

Empirical Tools and Competition Analysis: Past Progress and Current Problems

Ariel Pakes
Harvard University and the NBER

October 4, 2015

I review a subset of the empirical tools available for competition analysis. The tools discussed are those needed for the empirical analysis of; demand, production efficiency, product repositioning, and the evolution of market structure. Where relevant I start with a brief review of tools developed in the 1990's that have recently been incorporated into the analysis of actual policy. The focus is on providing an overview of new developments; both those that are easy to implement, and those that are not quite at that stage yet show promise.

*This is a revised version of my Laffont lecture, given at the 2015 CRESSE conference, Crete.

1 Introduction.

This paper reviews a set of tools that have been developed to enable us to empirically analyze market outcomes, focusing on their possible use in formulating and executing competition policy. I will be more detailed on developments that seem to be less broadly known by the antitrust community (often because they have only been recently introduced), and will emphasize problems that are current roadblocks to expanding the purview of empirical analysis of market outcomes further.

The common thread in the recent developments has been a focus on incorporating the institutional background into our empirical models that is needed to make sense of the data used in analyzing the issues of interest. In large part this was a response to prior developments in Industrial Organization theory which used simplified structures to illustrate how different phenomena could occur. The empirical literature tries to use data and institutional knowledge to narrow the set of possible responses to environmental or policy changes (or the interpretations of past responses to such changes). The field was moving from a description of responses that could occur, to those that were “likely” to occur given what the data could tell us about appropriate functional forms, behavioral assumptions, and environmental conditions.

In pretty much every setting this required incorporating

- heterogeneity of various forms into our empirical models,

and, when analyzing market responses it required

- use of equilibrium conditions to solve for variables that firms could change in response to the environmental change of being analyzed.

The difficulties encountered in incorporating sufficient heterogeneity and/or using equilibrium conditions differed between what was traditionally viewed as “static” and “dynamic” models. The textbook distinction between these two was that: (i) static models analyze price (or quantity) setting equilibria and solve for the resultant profits conditional on state variables, and (ii) dynamics analyzes the evolution of those state variables (and through that the evolution of market structure). The phrase “state variables” typically referred to properties of the market that could only be changed in the medium to long run: the characteristics of the products marketed, the determinants of costs, the distribution of consumer

characteristics, any regulatory or other rules the agents must abide by, and so on.

This distinction has been blurred recently by several authors who note that there are a number of industries in which products of firms already in the market can be "repositioned" (can change one or more characteristic) as easily (or almost as easily) as they can change prices, and then shown that in these industries static analysis of an environmental change that does not take into account this repositioning is likely to be misleading, even in the very short run. The product repositioning literature also employs different empirical tools than had been used to date, and the tools are relatively easy for competition analysis to access and use. So between the static and dynamic sections I will include a section on "product repositioning" in which; (iii) incumbents are allowed to reposition characteristics of the products they market.

2 Static Models.

The empirical methodology for the static analysis typically relied on earlier work by our game theory colleagues for the analytic frameworks we used. The assumptions we took from our theory colleagues included the following:

- Each agent's actions affect all agents' payoffs, and
- At the "equilibrium" or "rest point"
 - (i) agents have "consistent" perceptions¹, and
 - (ii) each agent does the best they can conditional on their perceptions of competitors' and nature's behavior.

Our contribution was the development of an ability to incorporate heterogeneity into the analysis and then adapt the framework to the richness of different real world institutions.

The heterogeneity appeared in different forms depending on the issue being analyzed. When analyzing production efficiency the heterogeneity was needed to rationalize the persistent differences in plant (or firm) level productivity that emerged from the growing number of data sets with plant (or firm) level data on production inputs and outputs. The productivity differences generate both

¹Though the form in which the consistency condition was required to hold differed across applications.

a simultaneity problem resulting from input choices being related to the (unobserved) productivity term and a selection problem as firms who exit are disproportionately those with lower productivity. However once these problems are accounted for, the new data sets enable an analysis of changes in the environment (e.g; a merger) on: (i) changes in the efficiency of the output allocation among firms, and (ii) the changes in productivities of individual producing units (Olley and Pakes, 1996). In demand analysis, allowing for observed and unobserved heterogeneity in both product characteristics and in consumer preferences enabled us to adapt McFadden's (1974) seminal random utility model to problems that arise in the analysis of market demand. In particular allowing for unobserved product characteristics enabled us to account for price endogeneity in a model with rich distributions of unobserved consumer characteristics. These rich distributions enabled us to both mimic actual choice patterns in the emerging data sets on individual (or household) choices and to seamlessly combine that data with aggregate data on market shares prices and distributional information on household attributes (Berry, Levinsohn and Pakes; 1995, 2004).

A major role of equilibrium conditions was to enable an analysis of counterfactuals (e.g.; proposed mergers, regulatory or tax changes,). Few believed that a market's response to a change in environmental conditions was to immediately "jump" to a new equilibrium. On the other hand most also believed that the agents would keep on changing their choice variables until an equilibrium condition of some form was reached; i.e. until each agent was doing the best they could given the actions of the other agents. So the Nash condition was viewed as a "rest point" to a model of interacting agents. Moreover at least without a generally accepted model of how firms learn and adapt to changes in the environment, it was natural to analyze the counterfactual in terms of such a rest point².

2.1 Demand Models and Nash Pricing Behavior.

This material is reasonably well known so I will be brief. The advances made in demand analysis were focused on models where products were described by their characteristics. I start with Berry Levinsohn and Pakes (1995; henceforth

²The equilibrium conditions were also used to: (i) compare the fit of various behavioral models for setting prices of quantities (e.g.; Nevo), and to (ii) generate additional constraints on the relationship between the primitives and control variables that were often a great help in estimation (e.g. Berry Levinsohn and Pakes, 1995).

BLP) model for the utility of agent i would obtain from product j or

$$U_{i,j} = \sum_k \beta_{i,k} x_{j,k} + \alpha_i p_j + \xi_j + \epsilon_{i,j}.$$

Here the $x_{j,k}$ and the ξ_j are product characteristics. The difference between them is that the $x_{j,k}$ are observed by the analyst and the ξ_j are not. The $\beta_{i,k}$ and α_i are household characteristics and the fact that they are indexed by i allows different households to value the same characteristic differently.

Each individual (or household) is assumed to chose the product (the j) that maximized its utility (its $U_{i,j}(\cdot)$), and aggregate demand was obtained by simply summing over the choices of the individuals in the market (a tactic made possible by the speed of computers and the advent of simulation estimators; see McFadden, 1989, and Pakes and Pollard, 1989). The main advantages of this framework are that

- Characteristic space enables us to overcome the "too-many parameters" problem in models where the demand for a given product depended on the prices of all products marketed (their if J is the number of products, the number of price coefficients needed grows like J^2). In characteristic based models the number of parameters estimated grows with the number of characteristics.
- The $\{\xi_j\}$ account for unobserved product characteristics and enable an analysis which does not need to condition on all characteristics that matter to the consumer.
- The $\{\beta_{i,k}\}$ and $\{\alpha_i\}$ enable us to generate own and cross price elasticities that depend on the similarity of the characteristics and prices of different products. They can be made a function of income and other observed characteristics of the consuming unit as well as an unobserved determinant of tastes.

