Empirical methods in Industrial Organization- Science or Art?

Michael Waterson, University of Warwick

September 2013

Abstract

This paper considers some key developments in empirical work in industrial organization, contrasting alternative approaches and their strengths and weaknesses in discovering answers to important questions about behaviour of firms. Over time, there have been clear empirical advances, but perhaps paradoxically, these have not necessarily been in the direction of increasing the science of the approach, but sometimes instead the art of finding good sources of variation in the data. One danger lies in valuing this art above the value of the question being examined; here contrasts are drawn. Some of these empirical issues are then examined in the context of my own work, covering both approaches.

Introduction

Since the early empirical days of Industrial Organization, in the 1950s to 1970s, the analysis has become significantly more sophisticated. This is not just because of improved calculating ability. Essentially, early empirical work can be characterised as focusing on establishing statistical associations between variables where causation was asserted by fiat: The predominant empirical approach was the "structure-(conduct)- performance" paradigm, whereby structural features were asserted to cause performance moves, structure being viewed as more immutable than performance. Although there may be an element of truth in this characterisation, it is not a very scientific approach! Moreover, the policy interpretation put on such studies actually required the answer to a different question than that normally examined empirically.

Since then, researchers have developed more focused analyses, generally of narrower but better formulated questions. They have treated potential endogeneity seriously and have used more sophisticated and nuanced methods better designed to analyse the question of interest. So, more recent studies are clearly more scientifically based. Yet at the same time, they contain elements which are more art than science, and they are often controversial. Essentially, I want to illustrate these points about the more recent developments through analysis of a debate about methods, together with some empirical examples of my own (which I choose not because I view them as being better than any alternatives, rather because they are examples with which I am most familiar!).

The early days

Early industrial organization economists, up until the late 1970s, were keen on examining big questions. Essentially, they were problem-oriented. The older ones amongst them may well have been powerfully influenced by experiences in the Depression era of the 1930s. The potential problem they saw was the role of big business, a problem that still attracts some attention today, of course. Their concern was to examine the impact of big business on pricing and firm performance more generally. As late as 1980, it was possible to ask a really broad question like "How serious are the performance deficiencies resulting from monopolistic structure and conduct in the United States?" (Scherer, 1980, p. 459) and anticipate generating a serious answer. In some sense, this is more exciting than asking "Is a merger between firms X and Y in a particular industry likely to raise prices to a greater extent than it reduces costs, so leading to a net decrease in overall welfare?" But it is much less easy to get even an approximate answer to the first question in a way that is scientifically justifiable.

Let me formalise what I mean somewhat. The big question can be characterised as: Consider the estimate of *b* in the equation below:

$$Y_i = a + bX_i + u_i$$

Is it statistically significant, and is it large in magnitude? Here the variable *X* is viewed as the thing determining *Y*. Therefore, a key issue is whether *X* is in fact exogenous or is itself influenced in turn by *Y*. A second issue is whether there are other missing variables, themselves correlated with *X*, that should in fact be included in the regression. The fact they are not will bias the estimate of *b*. A third issue relates to the source of variation across which the relationship is estimated. This is linked

with the second issue, because the broader the data, for example if it is cross-industry, the more likely there are to be other things for which we should control. A fourth issue is whether what we actually need in order to get good estimates.

The more recent empirical studies have had more modest ambitions. We now accept that "Individual industries are taken to have important idiosyncrasies" (Bresnahan, 1989, p. 1012), a clear answer that the second issue is very important.[1] Again, starting in the 1980s, the idea came that the question: "what determines the number of firms?" (Stiglitz, 1986, p. ix) was important; "The objective of this line of research is to identify *exogenous* variables … [which] determine the market structure, which is thus viewed as *endogenous*." Hence, at least one additional equation is required (in the medium run) concerned with explaining *X*.

More recent developments

Once questions move to a smaller canvas, fewer things require to be controlled for, so the second issue changes in nature. The obvious source of variation seemed, to the early investigators, to be cross-industry. This was superficially convenient, because publicly available data coming from Censuses of Production and the like could be employed in analysis. The statisticians' definitions of industries might be somewhat arbitrary, the variables measuring in particular performance might be crude, but analysis could be done with reasonable dispatch. The census is not a particularly good tool for analysis within the industry (though see for example, Clarke et al, 1983). This meant that, increasingly, authors had to do the hard work of digging out, constructing and combining their own databases before starting the econometric analysis. But these may lack information on some key elements.

