and Pischke ask: Who should we trust when it comes to antitrust? A model grounded in economic theory, estimated using careful design? Or a statistical model that is based on a few observations of previous, quite different mergers, where exogeneity may be questionable?

Comparing the Two Approaches

To highlight the flaws in Angrist and Pischke's argument, we have so far highlighted the problems with the treatment effect approach to predicting merger effects and deliberately overemphasized the benefits of structural simulation analysis. While one-sided comments may make good controversy, they probably don't make for good economics.

We do in fact believe that the treatment effect approach will sometimes prove useful for predicting a merger's effects. Even if it is unlikely that we will be able to obtain credible evidence on a wide range of merger treatment effects given the

many possible circumstances of mergers, it may prove fruitful to focus efforts on examining the effects of certain types of mergers. For example, Ashenfelter and Hosken (2008) suggest focusing on mergers that are on the margin of current enforcement practice, where evidence is likely to be most useful. Particular industries with extensive merger histories and credible inference possibilities might also be targeted.

We also believe that merger simulation has limitations. First, while in principle the first step, demand estimation, can incorporate credible inference, in practice a typical exercise may rely on less-than-ideal instrumental variables. For example, following Hausman (1997), Nevo (2000) uses prices in other markets as instruments for price when estimating demand. Angrist and Pischke refer to the assumptions that justify these instruments as "arbitrary." While we are somewhat more positive about the validity of these instruments, we are sympathetic to the concerns.[5] These instruments are not the only ones used, or even the most popular, and the validity of the assumptions justifying the instruments will vary on a case-by-case basis. In general, we think it is fair to say that in many cases the instruments are less than ideal. In our view, rather than invalidating the entire approach, this concern merely highlights the importance of ongoing work that explores additional instruments and different inference methods. For example, Nevo and Rosen (2009) study the above instrumental variables and propose a way to (set) identify the parameters of interest even if the standard orthogonality conditions fail.

Second, one needs a good model of a variable's determination to predict accurately how a merger will change it. Thus, current merger simulations focus mostly on predicting price changes holding the current set of products, firms, and pricing behavior fixed. Effects on prices due to changes in long-run investments, research and development, and entry are typically ignored at present, as are effects on available product offerings. (This is one advantage of the treatment effect approach, where it is feasible, since in principle it can capture some of these additional effects of a merger.) In addition, merger simulation relies on assumptions about how the merger will change behavior, often based on static Nash equilibrium before and after the merger. Richer models of how behavior changes (for example, models of collusion) have seen little use. These limitations are potentially serious, although this is an active area of research and we expect economist's abilities on both fronts to improve over time.

Another concern often raised with merger simulation and structural work more generally is of an "elaborate superstructure," to use the words of Angrist and Pischke. There is a feeling that results are driven by nontransparent complicated models and not by data per se. This is a concern to be taken seriously, because estimates driven by functional form rather than credible sources of identification

---

[5] Nevo (2001) provides a discussion of these instruments, including cases where they might fail, and shows that in a more limited model they yield results almost identical to those obtained from using cost variation as the exogenous source of variation. In his full model, the cost variation cannot be used as the sole source of exogenous variation due to a dimensionality problem: there are too many parameters to estimate.

in the data are unlikely to produce useful predictions. Yet, while sometimes this might be the case, often the so-called "complicated models" are introduced exactly to relax the reliance on functional forms. For example, the "complicated" demand model of Berry, Levinsohn, and Pakes (1995) relaxes some of the strong implications of the much simpler multinomial logit model. More recent work that explores nonparametric identification and estimation of this model (for example, Berry and Haile, 2009) even further relaxes some of the imposed structure. We therefore believe that these concerns are due at least in part to a lack of familiarity and comfort with the models used in industrial organization.

In sum, *both* merger simulation and the merger treatment effect approach seem likely to be useful in some cases and fail in others. Depending on the question being addressed, and the availability of data, one approach might dominate.