Assume we are studying consumer demand and pricing behavior in a retail market, and for simplicity consider a constant marginal cost firm which has only one product to sell in this market. It sets its prices to maximize its profits from the sale of this product. Profits are given by $p_1 q_1(\cdot)(p_1, \dots, p_J) - mc_1 q_1(\cdot)$ where $q_1(\cdot)$ is the quantity demanded (which depends on the prices and characteristics of all the competing firms) and mc_1 is its marginal cost. As is well

known this implies the price setting equation

$$p_1 = mc_1 + \frac{1}{[\partial q_1 / \partial p_1] / q_1} = x_1 \beta + w_1 \gamma + \omega_1 + \frac{1}{[\partial q_1 / \partial p_1] / q_1}, \quad (1)$$

where $mc_1 = x_1 \beta + w_1 \gamma + \omega_1$, x_1 is the characteristics of the product, w_1 observed market wage rates, and ω_1 is unobserved "productivity"³.

There are a couple of points to keep in mind about equation (1). The markup above costs here is, $\frac{1}{[\partial q_1 / \partial p_1] / q_1}$. It can be calculated directly from the demand estimates (we do not need to estimate a pricing equation to get it). Also this markup has a coefficient of one in the pricing equation which; (i) is a testable restriction, and (ii) enables us to back out marginal costs conditional on the model⁴.

As an example of how well this equation fits data, I asked Tom Wollman to use the data underlying his thesis on the commercial truck market to do the following: (i) use his demand estimates to construct the markup term above, (ii) regress the estimated markup on "instruments" to obtain a predicted markup which is not correlated with the product and cost specific unobservables (ξ and ω)⁵, and (iii) regress the observed price on the observed cost determinants and this predicted markup.

A summary of his results is reported in table one. There are two major points to note. First, the fit of the equation is extraordinary for a behavioral model in economics (that is we are fitting a variable which is chosen by the firm). Second the coefficient of the markup is very close to one (it is within one standard deviation of one). It is also interesting that level shifts in demand over time that have a common effect on all firms' prices account for about 60% of the unexplained variance. This latter point is indicative of the fact that our pricing models do not do as well in fitting shifts in prices over time as they do in explaining the cross-sectional differences in prices. For example, here the R^2 for the cross sectional variance is between .85 and .95 whereas when we used the same data to look at price changes of a given product over time we got an

³The equilibrium pricing equation has an extra term in models with adverse selection, as is frequently the case in financial and insurance markets. Adverse (or beneficial) selection causes costs to be a function of price and so will cause a derivative of costs with respect to price to enter the price setting equation; see Einav, Jenkins and Levin, 2012.

⁴Of course, if other sources of cost data are available, and the competition authorities can sometimes requisition internal firm reports on costs, they should be incorporated into the analysis.

⁵The instruments we used were those used by BLP (1995). For a discussion of the performance of alternative instruments see Reynaert and Verboven, 2014

Table 1: Fit of Pricing Equilibrium.

	Price	(S.E.)	Price	(S.E.)
Gross Weight	.36	(0.01)	.36	(.003)
Cab-over	.13	(0.01)	.13	(0.01)
Compact front	-.19	(0.04)	0.21	(0.03)
long cab	-.01	(0.04)	0.03	(0.03)
Wage	.08	(.003)	0.08	(.003)
$\hat{Markup}$	.92	(0.31)	1.12	(0.22)
Time dummies?	No	n.r.	Yes	n.r.
R^2	0.86	n.r.	0.94	n.r.

Note. There are 1,777 observations; 16 firms over the period 1992-2012. S.E.=Standard error.

R^2 of .5 to .6 (which though smaller, is still quite large for a behavioral model in economics). The difference between the cross sectional and time series fit is indicative of the fact that we fit "rest points" better then we fit the transitions between rest points.

Horizontal Mergers and the Unilateral Effects Model. The unilateral effects price for a single product firm with constant marginal cost is given by the Nash pricing equilibrium formalized in equation (1) above. The intuition behind that equation is the following: if the firm increases its price by a dollar it gets an extra dollar from the consumers who keep purchasing the firm's product, but it loses the markup from the consumers who switch out of purchasing the good because of the price increase. The firm keeps increasing its price until the two forces' contribution to profits negate each other.

When firm 1 merges with another single product firm, our firm 2, this logic needs to be modified. Now if the firm increases its price it still gets the extra dollar from those who stay, but it no longer loses the markup from every consumer who switches out. This because some of the consumers who leave go to what used to be firm 2's product, and the merged firm captures the markup from those who switch to that product. This causes price to increase beyond the price charged by firm 1 when it only marketed one product. The extent of increase depends on the fraction of the customer's who switch out of firm 1's product that switch to what was firm 2's product, and the markup on the second

product. Formally the post merger price (p^m) for firm 2 is written as

$$p_1^m = mc_1 + \frac{1}{[\partial q_1 / \partial p_1] / q_1} + (p_2^m - mc_2) \frac{\partial q_2}{\partial p_1} / \frac{\partial q_1}{\partial p_1}. \quad (2)$$

The difference between the post merger price in equation (2), our p^m , and the pre merger price in (1), our p , has been labelled the "upward pricing pressure" or "UPP" (see DOJ and FTC, 2010; and Farrell and Shapiro, 1990). If we let the merger create efficiency gains of $E_1\%$ in marginal costs we have the following approximation⁶

$$UPP_1 \approx p_1^m - p_1 = (p_2 - mc_2) \frac{\partial q_2}{\partial p_1} / \frac{\partial q_1}{\partial p_1} - E_1 mc_1.$$

Here $\frac{\partial q_2}{\partial p_1} / \frac{\partial q_1}{\partial p_1}$ is the fraction of consumers who leave the good owned by firm 1 that actually switch to firm 2 and is often called the "diversion ration".

The concept of "UPP" and diversion ratios in merger analysis, like many other ideas taken from Industrial Organization into antitrust analysis, has been in graduate class notes for many years. Partly as a result we have had time to consider both when it is appropriate, and the data requirements needed to use it. A full discussion of these issues is beyond what I wanted to do here⁷, but I did want to stress that the UPP analysis is not always appropriate. This is rather obvious when co-ordinated (in contrast to unilateral) behavior is a possibility, so I am going to focus on two more subtle features of the environment which, when present, are likely to make UPP analysis inaccurate. I focus on these two because a more telling analysis of both of them is now possible. They are (i) mergers when product repositioning is likely (I return to this below), and (ii) mergers in the upstream portion of vertical markets in which there are a small number of buyers and sellers.

The basic difference between markets with a small number of buyers and sellers (a buyer-seller network) and retail markets is that in a retail market we generally assume that buyers only options are to buy or not buy at the current price. In a buyer-seller network the buyer can make a counteroffer to any given price; i.e. the buyer can decline to buy at price p but offer to buy at a lower price. Negotiations ensue and if a contract is formed it determines how the

⁶This is an approximation because the merger will cause a change also in the price of the second good which should be factored into an equilibrium analysis. This modifies the formulae (see Pakes, 2011), but not the intuition.