Reality is complex, and undoubtedly the econometrician knows less about reality than they would like. Consider the market for cars. A car is a significant purchase, and there are undoubtedly significant factors, idiosyncratic to the individual or the household, that influence their choice. Amongst these are what they may seek to use the vehicle for (something that might be captured in part by consumer demographical information). They may also be influenced by hard to measure factors like style- if not then the sellers of stylish but notoriously less reliable vehicles would have a difficult time. More complex still is the supply side: the production technology and organisational choices made by car producers are difficult to model in any detail.

A structural approach

It is here that different research strategies become apparent. One way of tackling a complex reality is to put some structure into the problem. Let us consider a common issue, analysis of a merger between two firms in the same industry. Most products are differentiated and, as already said, some characteristics of those products that will undoubtedly influence demand have to go un-modelled. The predominant approach here (seen most obviously in Berry et al, 2005) tackles this issue head on by allowing for unmeasured factors influencing purchases, but doing so in such a way that this drops out as the error term. Translating from demand to market share, the demand for

[1] At this point, I should note that there is another approach entirely, although not one that has found broad favour, for whatever reason. The approach taken by Sutton (1991) is first to look for things that can be maintained quite broadly, for example that in larger markets, there are likely to be more firms. This then forms a firm foundation for the next stage of analysis.

product j can be written in terms of its share s_i relative to that of the outside good, s_o, along the following lines:

$$\ln(s_i) - \ln(s_o) = x_j\beta - p_j\alpha + \sigma \ln(\bar{s}_{j/g}) + \xi_j$$

where the x is a vector of characteristics of the product j, p_j is its price, the parameters α and β are allowed to vary with consumer characteristics, and the term that has the parameter σ relates to the share of the group of products in which j is an element. The final term is the error relating to unobserved characteristics.

So much for the demand side. Then on the supply side, something along the lines of a supply relationship coming from first-order maximisation conditions

$$p_j = c_j(w_j) - q_j P'(Q).\theta_j$$

is common. Here marginal costs c are a function of factor prices w and assuming some degree of imperfect competition, supply will relate to market demand elasticity through the second element representing the markup. In this part, θ is the conduct parameter.

Two of the most controversial aspects of this are (i) instruments for prices in the demand function and (ii) determination of the conduct parameter. Within the products produced, the demand for one product will be influenced by the prices of others, so prices are not exogenous, but good instruments are seldom available. Recall that a good instrument is one that is uncorrelated with the error, but correlated (preferably well) with the variable for which it is instrumenting. A common ploy is to argue for the use of prices in other locations- highly correlated with prices in this location but unrelated to the error term. This makes the assumption that markets in different places are independent, which may well not be true. It is here where some art creeps into the science, in terms of the arguments as to why, in a particular context, certain instruments are indeed valid in terms of being uncorrelated with the error.[2] Fundamentally, this aspect is problematic to test, because the error term by definition is unobserved. It is impossible to test if the equation in question is exactly rather than over-identified.

The conduct parameter is also something of a nuisance. If true marginal costs were observable, then the possibility exists that conduct may be estimated (Genesove and Mullin, 1968). But in general, for complex technologies, this cannot be done, so that the estimates in effect become joint estimates of marginal cost and conduct. But for example the conduct parameter may be influenced by the state of demand, so that a given outcome may be interpretable, say either as a large markup on marginal cost, or by a high marginal cost together with a low markup (Corts, 1999). Identification may depend subtly on the assumptions regarding the form that a shift in demand may take (Bresnahan, 1982). Clearly, relying for identification of the supply side on a particular assumed structure of demand is controversial. On some occasions, the assumptions used are clearly identified and justified (Nevo, 2000), but nevertheless, Angrist and Pischke (2010) are harshly critical, entitling the section of their paper dealing with the structural approach in IO, Industrial Disorganization!

[2] Suppose the company in question engages in a nationwide marketing campaign. Then even if the markets concerned are well separated, the prices across markets may be subject to common influences, so that the assumption of common errors cannot be ruled out.

So, if the structural econometrician has not been artfully able to convince readers that the assumptions underlying their estimates are valid, what then?

The (quasi-) experimental approach

The ideal clinical trial for a new drug involves a randomly-selected group of individuals being matched with another randomly selected group, alike in say having a particular condition, but with one group being subject to treatment, another group subjected to a placebo treatment (or indeed, no treatment- the outcomes may not be identical); they should not know to which group they are assigned. Ideally also, the researchers assessing the effect of the treatment should not know the assignment, so the trial is "double-blind".