Other Uses of Retrospective Estimates of Merger Treatment Effects

While using estimates of merger treatment effects to predict the effect of a given merger has some serious limitations, the estimates can be very useful for addressing other questions. Indeed, we have been encouraging retrospective merger studies for a while (for example, Nevo, 2000; Whinston, 2007a, b).

First, whatever methods are used for predicting the effects of mergers before they occur, retrospective studies of merger effects can be useful for judging the accuracy of those methods. This use of retrospective studies is fairly recent.[6] For example, Peters (2006) examines structural merger simulation methods as applied to a set of airline mergers in the 1980s. He finds that the merger simulations fail to predict accurately the magnitude of price changes in several of the mergers. Peters also explores the sources of the errors in his merger simulations (for example, post-merger changes in product offerings, shifts in demand, or changes in behavior). Of course, the perhaps more relevant issue is how the simulation method does *compared to other possibilities*, such as prediction based on treatment effects computed from *other* past mergers. Indeed, we can imagine future studies comparing structural merger simulation methods to the treatment effect approach championed by Angrist and Pischke, as well as other methods. (Peters, for example, compares the structural merger simulation predictions to the predictions from a reduced-form regression with industry concentration as the independent variable and price as the dependent variable.)

Second, the problem of optimal legal (and regulatory) review has one important feature we have not mentioned: the costliness of the proceedings makes it optimal to have screens based on limited evidence. For example, there may be "safe harbors" granting approval to certain mergers without a full review. (This is in effect what happens when the U.S. antitrust agencies decide not to issue a "second request" for additional information about mergers that are reported under the Hart–Scott–Rodino merger filing law.) For this purpose, knowing the average

---

[6] See the discussion in Nevo (2000, page 416) in the context of mergers, or Hausman and Leonard (2002), who evaluate the ability of structural models to accurately predict the gains from new goods.

effect in a wide class of mergers (for example, those in industries with concentration below some level) would be useful information.[7] However, determining this average effect is nontrivial: in particular, it will not equal the average treatment effect of *approved* mergers because the sample of approved mergers is a selected sample, where the selection is based on additional information that was at the agency or court's disposal.

How and Why Are Industrial Organization and Labor Different?

Empirical work in industrial organization does differ in some striking ways from that in labor (and other fields that emphasize estimation of treatment effects). We have discussed extensively one important difference—the heavier reliance on structural modeling (and greater attention to issues this raises) in industrial organization—but this is not the only difference.

Empirical papers in industrial organization are also less likely than are papers in labor to focus on pinning down a particular "number"—like an elasticity or a price effect. Many structural papers in industrial organization, for example, are focused on showing that an approach to answering a question is feasible. And even nonstructural "reduced form" papers whose methods resemble the treatment effect approach often focus on testing a general prediction of a class of theoretical models rather than producing an estimate of a treatment effect. For example, Borenstein and Shepard (1996) study cyclical pricing in the gasoline market, using what is clearly not a structural approach, yet their focus is on providing evidence in support of collusive pricing and not recovering a particular number. Indeed, even Hastings (2004) seems to focus as much or more on the sign of the price effects arising from the ARCO/Thrifty merger and what they imply about the structure of retail gasoline competition than on the exact magnitude of those effects.

An interesting question is why these differences across fields exist. Several possible explanations suggest themselves. As our discussion of merger analysis illustrates, industrial organization economists seem far more concerned than labor economists that environmental changes are heterogeneous, so that useful estimates of average treatment effects in similar situations are not likely to be available. We are unsure whether the typical merger is more distinctive than is the typical labor market or education policy intervention, but Angrist and Pischke's discussion of class size studies suggests that this may be the case. In addition, Angrist and Pischke's discussion also suggests that the data available to labor economists may be more likely than that in industrial organization to contain many examples of similar changes, as well as a richer set of directly observable controls. To the extent that either of these differences is present, it creates good reason for industrial

---

[7] Actually, knowing the full distribution of effects, as well as how those distributions would be narrowed with more information, would also be helpful.

organization economists to rely more explicitly on theory to predict responses than labor economists do.