⁷My own discussion of these issues can be found at <http://scholar.harvard.edu/pakes/presentations/comments-onupward-pricing-pressure-new-merger-guidelines;>

profits from the relationship are split⁸. There is no general agreement on the appropriate equilibrium concept for these markets, though there is agreement that each agents' "outside option" (i.e. the profits it would earn were it not to enter into an agreement) will be a determinant of both which relationships are established and the split of the profits among those that are. The implication of this that is not obvious from the UPP formula (and partly as a result is often ignored) is that any analysis of a change in the upstream market must take account of the structure of the downstream market and visa versa.

The empirical literature has found it most convenient to use Horn and Wolinsky's (1988) "Nash in Nash" assumptions as a basis for the analysis of these situations. In that framework negotiations are bilateral; if two agents contract the split of the profits between them is determined by Nash bargaining, and a Nash condition is imposed on who contracts with whom. The Nash condition is that if an agent contracts with another it must earn more profits from the situation generated by the contract than it would absent the contract, and visa versa for at least one agent were the agents not to contract. One questionable assumption is that the deviations are evaluated under the assumption that no other changes in the contracting situation would occur in response to the deviation.

The paper that laid the groundwork for empirical work in this area is Crawford and Yorukoglu (2013), who analyzed the likely impact of forced de-bundling of the tiers set by Cable TV providers. The estimate of the resulting change in consumer welfare differed dramatically by whether the upstream market (the market between the content providers and the cable networks) was allowed to adjust to the new downstream situation (and many content providers would not survive if no realignment of upstream contracts was allowed). Interestingly a variant of their model was also used by the FCC to analyze the likely impacts of the Comcast-NBCU merger⁹. Since that time related frameworks have been used to analyze the likely impacts of hospital mergers (Gowrisankaran, G., A. Nevo, and R. Town, 2015), whose impacts on insurance premiums would depend on the nature of the downstream premium setting game between insurers, and health insurer mergers (Ho, K. and R. Lee, 2015) whose impact on hospital

⁸There has been a good deal of empirical work recently on other markets where a Nash in price (or quantity) assumption that underlies the UPP analysis is inappropriate. Two others, are markets with repeated auctions such as those often used by procurement agencies (see Porter, 2015, for an extended discussion). and markets where there is a centralized matching mechanism, such as the medical match (see Agarwal, 2015, for a start at analyzing market structure issues in these markets).

⁹See https://apps.fcc.gov/edocs_public/attachmatch/FCC-11-4A1.pdf.

prices depended on the nature of the upstream market between insurers and hospitals¹⁰. The current structural changes in the nature of health care provision in the United States make this sort of analysis both timely, and economically important.

2.2 Firm and Industry Productivity.

Partly due to various public agencies providing access to Census like files, there has been a marked increase in the availability of data on the inputs and outputs of production units (firms and/or plants). The advantage over industry level data on inputs and outputs is that the micro data has opened up the possibility of separating the analysis of the; (i) sources of productivity increases (or cost decreases) in individual plants, from (ii) the role of inter-firm output allocations, or market structure, in determining industry productivity. So we now are more able to analyze such traditional IO questions as; the role of market structure in inducing competitive efficiencies¹¹, the role of economies of scale and scope in determining efficient inter-firm allocation of outputs, spillovers from inventive activity, the effect of infrastructure on the productivity of individual firms, etc.

Productivity is defined as a measure of output divided by an index of inputs. The input index is typically associated with a production function. Among the simplest such functions used is the Cobb-Douglas, written as

$$y_{i,t} = \beta_0 + \beta_l l_{i,t} + \beta_k k_{i,t} + \beta_m m_{i,t} + \omega_{i,t}, \quad (3)$$

where $y_{i,t}$ is the logarithm of the output measure for firm i in period t , and $k_{i,t}$, $m_{i,t}$ and $l_{i,t}$ are the logarithms of the labor, capital, and materials inputs respectively, while $\omega_{i,t}$ is the unobserved productivity term. So if we let upper case letters be antilogs of their lower case counterparts and $B = \exp[\beta_0]$, then productivity would be measured as

$$p_{i,t} = \frac{Y_{i,t}}{BL_{i,t}^{\beta_l} K_{i,t}^{\beta_k} M_{i,t}^{\beta_m}},$$

or output per unit of the input index.

¹⁰Interestingly, at the time I am writing, the framework set out in this paper being used to analyze mergers in the U.S. health insurance industry.

¹¹For a review of the empirical findings on this, see Holmes and Schmitz, 2010

To construct the productivity measures one needs estimates of the β parameters. There are two well known problems in obtaining those estimates. First the input measures are at least partially under the control of firms, and if per unit input costs are approximately the same across firms (or at least not negatively correlated with productivity), it is in the interest of more productive firms to use larger quantities of those inputs. This generates a positive correlation between input quantities and the unobserved productivity term which is generally referred to as "simultaneity".

Second, basic economics tells us that higher productivity firms and firms with more capital should be more highly valued and, at least absent a perfect market for those inputs, be less likely to exit. As a result highly capitalized firms who exit tend to have disproportionately low productivity and poorly capitalized firms who continue tend to have disproportionately high productivity. This selection of continuing firms will; (i) cause a correlation between their productivities and capital of continuing firms, and (ii) will generate a bias in our estimates of the productivity effects of any given environmental change. That is firms whose productivity was negatively impacted by the change exit and, as a result, the impact of the environmental change on their productivity is not measured. The selection effects are often particularly large when, as is often the case, we are investigating the effects of a major environmental change on productivity (example, the impacts of a major regulatory change).

Olley and Pakes (1996) show how use of the implications of equilibrium behavior, both for the choice of inputs and for the decisions of when to exit, can be used by the researcher to circumvent the resulting estimation problems¹². The implementation is made significantly easier and less subject to particular assumptions through use of the then recent advances in semi-parametric estimation techniques (see Powell, 1994, and the literature cited therein). Olley and Pakes (1996) also propose a way of decomposing industry wide productivity, defined as the share weighted average of firm level productivity, which has been used extensively. Industry productivity (say P_t) is shown to be the sum of (i) a term measuring the efficiency of the allocation of output among firms, and (ii) the average of individual firm productivities. Formally

$$P_t = \sum_i p_{i,t} s_{i,t} = \bar{p}_t + \sum_i \Delta p_{i,t} \Delta s_{i,t},$$

¹²For a detailed treatment of this estimation procedure, and the improvements that have been made to it, see the section 2 of Akerberg, Benkard, Berry, and Pakes (2007).

where: $\bar{p}_t$ is the (unweighted) average productivity, $\Delta p_{i,t} = p_{i,t} - \bar{p}_t$, and $\Delta s_{i,t}$ is the firm's share minus the average share. So the term measuring the efficiency of the output allocation is the covariance between share and productivity across firms in the industry (more precisely the numerator of that covariance)¹³.

The empirical content of the Olley and Pakes (1996) article is also of some interest to the study of antitrust. They analyze the evolution of productivity in the telecommunication equipment industry before and after the breakup of A.T.&T and the divestment of its wholly owned equipment subsidiary, Western Electric (the consent decree was filed in 1982 and the breakup and divestment were complete by 1984). They find industry wide productivity growth between 3 to 4% during the period from 1978 to 1987 (this includes a large fall in productivity during an adjustment period of 1981-83). Perhaps more important to regulatory economics, the industry's productivity increases, which were accelerating at the end of the sample period, were: (i) almost entirely caused by the covariance term in the above formula and, (ii) could largely be attributed to the reallocation of capital to more productive firms (often as a result of the exit of unproductive firms). That is the entry and selection effects of the increased competition that resulted from deregulation accounted for most of the productivity effects that occurred just after the breakup.