The aim of the quasi-experimental researcher in economics may be to emulate this ideal. Those who engage in laboratory experiments can carefully control the environment and may be able to do so, although one clear moot point is whether their subjects are in some sense "typical", since they are commonly recruited from resident student populations.

Outside the laboratory, a clever research design can emulate an experiment. The aim is to capture the essence of an experimental approach. Alternatively and less commonly, and here the persuasive art of a researcher comes in, they may be able to persuade a firm, say, to adopt some randomisation strategy in order to discover the impact, or indeed the firm may itself wish to discover something about the nature of demand and might be persuaded to give access to the data thereby generated to researchers.

In principle, this is an ideal approach in most respects. Rather than trying to model complex effects, all possible influences except one are maintained constant in the (quasi)- experiment and the effect of the experiment can then be evaluated. To be sure, it is only the net effect that is found- for example if a price is changed for some consumers not others, one might discover the elasticity of demand but not be able to unpick income and substitution effects of the price change.[3]

Unfortunately, there are many pitfalls. To illustrate, I take the example of an essentially failed experiment in which I was involved, together with a PhD student. The experiment was designed to evaluate the use of "smart meters" on electricity consumption. There were several alternative treatments- different meter types, different degrees of "smartness", different incentive schemes. It was sophisticated, but at the same time fundamentally flawed; the flaw was present before we were involved, but only became apparent some way into the process. The experimenters had contacted consumers to take part in the experiment. However, they had kept no record of those consumers who had refused to take part. Clearly then, the consumers who participated would be likely to be a biased sample of the population. For example, a consumer is more likely to sign up to an experiment where they are incentivised for using less energy than before if they know that household consumption is likely to go down. So we could only make comparisons across

[3] There is an irony here. Some of the most traumatic events, for example the "9/11" attack on New York, have one beneficial side-effect: of acting as natural experiments (Bloom 2009).

differently treated groups, not comparisons with a control group. Hence it was impossible to answer a key question like: Does having a smart meter lead to reduced consumption?

Perhaps this illustrates an unusual example of bad luck. The main problem is that however much art the researcher puts into the project, reality is too messy to make a good experiment. Important questions may be deemed "too important" to experiment with consumers (yet subject to Ministerial interference based on pure prejudice). If you want to know what the impact a particular merger might have, it is only unusually the case that you have a convenient experiment along the lines illustrated in the Staples- Office Depot case, a case well known mainly because of its rarity. Yet even so, this case is not without controversy as a result of potential control issues. (Niskanen, 2012).

Thus alternatively, a researcher may spend a significant amount of time analysing a good experiment, not because it is important, but simply because it is bullet-proof. This is a poor target at which to aim.

Two examples

What I want to do now is to look at two examples with which I am familiar, because I have been heavily involved as an author in both. I do not pretend these are perfect, but they simply illustrate the uses of the two alternative approaches to empirics, both of which have potential flaws.

A natural experiment relating to energy storage

How much energy storage should a country maintain? This question can be asked two ways, the first being how much it should maintain for strategic purposes, for example in case of blockades or politically motivated supply curtailments, the second being what would the amount of energy storage that maximises monetary returns. To a first approximation, a prior policy decision to leave storage to the market implies choosing the latter rather than the former as the appropriate goal. To put it another way, storage is like an insurance policy. Having a large component of storage means (a) you are less likely to need to ration energy in time of war, (b) users of energy are less likely to face significantly higher prices/ shortages in case of sudden accidents. In a market-based system of storage, it is the latter peril that is being insured against. Storage is costly, so that having too much is like over-insuring your property.

Non-economists do not usually couch the argument in these terms. Heavy industrial users of energy see more clearly the potential problems they face when energy runs short than the constant cost of maintaining significant storage. Various techniques of varying degrees of sophistication have been used in attempts to value storage. We took an experimental approach. One way of valuing storage is to see what the impact is when the storage is no longer available. Luckily (in one sense), but context specific of course, just such an experiment took place in the UK in the natural gas industry in 2006.

The context is important here. Many people had claimed that the UK was short of natural gas storage, for historical reasons, and by comparison with other major European economies such as France and Germany. A crude tabulation as in Table 1

shows this to be true. But, how valuable is the storage we have? In the late Winter of 2005/06, we were about to find out. A small accidental fire on 16th February 2006, at first seemingly a minor incident, rendered 80% of the UK's natural gas immediately inoperative until well into the Summer. This might not have mattered, had the Spring been exceptionally mild, but in fact there was an extended cold spell in March, when gas would certainly have been drawn from storage. What happened instead was that LNG was called forth from ships coming from the Gulf and gas was transferred by pipelines to the UK, of course at a price. The impact is immediately apparent from the data, as we see in Figure 1 below. Making allowance for temperature, prices were markedly higher in the period affected by the fire (the dark points in the figure).