Another factor may relate to differences between the data available to researchers and the data available to policymakers in the two fields. For example, when an antitrust agency examines a merger, it is likely to have much more information than would the typical researcher studying the same issue. In contrast, a policymaker approaching a labor question probably has no more information than does an outside researcher. As a result, it may be most useful for industrial organization economists to identify techniques for policymakers to use, while labor economists are most useful when they estimate effects, pinning down numbers such as the elasticity of labor supply or the effect of smaller class sizes.

Still another difference may be related to the nature of the models used in the different fields. In general, the models used by industrial organization economists tend to be more complicated than those used by labor economists. Consider, for example, demand analysis. A labor economist might study how technical change (perhaps the advent of computers) affects the demand for skilled and unskilled labor. Doing so involves a fairly simple demand system. In contrast, an industrial organization economist looking at how a change in the price of gasoline affects consumer demand for cars would often be concerned with estimating reasonable elasticities for many different car models. This leads to a much more complicated model and estimation problem, as in Berry, Levinsohn, and Pakes (1995). Moreover, in many of the problems studied by industrial organization economists, strategic interaction between agents is of first-order importance, requiring tools beyond simple supply and demand analysis. There are, of course, subfields of labor where more complicated theoretical models arise, such as in studying search, but these represent a minority of current work in labor.

That said, we suspect that some of the differences in the styles of empirical work may be due more to cultural differences than to the actual economic problems, suggesting that the differences are greater than they should be. For example, in the demand estimation problems discussed in the previous paragraph: Should labor economists distinguish among many different types of skilled and unskilled labor? Should industrial organization economists use simpler, more aggregated demand structures for cars? It probably depends on the question, but cultural differences may now be driving these choices to some degree.

Indeed, a typical scholar of industrial organization is exposed to theory earlier and more often in his or her career than is the typical labor economist, and is therefore more likely to want, and be able, to relate to economic theory in empirical work. The industrial organization researcher may also be more concerned about the exact circumstances surrounding a policy intervention or exogenous event, having been trained to think they are likely to be important (recall the merger versus class size discussion). The exposure to theory could also be driving a desire to not just measure an effect but to understand the mechanism at work—even if there is no policy relevance.

As this point and the discussion above suggest, we suspect that researchers in industrial organization and those in fields where treatment effect methods are dominant would both do well to ask themselves where adoption of each others' approaches could prove useful, while respecting the fact that differences in the markets, data, and questions considered in different fields will call for differing approaches.

Belaboring the Obvious

Our view is that the future of econometrics and applied microeconomic work is in combining careful design, credible inference, robust estimation methods, and thoughtful modeling. Therefore, any serious empirical researcher should build a toolkit consisting of different methods, to be used according to the specifics of the question being studied and the available data. That this should not be an either–or proposition seems quite obvious to us.

References

Armstrong, Mark, and John Vickers. Forthcoming. "A Model of Delegated Project Choice." *Econometrica.*

Ashenfelter, Orley, and Daniel Hosken. 2008. "The Effect of Mergers on Consumer Prices: Evidence from Five Selected Case Studies." NBER Working Paper 13859.

Ashenfelter, Orley, Daniel Hosken, and Matthew Weinberg. 2009. "Generating Evidence to Guide Merger Enforcement?" NBER Working Paper 14798.

Berry, Steven, and Phillip Haile. 2009. "Nonparametric Identification of Multinomial Choice Demand Models with Heterogeneous Consumers." Cowles Foundation Discussion Paper 1718.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica,* 63(4): 841–90.

Borenstein, Severin, and Andrea Shepard. 1996. "Dynamic Pricing in Retail Gasoline Markets." *Rand Journal of Economics,* 27(3): 429–51.