The production data from most data sets contains information from multi-product producing units, a fact which underlies the widespread use of sales revenue as the output measure in the productivity analysis (indeed this led Olley and Pakes to relabel the "production function" in equation (3) a "sales generating function"). Though it is of interest to know how environmental changes effect sales per unit of inputs, it would clearly also be of more interest to understand how those changes effect price and quantity separately, but for this we would need to know separate quantities and prices of the different outputs produced by the firm. Jan De Loecker and his coauthors have used different methods to circumvent this problem, and then pushed ahead with decomposing sales productivity growth into; (i) growth in markups and, (ii) growth in the productivity of inputs in decreasing costs (or a close approximations thereof). In particular De Loecker and Warzynski (2012) use the input equilibrium conditions for a perfectly variable factor of production in a price taking input market to add a step to the estimation procedure that allows them to separate changes

¹³For a dynamic version of this decomposition which explicitly considers the contributions of entrants and exitors in change in P_t over time see Melitz and Polanec, 2015.

in revenue into: (i) changes in markups, and (ii) productivity related changes in costs.

A good example of the usefulness of this type of analysis appears in a subsequent paper by De Loecker, J., P. Goldberg, A. Khandelwal and N. Pavcnik, (forthcoming). They study the impacts of a trade liberalization in India on productivity prices and costs (the liberalization began in 1991, and they study data from 1989 and 1997). There have been a number of prior analysis of the productivity impacts of trade liberalization but most used a framework in which markups were not allowed to change. The new analysis allowed firm specific markups, and analyzed how the markups varied after the tariff change. The tariff declines lowered both input and output prices. The input price declines resulted in disproportionate increases in markups, rather than to reductions in consumer prices; i.e. there was limited pass through of the fall in producer costs to consumer prices (at least during the time period they analyze). The markup increase did induce entry and an increase in the variety of goods available in the market. An analysis of why this occurred would undoubtedly require more institutional information on the industries studied and more detailed analysis of the equilibrium generating the pass through they find. However their study is clearly an important first step in our understanding of how the impact of a major change in the economic environment impacted a large share of the world's population.

3 Product Repositioning.

In some industries incumbents can "reposition" their products as fast, or almost as fast, as prices can be changed. The most dramatic illustration of this that I am aware of is contained in a series of figures in Chris Nosko's thesis (see Nosko 2014)¹⁴.

Nosko analyzes the response of the market for CPU's in desktop computers to Intel's introduction of the Core 2 Duo generation of chips. Two firms dominate this market: Intel and AMD. Nosko explains that it is relatively easy to change chip performance provided it is lower than the best performance from the current generation of chips. Indeed he shows that chips sold at a given price typically change their characteristics about as often as price changes on a given

¹⁴I thank him for allowing me to reproduce two of them.

set of characteristics.

The first figure provides benchmark scores and prices for the products offered in June 2006, just prior to the introduction of the Core 2 Duo. The red and blue dots represent AMD's and Intel's offerings, respectively. Note that in June there was intense competition for high performance chips with AMD selling the highest priced product at just over \$1000. Seven chips sold at prices between \$1000 and \$600, and another five between \$600 and \$400. July saw the introduction of the Core 2 Duo and figure 2 shows that by October; (i) AMD no longer markets any high performance chips (their highest price chip in October is just over two hundred dollars), and (ii) there are no chips offered between \$1000 and \$600 dollars and only two between \$600 and \$400 dollars. Shortly thereafter Intel replaces the non-Core 2 Duo chips with Core 2 Duo's.

Nosko goes on to explain how the returns from the research that went into the Core 2 Duo came primarily from the markups Intel was able to earn as a result of emptying out the space of middle priced chips and dominating the high priced end of the spectrum. He also shows, through a counterfactual merger simulation, that a similar phenomena would likely occur if AMD were to merge with Intel.

Other recent articles that employ an analysis of product repositioning include Eizenberg (2014) and Wollman (2015). Eizenberg studies the introduction of the Pentium 4 chip in PC's and notebooks. He focuses on the decisions to stop the production of products with older chips (and lower prices), a decision relative easy to implement almost immediately, and the implications of those decisions on consumer welfare. He finds distributional effects to withdrawing goods; though overall welfare is likely higher with the withdrawals poorer consumers are better off without them. Wollman (2015) considers the bailout of GM and Chrysler during the recent financial crisis, and asks what would have happened had GM and Chrysler not been bailed out, but rather exited the commercial truck market. He notes that part of the commercial truck production process is modular (it is possible to connect different cab types to different trailers), so some product repositioning would have been almost immediate. In addition to "pure exit" he considers the possibility of mergers between the "exiting" truck makers and the remaining incumbents, and he calculates all his counterfactuals twice; once just allowing the prices of the products offered by the participants who remained after the counterfactual to adjust, and once

allowing their product offerings as well as their prices to adjust. Allowing for changes in characteristics had policy relevant differences on the results.

The analytic framework used in these papers is that of a two period model, with product offerings set in the first stage and prices set in the second. Detailed structural models, in the spirit of the models referred to above, are used to estimate demand and cost, and a Nash pricing equilibrium is used to set prices when needed. This provides the structure used in the second stage which needs to calculate profits for different sets of products. The analysis of the first stage, i.e. the calculation of whether to add (or delete) different products, requires an estimate of an additional parameter: the fixed costs of adding or dropping a product from the market. It is estimated using the profit inequality approach proposed in Pakes, Porter, Ho and Ishii (2015; henceforth PPHI) and Pakes (2010).

The profit inequality approach is analogous to the revealed preference approach to demand analysis when the choice set is discrete. To understand how it works assume the fixed costs are the same for adding any product., and let $x_j = [1, 0]^L$ where L is the number of products that firm j could offer and the vector x_j has a 1 for products offered and a 0 otherwise. Let e_z be an L -dimensional vector with one in the "z" spot and zero elsewhere, and assume that the z^{th} product was just added. We use the framework above to compute both the actual profits and the implied profits had the product not been added¹⁵. Let $(\pi_j(x_j, x_{-j}), \pi_j(x_j - e_z, x_{-j}))$, be the profits with and without marketing the z^{th} product, $\mathcal{I}_j$ be the agent's information set, and $E[\cdot|\mathcal{I}_j]$ deliver expectations conditional on $\mathcal{I}_j$. The behavioral assumption is that z was added because


$$E[\pi_j(x_j, x_{-j}) - \pi_j(x_j - e_z, x_{-j})|\mathcal{I}_j] \geq F,$$

where F is the fixed cost of adding a product. If we average over all the products introduced and assume expectations are unbiased, we get a consistent lower bound for F . On the other hand if the z^{th} product is a feasible addition that was *not* offered, and $(\pi_j(x_j, x_{-j}), \pi_j(x_j + e_z, x_{-j}))$ were the profits without and with that product, then


$$E[\pi_j(x_j + e_z, x_{-j}) - \pi_j(x_j, x_{-j}), |\mathcal{I}_j] \leq F,$$

which gives us an upper bound to F ,


¹⁵This would have been a unilateral deviation in the first stage simultaneous move game, and hence not changed the products marketed by other firms.