Table 1. Natural Gas Storage, end 2009. Selected countries.

Country	Annual consumption (mcm)	Working capacity (mcm)	Peak output (mcm/day)	Implied average days' supply	Max extraction rate
Austria	8802	4639	54.8	192	2.27
Belgium	17188	709	24	15	0.51
France	44507	12395	248.4	102	2.04
Germany	92646	19866	465.1	78	1.83
Italy	78051	14295	271.1	67	1.27
Netherlands	48796	5078	177	38	1.32
Spain	33884	2726	14.5	29	0.16
UK	90759	4310	113	17	0.45

Source: Derived from IEA Statistics: Natural Gas Information 2010

Note: mcm stands for million cubic meters

Taken from Giulietti et al (2012)

Conceptually, our framework (reported in Giulietti et al, 2012) for the gas market is as follows:

$$D_t = D(T_t, P_t); D_T, D_P < 0$$

$$S_t = S(P_t); S' > 0; S \le \bar{S}$$

$$= S_F(P_t); S' > 0; S > \bar{S}$$

$$S_t = D_t$$

where D is demand, S is supply, $\bar{S}$ is the point at which the supply curve bifurcates, S_F is supply in the period of the fire outage, T temperature, t time and P price.

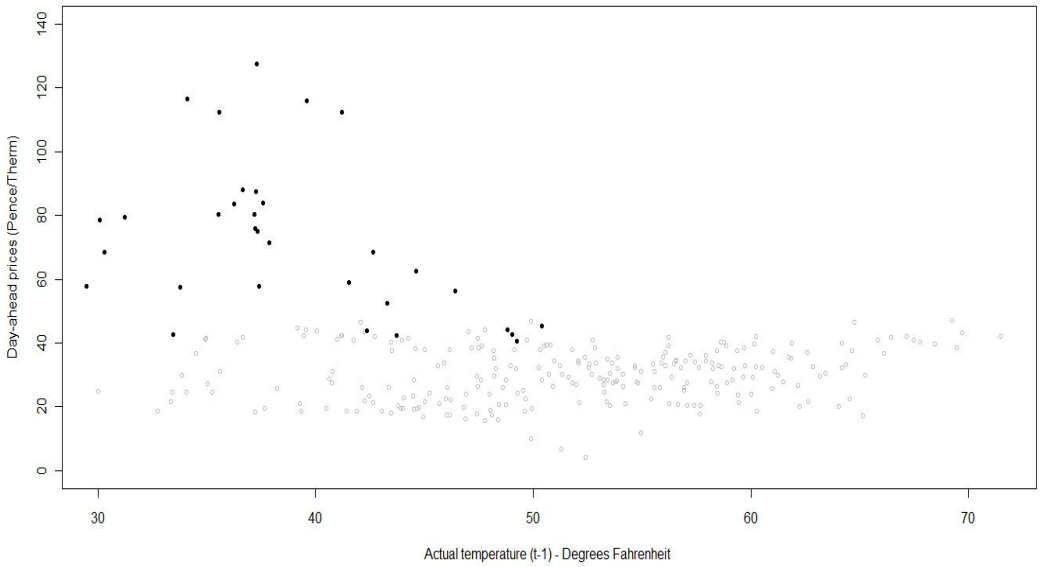

Figure 1: Prices and temperature (at time t-1) from the fire to March 2007

Taken from Giulietti et al (2012)

Our model was as simple as possible, given the problem. The empirical investigation we undertook focused on two points. First, in order to value storage, it was important to establish that the demand side was unaffected by the outage. Given this, we estimated the impact on supply. Supply and demand are of course interlinked, both relate quantity to price.[4] Hence we needed to instrument for the quantity demanded/ supplied in the equation explaining the supply price. Temperature is an almost ideal instrument for demand. The relationship between demand for gas and temperature is extremely close. Temperature does not affect supply (within the range of temperatures experienced in the UK) and is clearly exogenous, that is unrelated to the error term in the supply function. Therefore in the second stage, we estimate the relationship between price and supply quantity with temperature as an instrument.