Deaton, Angus. 2009. "Instruments of Development: Randomization in the Tropics, and the Search for the Elusive Keys to Economic Development." NBER Working Paper 14690.

Greene, William H. 2003. *Econometric Analysis,* 5th ed. New Jersey: Prentice Hall.

Hastings, Justine S. 2004. "Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California." *The American Economic Review,* 94(1): 317–28.

Hausman, Jerry A. 1997. "Valuation of New Goods under Perfect and Imperfect Competition." In *The Economics of New Goods,* ed. Timothy F. Bresnahan and Robert J. Gordon, 209–248. Chicago: National Bureau of Economic Research.

Hausman, Jerry A., and Gregory K. Leonard. 2002. "The Competitive Effects of a New Product Introduction: A Case Study." *The Journal of Industrial Economics,* 50(3): 237–63.

Heckman, James J. 2000. "Causal Parameters and Policy Analysis in Economics: A Twentieth Century Retrospective." *The Quarterly Journal of Economics*, 115(1): 45–97.

Leamer, Edward. 1983. "Let's Take the Con Out of Econometrics." *The American Economic Review*, 73(1): 31–43.

Lyons, Bruce. 2002. "Could Politicians Be More Right than Economists? A Theory of Merger Standards." Unpublished paper, University of East Anglia.

Manski, Charles, F. 2003. *Partial Identification of Probability Distributions.* New York: Springer.

Marschak, Jacob. 1953. "Economic Measurements for Policy and Prediction." In: *Studies in Econometric Method,* eds. W. C. Hood and T. C. Koopmans, pp. 1–26. New York: Wiley

Nevo, Aviv. 2000. "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry." *The RAND Journal of Economics*, 31(3): 395–442**.**

Nevo, Aviv. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica,* 69(2): 307–342.

Nevo, Aviv, and Adam Rosen. 2009. "Identification with Imperfect Instruments." NBER Working Paper 14434.

Nocke, Volker, and Michael D. Whinston. 2008. "Dynamic Merger Review." NBER Working Paper 14526.

Nocke, Volker, and Michael D. Whinston. 2010. "Merger Policy with Merger Choice." Unpublished paper.

Peters, Craig. 2006. "Evaluating the Performance of Merger Simulation: Evidence from the U.S. Airline Industry." *The Journal of Law and Economics,* 49(2): 627–49.

Powell, James L. 1994. "Estimation of Semiparametric Models." Chap. 41 in *Handbook of Econometrics,* vol. 4, ed. Robert Engle and Daniel McFadden. Amsterdam: North Holland.

Tamer, Elie. Forthcoming. "Partial Identification in Econometrics." *Annual Reviews of Economics,* vol. 2.

Whinston, Michael D. 2007a. "Antitrust Policy towards Horizontal Mergers." In *Handbook of Industrial Organization,* vol. 3, ed. Mark Armstrong and Robert Porter, 2369–2440. Amsterdam: Elsevier.

Whinston, Michael D. 2007b. *Lectures on Antitrust Economics.* Cambridge, MA: MIT Press.

White, Halbert L. 1980. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." *Econometrica*, 48(4): 817–38.

This article has been cited by:

1. Judea Pearl. 2011. Statistics and Causality: Separated to Reunite-Commentary on Bryan Dowd's "Separated at Birth". *Health Services Research* **46**:2, 421-429. [CrossRef]

2. Henk Folmer, Olof Johansson-Stenman. 2011. Does Environmental Economics Produce Aeroplanes Without Engines? On the Need for an Environmental Social Science. *Environmental and Resource Economics* **48**:3, 337-361. [CrossRef]

3. James J. Heckman. 2010. Building Bridges between Structural and Program Evaluation Approaches to Evaluating PolicyBuilding Bridges between Structural and Program Evaluation Approaches to Evaluating Policy. *Journal of Economic Literature* **48**:2, 356-398. [Abstract] [View PDF article] [PDF with links]