Price/Performance – June 2006



Price/Performance – July 2006



Price/Performance – Oct 2006



Price/Performance – January 2008

Notice that it is the power of the structural approach to make out of sample predictions, the same aspect of that approach used in the analysis of mergers, that lets us estimate fixed costs. There are different issues that can arise here. First we may want to allow for differential fixed costs across products or locations. A straightforward generalization of what I just described can allow fixed costs to differ as a function of observable variables (like product characteristics)¹⁶. Things do get more complicated when one allows for unobservable factors that generate differences in fixed costs and were known to the agents when they made their product choice decision. This because the products that are provided are likely to have been partially selected on the basis of having unobservable fixed costs that were lower than average, and those that are not may have been partially selected for having higher than average fixed costs. Possible ways for correcting for any bias in the bounds that selection generates are provided in PPHI and in Manski (2003); see Eizenberg (2014) for an application.

One might also worry that the two period model misses the dynamic aspects of marketing decisions that would arise if F represented sunk, rather than fixed, costs. A more complete analysis of this situation would require the sequential Markov equilibrium discussed in the next section. However there is a way to partially circumvent the complexity of Markov equilibria. If it was feasible to market z but it was not offered and we are willing to assume the firm can credibly commit to withdrawing it in the next period *before* competitors next period decisions are taken, then we can still get an upper bound. The upper bound is now to the cost of marketing and then withdrawing the product. I.e. our behavioral assumption now imply that the difference in value between (i) adding a product not added and then withdrawing it in the next period and (ii) marketing the products actually marketed, is less than zero. This implies

$$E[\pi_j(x_j + e_z, x_{-j}) - \pi_j(x_j, x_{-j}) | \mathcal{I}_j] \leq F + \beta W.$$

where $W \geq 0$ is the cost of withdrawing and β be the discount rate. Lower bounds require further assumptions, but the upper bound ought to be enough for examining extremely profitable repositioning moves following environmental changes (like those discussed in Nosko (2011)).

Given these bounds we turn to considering likely product repositioning decisions that follow an environmental change. Typically if we allow the product

¹⁶Though this does require use of estimation procedures from the literature on moment inequality estimators. For a review of estimation using moment inequalities see Tamer (2010).

mix of incumbents to adjust as a result of say a merger or an exit there will be a number of market structures that are Nash equilibria in the counterfactual environment (though that number will be limited by profitability implications of the investments in place).

There are a number of ways to chose a "rest point" among the possible equilibria. One is to add a model for how firm's adjust to changes in their environment. This was done in Lee and Pakes (2009) and in Wollman (2015). We could assume that firms' assume that their competitors do not change their policies, and find a best response to those policies given the changed environment. Alternatively firms' could be thought to best respond to the best response to prior policies (an "iterated" best response), and there are other possibilities ¹⁷. Alternatively one could attempt to enumerate the possible equilibria (or at least those that seem likely on the basis of some exogenous criteria), and consider properties of all members of that set, as is done in Eizenberg (2014)¹⁸.

4 Dynamics and the Evolution of State Variables.

Empirical work on dynamic models proceeded in a similar way to the way we proceeded in static analysis; we took the analytic framework from our theory colleagues and tried to incorporate the institutions that seemed necessary to analyze actual markets. We focused on models where

1. state variables evolve as a Markov process,
2. and the equilibrium was some form of Markov Perfection (no agent has an incentive to deviate at any value of the state variables).

In these models firms chose "dynamic controls" (investments of different types) and these determine the likely evolution of their state variables. Implicit in the second condition above is that players' have perceptions of the controls' likely impact on the evolution of the state variables (their own and those of their competitors) and through that on their current and future profits, and that these

¹⁷In a world where there is randomness in the environment, the learning processes can lead to a probability distribution of possible equilibria. For a theoretical treatment of alternatives see Fudenberg and Levine, 1998, and for an attempt to distinguish between learning algorithms actually used in an industry, see Doraszelski, Lewis, and Pakes, 2015.

¹⁸If the fixed costs estimates bounds define a reasonably large interval of possible estimates, one would have to calculate these counterfactuals for a sufficiently rich set of values for that parameter.

perceptions are consistent with actual behavior (by nature, as well as by their competitors). The standard references here are Maskin and Tirole (1988a and b) for the equilibrium notion and Ericson and Pakes (1995) for the framework used in computational theory and empirical work.

The Markov assumption is both convenient and, except in situations involving active experimentation and learning, fits the data well¹⁹, so empirical work is likely to stick with it. The type of rationality built into Markov Perfection is more questionable. It has been useful in correcting for dynamic phenomena in empirical problems which do not require a full specification of the dynamic equilibrium (as in the Olley and Pakes example above),. The Markov Perfect framework also enabled applied theorists to explore possible dynamic outcomes in a structured way. This was part of the goal of the original Maskin and Tirole articles and computational advances have enabled Markov Perfection to be used in more detailed computational analysis of a number of other issues that were both, integral to antitrust analysis, and had not been possible before²⁰. Examples include; the relationship of collusion to consumer welfare when we endogenize investment entry and exit decisions (as well as price, see Fershtman and Pakes, 2000), understanding the multiplicity of possible equilibria in models with learning by doing (Besanko, Doraszelski, Kryukov and Satterthwaite, 2010), and dynamic market responses to merger policy (Mermelstein, Nocke, Satterthwaite, and Whinston 2014).

Perhaps not surprisingly, the applied theory indicated just how many different outcomes were possible by perturbing functional forms, timing, or just exploring the multiple equilibria of Markov perfect models. This provided additional incentives for empirical work, but though there has been some compelling dynamic empirical work using the Markov Perfect framework (see for e.g.; Benkard, 2004, Collard-Wexler, 2013, and Kalouptsi, 2014)²¹, it has not been as forthcoming as one might have hoped. This is because the framework becomes unweildly when confronted with the task of incorporating the institutional background that seemed relevant.

¹⁹It also does not impose unrealistic data access and retention conditions for decision making agents.

²⁰The computational advances enabled us to compute equilibria quicker and/or with less memory requirements than in the simple iterative procedure used in the original Pakes and McGuire (1994) article. The advances include; Judd's (1998) use of deterministic approximation techniques, Pakes and McGuire (2001)'s use of stochastic approximation, and Doraszelski and Judd (2011)'s use of continuous time to simplify computation of continuation values.

²¹A review by Doraszelski and Pakes, 2000 provides a more complete list of cites to that date

When we try to incorporate the institutional background that seems essential to understanding the dynamics of the market we often find that both the analyst, and the agents we are trying to model, are required to: (i) access a large amount of information (all state variables), and (ii) either compute or learn an unrealistic number of strategies. To see just how complicated the dynamics can become, consider a symmetric information Markov Perfect equilibrium where demand has a forward looking component; as it would if we were studying a durable, storable, experience, or network good.