We find that supply price was higher on average and more sensitive to fluctuations in demand caused by temperature differences during the period of the fire. Prices additionally experienced more variance than in the fire period. What was the impact of the increased prices[5]? To examine this, we took the counterfactual prices and quantities that would have occurred at those temperatures if the fire had not occurred and compared money paid by consumers under the counterfactual with money paid given the fire. The total "excess" payments were of the order of £253.6m.

There is a lingering worry about any such experiment. Is there something else that might have happened at about the same time which could have had an impact? Luckily in this case, the incident is relatively short-lived, its main effects being seen over the first month of cold weather, so that we were able to employ a "moving

[4] There is some flexibility in the relationship between supply quantity and demand quantity in gas, given that pipelines operate on a range of packing pressures, rather than at a single pressure, but to a first order of approximation, quantity demanded equals quantity supplied.

[5] The impact of variance is not fully evaluated in the paper.

window" technique. Taking a window of a month, we ask the question, if we move this window along in the data, when does the change appear to happen? The answer is at almost precisely the time of the fire, as figure 2 shows. The test statistics exhibit a clear jump at exactly the point where the fire occurs. This essentially rules out alternative explanations for what we observe.

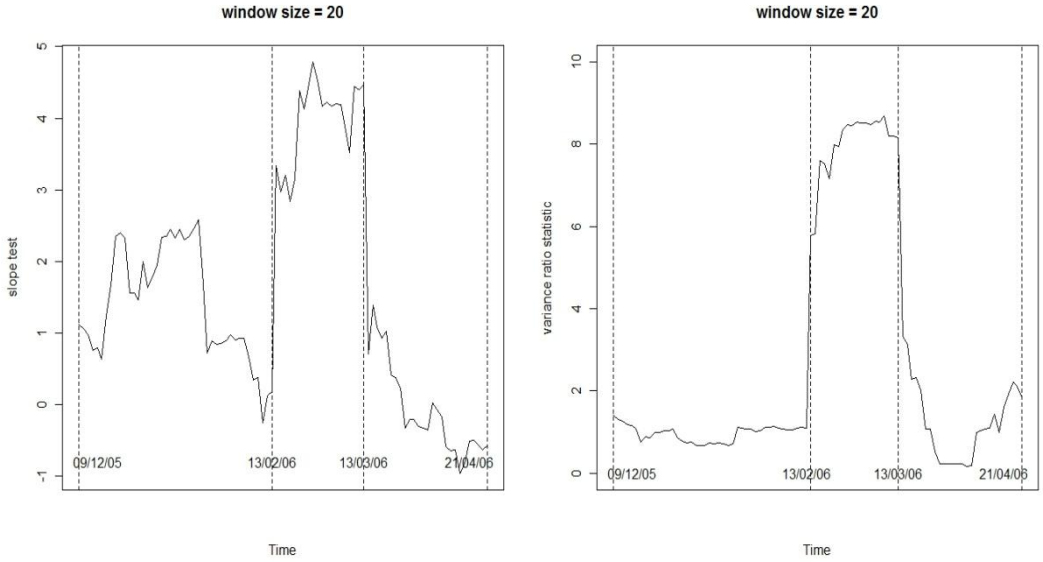

Figure 2: Test for the equality of slopes (left panel) and for the equality of residual variances (right panel). Window length w=20, number of iterations L=90.

Taken from Giulietti et al (2012)

A structural model of consumer search

Consumer search is a difficult topic for the quasi-experimental approach to empirical work, because there are few circumstances under which differential searches are experienced in a "clean" manner. A good example of an attempt to do this that avoids most of the potential pitfalls is Sorensen (2000), who examined pharmacies' pricing behaviour in relation to drugs that were prescribed for long term conditions (where search may be viewed as more important given that the consumer is paying more over the longer term) compared with their behaviour in response to "one off" ailments. Recent advances in technology have also enabled direct observation of search through devices such as RFID tags on supermarket trolleys, linked to cameras capturing movement through the store (a slightly scary thought!). These may well lead to new insights.

We (Giulietti, et al, 2013) examine implied search costs across regions of Great Britain between the start of 2002 and the end of 2005. This is a policy-relevant but relatively complex market. The approach we took was to develop a structural model through which to understand the relationship between firms' pricing behaviour and the implied search behaviour of individuals. This follows on the path-breaking example of Hong and Shum (2006). Essentially, it starts from a simple observation: If each consumer only visited one store when making a purchase, the resulting optimal price for each store to set would be the monopoly price. If a consumer visits two stores, then each has to trade off the increased chance of getting a sale if it sets lower

prices against the lower margin made per sale. If customers vary in their search behaviour, the same basic tradeoff occurs. Increasing the number of stores has two opposing effects: People are less likely to have encountered a low price before visiting your store, but they are also less likely to encounter your store in their search. The basic workhorse model here is due to Stahl (1989), but is substantially developed and extended to suit our case.