For specificity consider the market for a durable good. Both consumers and producers would hold in memory at the very least; (i) the cartesian product of the current distribution of holdings of the good across households crossed with household characteristics, and (ii) each firm's cost functions (one for producing existing products and one for the development of new products). Consumers would hold this information in memory, form a perception of the likely product characteristics and prices of future offerings, and compute a dynamic programme to determine their choices. Firms would use the same state variables, take consumers decisions as given, and compute their equilibrium pricing and product development strategies. Since these strategies would not generally be consistent with the perceptions that determined the consumers' decisions, the strategies would then have to be communicated back to consumers who would then have to recompute their dynamic program using the updated firm strategies. This process would need to be repeated until we found strategies where consumers do the best they can given correct perceptions of what producers would do and producers do the best they can given correct perceptions on what each consumer would do (a doubly nested fixed point calculation). Allowing for asymmetric information could reduce information requirements, but it would substantially increase the burden of computing optimal strategies. The additional burden results from the need to compute posteriors, as well as optimal policies; and the requirement that they be consistent with one another.

It is hard to believe that this is as good an approximation to actual behavior as we can come up with. So I want to point to other notions of equilibria which might both better approximate agents' behavior and enable empirical work on a broader set of dynamic issues. This work is just beginning, so I will focus on the concepts that underlie it and some indication of how far we have gotten.

4.1 Less Demanding Notions of Equilibria.

The framework I will focus on is "Restricted Experience Based Equilibrium" (or REBE; in Fershtman and Pakes, 2012). This is similar in spirit to "Self-confirming Equilibrium" introduced in Fudenberg and Levine (1983), and since refined in different ways (e.g., Battigalli, P. et. al, 2015). The major difference is that REBE uses a different "state space" (one appropriate for dynamic games), and as a result has to consider a different set of issues.

A REBE equilibrium satisfies two conditions which seem natural for a "rest point" to a dynamical system. They are

1. agents perceive that they are doing the best they can conditional on the information that they condition their actions on, and that
2. if the information set that they condition on has been visited repeatedly, these perceptions are consistent with what they have observed in the past.

Notice that this does not assume that agents form their perceptions in any particular way; just that they are consistent with what they have observed in the past at conditioning sets that are observed repeatedly²².

The state space consists of the information sets the agents playing the game condition their actions on. The information set of firm i in period t is denoted by $J_{i,t} = \{\xi_t, \omega_{i,t}\}$, where ξ_t is public information observed by all, and $\omega_{i,t}$ is private information. The private information is often information on production costs or investment activity (and/or its outcomes), and is allowed to be serially correlated (to enhance the model's ability to mimic data). The public information varies with the structure of the market. It can contain publicly observed exogenous processes (e.g. information on factor price and demand), past publicly observed choices made by participants (e.g. past prices), and whatever has been revealed over time on past values of $\omega_{-i,t}$.

Firms chose their "controls" as a function of $J_{i,t}$. Since we have allowed for private information, these decisions need not be a function of all the variables that determine the evolution of the market (it will not depend on the private information of competitors²³). Relative to a symmetric information Markov

²²It might be reasonable to assume more than this, for example that agents know and/or explore properties of outcomes of states not visited repeatedly, or to impose restrictions that are consistent with data on the industry of interest. If so this would mitigate the multiplicity issues discussed below.

²³The model also allows agents to chose to ignore information that they think is less relevant. We come back to how the empirical researcher determines $J_{i,t}$ below.

equilibrium, this reduces both what the agent needs to keep track of, and the number of distinct policies the agent needs to form. Typically potential entrants will choose whether to enter, and incumbents will choose whether to remain active and if so their prices and investments (in capital, R&D, advertising, ...). These choices are made to maximize their perceptions of the expected discounted value of future net cash flows, but their perceptions need not be correct.

The specification for the outcomes of the investment and pricing process, and for the revelation of information, determines the next period's information set. For example assume costs are serially correlated and are private information, and that prices are a function of costs and observed by all participants. Then past price is a signal on current costs, and we would expect all prices in period t to be components of ξ_{t+1} , the public information in period $t + 1$. If, in addition, investment is not observed and it generates a probability distribution for reductions in cost, then the realization of the cost variable is a component of $\omega_{i,t+1}$, the private information in period $t + 1$.

Since agents' choices and states are determined by their information sets, the "state" of the industry, which we label as s_t , is determined by the collection of information sets of the firms within it

$$s_t = \{J_{1,t}, \dots, J_{n_t,t}\} \in \mathcal{S}.$$

Assumptions are made which insure that s_t evolves as a finite state Markov chain. This implies that, no matter the policies, s_t will wander into a recurrent subset of the possible states (i.e. of $\mathcal{S}$), and then remain within that subset forever (Freedman, 1971). Call that subset $\mathcal{R} \subset \mathcal{S}$. These states are determined by the primitives of the market (its size, feasible outcomes from investment, ...), and they are visited repeatedly. For example a small market may never see more than x firms simultaneously active, but in a larger market the maximum firms ever active may be $y > x$. Firms in the larger market may flounder and eventually exit, but it is possible that before the number of active firms ever falls to x there will be entry. Then the recurrent class for the larger market does not include $J_{i,t}$ that are a component of an s_t that has less than x firms active.

The $J_{i,t}$ which are the components of the s_t in $\mathcal{R}$ are the information sets at which the equilibrium conditions require accurate perceptions of the returns to feasible actions. So in industries that have been active for some time, neither the agent nor the analyst needs to calculate equilibrium values and policies for the information sets in all of $\mathcal{S}$, we only need them for those in $\mathcal{R}$, and $\mathcal{R}$ can

be much smaller than $\mathcal{S}$.

The Fershtman and Pakes (2012) article provide a learning algorithm, in the spirit of reinforcement learning²⁴ that enables the analyst to compute a REBE. Briefly, the algorithm starts at some initial s_t and has a formulaic method of generating initial perceptions of the expected discounted values of alternative actions at each possible $J_{i,t}$. Each agent choses the action which maximizes its initial perceptions of its values. The actions generate a probability distribution over outcomes, and a pseudo random draw from those distributions plus the rules on information access determine both the current profit and the new state (our $J_{i,t+1}$). Then the current profit and the new state are used to update the perceptions of the values for taking actions in the original state (at $J_{i,t}$). More precisely, the profits are taken as a random draw from the possible profits from the action taken at the old state, and the new state, or rather the initial perception of the value of the new state, is treated as a random draw from the possible continuation values from the action taken at $J_{i,t}$. The random draws on profits and continuation values are averaged with the old perceptions to form new perceptions. This process is then repeated from the new state. So the algorithm is iterative, but each iteration only updates the values and policies at one point (it is "asynchronous"). The policies at that point are used to simulate the next point, and the simulated output is used to update perceptions at the old point. Notice that firms could actually follow these steps to learn their optimal policies, but it is likely to do better as an approximation to how firms react to perturbations in their environment then to major changes in it²⁵.

This process has two computational advantages. First the simulated process eventually wanders into $\mathcal{R}$ and stays their. So the analyst never needs to compute values and policies on all possible states, and the states that are in $\mathcal{R}$ are updated repeatedly. Second the updating of perceptions never requires integration over all possible future states, it just requires averaging two numbers. On the negative side there is no guarantee that the algorithm will converge to a REBE. However Fershtman and Pakes program an easy to compute test for convergence into the algorithm, and if the test output satisfies the convergence

²⁴For an introduction to reinforcement learning see, Sutton, R. and Barto, A. (1998).

²⁵This because the current algorithm does not allow for experimentation, and (ii) can require many visits to a given point before reaching an equilibrium (especially when initial and equilibrium perceptions differ greatly). Doraszelski, Lewis and Pakes (2015) study firms learning policies in a new market and find that there is an initial stage where an analogue of this type of learning does not fit, but the learning algorithm does quite well after an initial period of experimentation.

criteria, the algorithm will have found a REBE.

Multiplicity. A Bayesian Perfect equilibrium satisfies the conditions of a REBE, but so do weaker notions of equilibrium. So REBE admits greater multiplicity than does Bayesian Perfect notions of equilibrium. To explain the major reason for the increase in equilibria partition the points in $\mathcal{R}$ into “interior” and “boundary” points²⁶. Points in $\mathcal{R}$ at which there are feasible (but non-optimal) strategies which can lead outside of $\mathcal{R}$ are boundary points. Interior points are points that can only transit to other points in $\mathcal{R}$ no matter which of the feasible policies are chosen.

The REBE conditions only ensure that perceptions of outcomes are consistent with the results from actual play at interior points. Perceptions of outcomes for feasible (but non-optimal) policies at boundary points need not be tied down by actual outcomes. As a result differing perceptions of discounted values at points outside of the recurrent class can support different equilibria. One can mitigate the multiplicity problem by adding either empirical information or by strengthening the behavioral assumptions.

The empirical information should help identify which equilibria has been played, but may not be of much use when attempting to analyze counterfactuals. The additional restrictions that may be appropriate include the possibility that prior knowledge or past experimentation will endow agents with realistic perceptions of the value of states outside, but close to, the recurrent class. In these cases we will want to impose conditions that insure that the equilibria we compute are consistent with this knowledge. To accommodate this possibility, Asker, Fershtman, Jeon, and Pakes (2015) propose an additional condition on equilibrium play that insures that agents’ perceptions of the outcomes from all feasible actions from points in the recurrent class are consistent with the outcomes that those actions would generate. They label the new condition “boundary consistency” and provide a computational simple test to determine whether the boundary consistency condition is satisfied for a given set of policies.

²⁶This partitioning is introduced in Pakes and McGuire, 2001. There is another type of multiplicity that may be encountered; there may be multiple recurrent classes for a given equilibrium policy vector. Sufficient condition for the policies to generate a process with a unique recurrent class are available (see Freedman, 1971, or Ericson and Pakes, 1995) but there are cases of interest where multiple separate recurrent classes are likely (see Besanko, Doraszeldi and Kryukov, 2014).

Empirical Challenges. In addition to an approximation to the static profit function (discussed in the earlier sections), empirical work on dynamics will require: (i) specifying the variables in J_i and (ii) estimates of the “dynamic” parameters (these usually include the costs of entry and exit, and parametric models for both the evolution of exogenous state variables, and the response of endogenous state variables to the agents’ actions).

There is nothing in our equilibrium conditions that forbids J_i from containing less variables than the decision maker has at its disposal, and a restricted J_i may well provide a better approximation to behavior. The empirical researcher’s specification of J_i should be designed to approximate how the dynamic controls are determined (in the simplest model this would include investment, entry and exit policies). This suggests choosing the variables in J_i through an empirical analysis of the determinants of those controls. Information from the actual decision makers or studies of the industry would help guide this process.

In principle we would like to be able to generate predictions for the dynamic controls that replicate the actual decisions up to a disturbance which is a sum of two components; (i) a “structural” disturbance (a disturbance which is a determinant of the agent’s choice, but we do not observe) which is *independently* distributed over time, and (ii) a measurement error component. The measurement error component should not be correlated with variables which are thought to be correctly measured, and the structural error should not be correlated with any variable dated prior to the period for which the control is being predicted. So the joint disturbance should be uncorrelated with past values of correctly measured variables. This provides one test of whether a particular specification for J_i is adequate. If the null is rejected computational and estimation procedures which allow for serially correlated errors should be adopted. This need not pose additional computational problems, but may well complicate the estimation issues we turn to now.

Estimates of dynamic parameters will typically be obtained from panel data on firms, and the increased availability of such data bodes well for this part of the problem. Many of the dynamic parameters can be estimated by careful analysis of the relationship between observables; i.e. without using any of the constructs that are defined by the equilibrium to the dynamic model (such as expected discounted values). For example if investment is observed and directed at improving a given variable, and (a possibly error prone) measure of that vari-

able is either observed or can be backed out of the profit function analysis, the parameters governing the impact of the control can be estimated directly.

However there often are some parameters that can only be estimated through their relationship to perceived discounted values (sunk and fixed costs often have this feature). There is a review of the literature on estimating these parameters in the third section of Akerberg et. al. (2007). It focuses on two step semi-parametric estimators which avoid computing the fixed point that defines the equilibrium at trial values of the parameter vector²⁷. In addition Pakes (2015) describes a “perturbation” estimator, similar to the Euler equation estimator for single agent dynamic problems proposed by Hansen and Singleton (1982). This estimator does not require the first step non-parametric estimator, but is only available for models with asymmetric information. Integrating serially correlated unobservables into these procedures can pose additional problems; particularly if the choice set is discrete. There has been recent work on discrete choice models that allow for serially correlated unobservables (see Arcidiano and Miller, 2011), but it has yet to be used in problems that involve estimating parameters that determine market dynamics.

5 Conclusion.

The advantage of using the tools discussed here to evaluate policies is that they let the data weigh in on the appropriateness of different functional forms and behavioral assumptions. Of course any actual empirical analysis will have to maintain some assumptions and omit some aspects of the institutional environment. The critical issue, however, is not whether the empirical exercise has all aspects of the environment modeled correctly, but rather whether empirics can do better at counterfactual policy analysis than the next best alternative available. Of course in comparing alternatives we must consider only those that; (i) use the information available at the time the decision is made and (ii) abide by the resource constraints of the policy maker. By now I think it is clear that in some cases empirical exercises can do better than the available alternatives, and that the proportion of such cases is increasing; greatly aided by advances in both computational power and the resourcefulness of the academic community

²⁷The relevant papers here are those of Bajari Benkard and Levin (2007), and Pakes Ostrovsky and Berry, (2007))

(particularly young Industrial Organization scholars).

References.