By 2002, all British consumers were able to choose their electricity supplier; if they decided not to exercise this choice then they remained with their incumbent supplier, which naturally gave that firm some advantage. But if a consumer decides to switch (i.e. they believe the switching costs are low enough that they will benefit from switching) then they can either switch at random, or search. If consumer *search* costs were zero, then we would still expect a difference in prices as between incumbent and entrants. But amongst entrants, the product is arguably homogeneous to consumers. Also, the major costs (of wholesale electricity from the market, plus distribution and transmission, which are regulated) are common to all entrants (and the incumbent). Therefore we would expect homogeneity of prices across entrants, and for these to tend towards marginal cost. This modified Bertrand outcome is not what we observe. Instead, we see very significant divergences between the prices the different entrants offered at, up to around 25% of the total bill. These divergences persist over time, but not in a way suggestive of product differentiation amongst them but rather with one firm being the cheapest for a while then becoming less competitive, as we see in Figure 3 below. Moreover, unlike in some markets, in this case (over the period we study) we can be assured that all the prices offered are actually faced by some consumers.

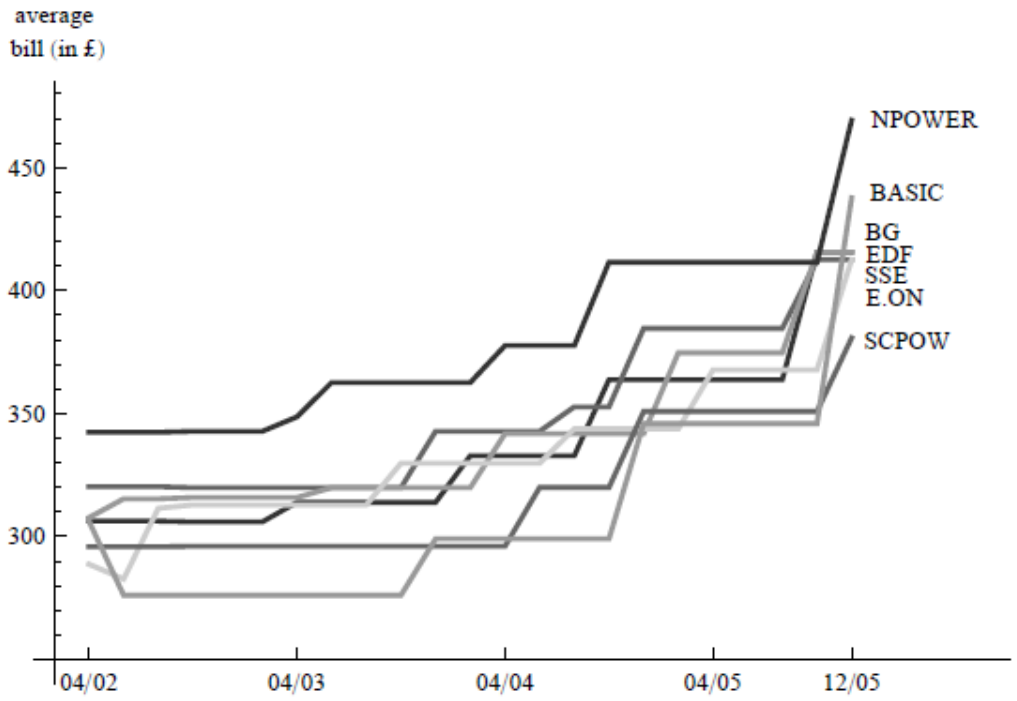

Figure 3: Pricing patterns amongst electricity suppliers in the Midlands area of the Great Britain. Npower is the incumbent, the remaining firms are entrants.
Figure taken from Giulietti et al (2013)

This pattern is suggestive of a mixed strategy equilibrium in pricing based upon search costs, of a type first characterised by Stahl (1989). Indeed, we cannot reject the hypothesis that price ranks are random across firms over time. Therefore, we build a structural model of sequential search (the technology most appropriate at this stage in the market's development). The model is somewhat complex, but can be broken down into capturing three considerations a market entrant should bear in mind in setting its price offer: local consumers who find it optimal to stay with the entrant, switchers from other entrants and switchers from the incumbent. Notice that these imply conflicting objectives: the firm may want to price keenly in order to attract consumers from the incumbent, or even more keenly in order to attract consumers who have chosen another entrant supplier, but may wish to keep prices relatively high to retain those consumers it already has attracted, thereby earning comfortable margins from them. At the time we study the market, price discrimination between new and old customers was not allowed within a region of the country. Following on from this strong competitive force though, once an entrant has attracted significant numbers of customers, its incentives may be to price less keenly, so long as it can retain most of them.