- Akerberg D, Benkard L, Berry S, Pakes A. (2007); "Econometric Tools for Analyzing Market Outcomes". In Heckman J, Leamer E, ed., *The Handbook of Econometrics* Amsterdam: North-Holland; 2007. pp. Chapter 63.
- Agarwal, N. (2015); "An Empirical Model of the Medical Match", *American Economic Review*, pp 1939-78, July 2015.
- Arcidiano, P. and R. Miller, 2011, "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity", *Econometrica*, Vol. 7, No. 6, 1823-1868
- Asker, J., C. Fershtman, J. Jeon, and A. Pakes (2014) "The Competitive Effects of Information Sharing" in process *Harvard University*.
- Bajari, P., L. Benkard, and J. Levin (2007), "Estimating Dynamic Models of Imperfect Competition," *Econometrica*.
- Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci. 2015. "Self-Confirming Equilibrium and Model Uncertainty." *American Economic Review*, 105(2): 646-77.
- Benkard, L., (2004) "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft" *Review of Economic Studies*, 71, 581-611.
- Berry S., J. Levinsohn, and A. Pakes (1995), "Automobile Prices in Market Equilibrium," *Econometrica*, vol. 63, no. 4, pp. 841-890.
- Berry S., J. Levinsohn, and A. Pakes (2004), "Estimating Differentiated Product Demand Systems from a Combination of Micro and Macro Data: The Market for New Vehicles," *Journal of Political Economy*, vol. 112, no. 1,1, pp. 68-105.
- Besanko, D., U. Doraszelski, Y. Kryukov and M. Satterthwaite (2010), "Learning by Doing, Organizational Forgetting and Industry Dynamics" *Econometrica*, 78 (2) pp.453-508.
- Besanko, D., U. Doraszelski, and Y. Kryukov: "The Economics of Predation: What Drives Pricing When There Is Learning-by-Doing?", *American Economic Review*, 104(3): 868-897.
- Collard-Wexler, A. (2013), "Demand Fluctuations in the Ready-Mix Concrete Industry", *Econometrica*, Vol. 81(3), pp.1003-1037.
- Crawford, G. and A. Yurukoglu (2013); "The Welfare Effects of Bundling in Multichannel Television Markets" *American Economic Review*, 102(2): 643-85.
- De Loecker J., and Warzynski (AER 2012) 2012. "Markups and Firm-Level Export Status." *American Economic Review*, 2437-71.
- De Loecker, J., P. Goldberg, A. Khandelwal and N. Pavcnik, (forthcoming), "Prices, Markups and Trade Reform" *Econometrica*.

- Doraszelski, U., and K. Judd (2011); "Avoiding the Curse of Dimensionality in Dynamic Stochastic Games", *Quantitative Economics*, 3 (1), 53 - 93.
- Doraszelski U., G. Lewis and A. Pakes (2014) "Just starting out: Learning and price competition in a new market" *mimeo* Harvard University.
- Doraszelski, U. and A. Pakes (2008), "A Framework for Applied Dynamic Analysis in I.O.," Chapter 33, pp. 2183-2262, in the *The Handbook of Industrial Organization*, M. Armstrong and R. Porter ed.s.
- Einav L, M. Jenkins and J. Levin (2012) "Contract Pricing in Consumer Credit Markets" *Econometrica*, 80 (4), 1387 - 1432.
- Eizenberg, A. (2014) "Upstream Innovation and Product Variety in the United States Home PC Market," *the Review of Economic Studies*, pp.1003-1045
- Ericson R. and A. Pakes (1995), "Markov Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, vol. 62, no. 1, pp. 53-82.
- Fershtman C. and A. Pakes (2000), "A Dynamic Game with Collusion and Price Wars," *RAND Journal of Economics*, vol. 31, no. 2, pp. 207-36.
- Fershtman C. and A. Pakes (2012); "Dynamic Games with Asymmetric Information: A Framework for Applied Work" *The Quarterly Journal of Economics*. pp 1611-62.
- Fudenberg D. and D. Levine (1983) "Subgame Perfect Equilibrium of Finite and Infinite Horizon Games" *Journal of Economic Theory* 31, 227-256.
- Fong K. and R. Lee (2013) "Markov Perfect Network Formation: An Applied Framework for Bilateral Oligopoly and Bargaining in Buyer-Seller Networks" *mimeo* Harvard University.
- Freedman D. (1971) Markov Chains, Holden Day series in probability and statistics.
- Gowrisankaran, G., A. Nevo, and Town, R. (2015); "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry," *American Economic Review*, 172-203.
- Hansen, L. and K. Singleton (1982) "Generalized instrumental variables estimation of nonlinear rational expectations models", *Econometrica*, vol. 50, 1982.
- Hendriks, K, and R. Porter (2015), "Empirical Analysis and Auction Design" *in process*, Northwestern University.
- Ho, K, and R. Lee (2015): "Insurer Competition in Health Care Markets" *mimeo Harvard University*
- Holmes and Schmitz, 2010); "Competition and Productivity: A Review of Evidence" *Annual Review of Economics*, pp. 619-42.
- Horn H. and A. Wolinsky (1988) "Bilateral monopolies and incentives for merger". *RAND Journal of Economics*. Vol. 19, No. 3, Autumn 1988.
- Judd, K. (1998), Numerical Methods in Economics, MIT Press: Cambridge, MA., 1998.

- Kalouptsi, M (2014), "Time to Build and Fluctuations in Bulk Shipping", *American Economic Review*, 104(2): 564-608.
- Manski (2003): *Partial Identification of Probability Distributions*, Berlin, Heidelberg, New York: Springer-Verlag.
- Maskin, E. and J. Tirole (1988), "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs" *Econometrica*, 56, 549-570.
- McFadden D. (1974), "Conditional Logit Analysis of Qualitative Choice Behavior", in P. Zarembka (ed.) *Frontiers in Econometrics*. Academic Press, New York.
- McFadden D. (1989) "A Method of Simulated Moments for Estimation of Discrete Choice Models", in *Econometrica*, Volume 57, 995-1026.
- Melitz, M., and S. Polanec, (2015), "Dynamic Olley-Pakes Productivity Decomposition with Entry and Exit." *Rand Journal of Economics*.
- B. Mermelstein, V. Nocke, M. Satterthwaite, and M. Whinston (2014), "Internal versus External Growth in Industries with Scale Economies: A Computational Model of Optimal Merger Policy," NBER Working Papers 20051.
- Nosko C. (2013); "Competition and Quality Choice in the CPU Market" Phd dissertation, Harvard University.
- S. Olley and A. Pakes (1996), "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, vol. 64, no. 6, pp. 1263-98.
- A. Pakes and D. Pollard (1989), "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, vol. 57, no. 5, pp. 1027-1057.
- A. Pakes and P. McGuire (1994), "Computing Markov Perfect Nash Equilibrium: Numerical Implications of a Dynamic Differentiated Product Model," *RAND Journal of Economics*, vol. 25, no. 4, pp. 555-589.
- A. Pakes and P. McGuire (2001), "Stochastic Algorithms, Symmetric Markov Perfect Equilibria, and the 'Curse' of Dimensionality", *Econometrica*, vol. 69, no. 5, pp. 1261-81.
- A. Pakes, Ostrovsky M., and Berry S. (2007); "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry-Exit Examples)" *RAND Journal of Economics*, Vol.38 (2) pp 373-399.
- A. Pakes (2010), "Alternative Models for Moment Inequalities". *Econometrica*, Vol. 78, No. 6 (November), pp 1783-1822.
- A. Pakes, J. Porter, K. Ho and J. Ishii, (2015) "Moment Inequalities and Their Application", *Econometrica*, pp 315-334.
- Pakes A, (2015) "Methodological Issues in Analyzing Market Dynamics" working paper, *Harvard University*.
- Powell J. (1994), "Estimation of Semiparametric Models", in the *Handbook of Econometrics*, Volume 4, R. Engle and D. McFadden (ed.s).

- Reynaert M, Verboven F (2014), “Improving the performance of random coefficients demand models: the role of optimal instruments” *Journal of Econometrics*, vol. 179, no. 1, pp. 83 - 98.
- R. Sutton and A. Barto (1998), Reinforcement Learning: An Introduction, MIT Press, Cambridge MA.
- Tamer (2010; “Partial Identification in Econometrics,” *Annual Review of Economics*, 2, 167-195).
- Wollman, T. (2014): “Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles,” *mimeo*, Harvard University.