Our estimates suggest that implied search costs started very high, as a proportion of the bill, but declined rather significantly over the period, as we see in Figure 4 which shows the distribution across consumers. Nevertheless, margins did not decline. These findings are reconciled by the point that the entrant firm's keenness to price near to cost is indeed reduced once it has gained a certain market presence. So the use of a structural framework enables us to break down a puzzling observation into two opposing but clear mechanisms. I suggest this is an advantage of the Structural approach, and that this analysis could not be carried out using a quasi-experimental framework.[6]

On the other hand, there is the criticism that the elaborate structure (which our paper certainly has) is only as good as the assumptions underlying it. For example, is it believable that firms practice mixed strategy pricing? Actually, we are able to provide something by way of a test for this. We are unable to reject the hypothesis that concordance in firm ranks across regions of the country and across time is random. We do not claim that firms practice mixed strategies, simply that the data are consistent with this view. We also consider various other key assumptions, for example that products are homogeneous, and find a similar pattern of results in this case.

[6] It would have been possible to ask consumers by questionnaire about their search and switching behaviour and planned search and switching behaviour- this is something we carried out in an earlier study on switching in the gas market (Giulietti et al, 2005). But that is subject to the common economist's criticism that what people say they do and what they actually do are two different things.

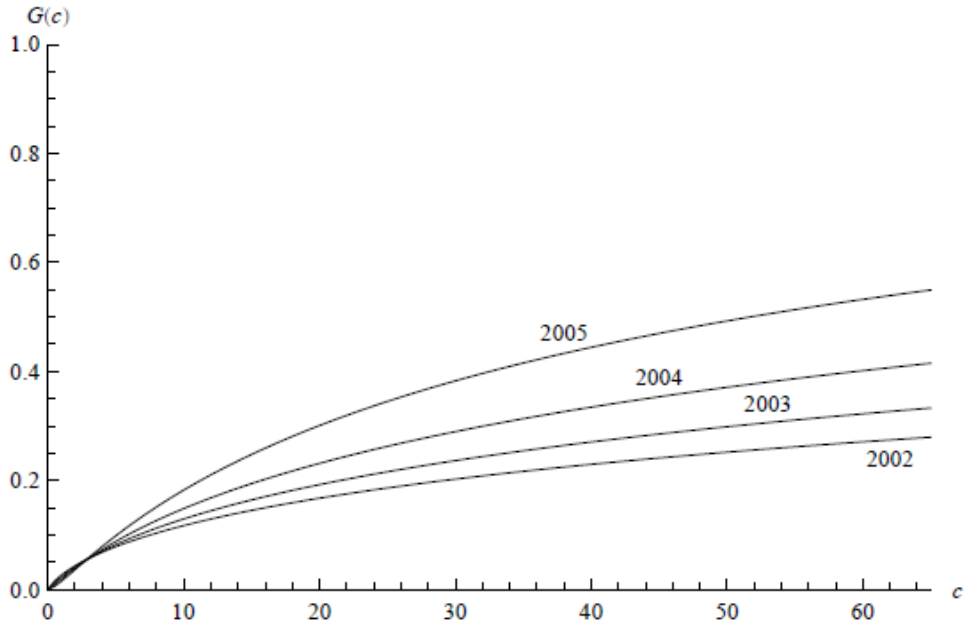

Figure 4: Estimated search cost CDF, in £.
Figure taken from Giulietti et al (2013).

Unusually, we are able to make use of some independent data, not used in estimation, in order to see (a) whether the total number of switchers we predict is roughly in line with the number actually switching and (b) whether market share ranks across firms are consistent with the predictions of the model. In other words, we engage in some post-estimation prediction, albeit needing to make some additional side assumptions in order to do so. The second is a particularly tough test, since the experiences of the firms in the market are rather different from each other. To a substantial degree, our predictions are borne out in the case of both tests, lending confidence to the underlying estimates.

Concluding remarks

To summarize my view, it is that different approaches will tackle different empirical questions within industrial organization. They both have weaknesses. For example the experimental approach has the weakness that a really neat experiment and an important application do not necessarily coincide- the neat experiment may involve a trivial issue; the important application may have too many conflicting influences to make a clean experiment possible. The structural approach has the weakness that it relies on an elaborate superstructure that incorporates maintained assumptions which are unlikely to be valid. But they both also have strengths, as is made clear by their advocates. They are both more scientific than the early work of the 1950s to 1970s. But at the same time, they both involve an element of art as well as science. The art in the experimental approach is in finding an appropriate experiment. The art in the structural approach is in making assumptions that are reasonable and at the same time facilitate the use of the particular model.

This is of course a survey paper, so it is difficult to summarise the socially relevant findings overall. Instead, I discuss this in relation to the two papers of mine, summaries of which I have included in this paper. On the Rough paper, the key

finding is that, although the relative lack of storage did indeed cost a substantial amount of money, in the hundreds of millions of pounds, of course constructing storage is not free. Some crude calculations in the paper suggest that the cost in terms of higher prices identified is probably worth bearing, relative to the cost of construction and operation of storage, at least given conditions at the time. On the consumer search paper, the key socially relevant finding is that, even in a market for what appears like a very homogeneous product, and even given a straightforward nature of search and switching (deliberately engineered to be straightforward), nevertheless the results for consumers can be poor. It is clear here that the market search processes are not leading to a very positive outcome for consumers, in what is in fact a very socially sensitive market where a significant proportion of consumers find themselves unable to afford to heat their homes adequately.

References

Angrist, JD. and Pischke, J-S. (2010) "The Credibility Revolution in Empirical Economics: How Better Research Design is Taking the Con out of Econometrics", *Journal of Economic Perspectives* 24(2), 3-30.

Berry, S., Levinsohn, J. and Pakes, A. (1995) "Automobile prices in market equilibrium", *Econometrica* 63, 841-90.

Bloom, N. (2009) "The impact of uncertainty shocks", *Econometrica* 77(3), 623-685.

Bresnahan, T (1982) "The oligopoly solution concept is identified", *Economics Letters* 10, 87-92.

Bresnahan, T. (1989) "Empirical studies of industries with market power" in *The Handbook of Industrial Organization, vol 2* (R Schmalensee and R Willig, eds), Elsevier, Amsterdam.

Clarke, R., Davies, SW. and Waterson, M. (1984) "The Profitability-Concentration Relation: Market Power or Efficiency?", *Journal of Industrial Economics*, June, pp. 435-450.

Corts, K. (1999) "Conduct parameters and the measurement of market power", *Journal of Econometrics* 88, 227-250.

Genesove, D. and Mullin, W. (1998), "Testing static oligopoly models: conduct and cost in the sugar industry, 1890-1914." Rand Journal of Economics 29, 355-77.

Giulietti, M., Waddams, C. and Waterson, M. (2005) "Consumer choice and Competition Policy: a study of UK Energy Markets" *Economic Journal*, 115, October, 949-968.

Giulietti M, Grossi, L. and Waterson, M. (2012), 'A Rough evaluation of the cost of gas storage' *Energy Journal*, Volume 33 (4), 119-141.

Giulietti, M, Waterson, M and Wildenbeest, M, (2013) "Estimations of search frictions in the British electricity market" Paper accepted for publication in *Journal of Industrial Economics.*

Hong, H. and Shum, M., 2006, 'Using Price Distributions to Estimate Search Costs,' *Rand Journal of Economics* 37, pp. 257-275.

Nevo, A. (2000), "Mergers with differentiated products: the case of the ready-to-eat cereal industry", *Rand Journal of Economics* 31, 395-421.

Niskanen, WA. (2012) http://www.cato.org/publications/commentary/antitrust-staplesoffice-depot-merger

Scherer, FM. (1980) *Industrial Market Structure and Economic Performance*, (2nd edition) Rand McNally.

Sorensen, A. (2000), "Equilibrium price dispersion in retail markets for prescription drugs", *Journal of Political Economy* 108, 833-850.

Stahl, DO. (1989), "Oligopolistic pricing with sequential consumer search", *American Economic Review* 79, 700-712.

Stiglitz, JE. (1986), "Introduction" in JE Stiglitz and GF Mathewson (eds.) *New Developments in the Analysis of Market Structure*, Macmillan, for International Economic Association.

Sutton, J. (1998) *Technology and Market Structure*, MIT press.