

A Helicopter Tour of Empirical Industrial Organization¹.

Ariel Pakes, Harvard University.

April, 2020.

This is the first version of this paper, so comments can still be incorporated and are welcome.

¹A more accurate title might insert “personal” before helicopter, as I will focus on the parts of Industrial Organization I know best; topics that either I or my students have worked on. Probably the most notable omission is the empirical work on auctions. I owe much of my knowledge on the topics covered here to my students and co-authors, as well as two who taught me related material from other fields and have sadly passed away, Gary Chamberlain and Zvi Griliches.

Industrial organization is the study of market responses to environmental and/or policy change. It recognizes that very few, if any, markets are either monopolies (with one firm supplying a product that does not have competitors) or perfectly competitive (where firms have no ability to influence price). Rather virtually all markets are imperfectly competitive; markets where each firm knows its actions affect the payoffs, and hence the actions, of its competitors.

In the 1980's theorists took up the challenge of analyzing responses to environmental and policy changes in markets where each firms' actions affected all firms outcomes, and this lead to a revolution in the way markets were analyzed. The work was based on notions of Nash equilibria (1956); or "rest points" where each agent (say firm or regulator) was doing the best it could given the actions of the other agents. At an equilibrium no agent has an incentive to change its behavior and in markets not in equilibrium at least one agent does have an incentive to change. So it was natural to analyze changes in terms of the equilibria they could generate. Of course this misses the process which brings the market to equilibria, a topic I return to below.

The applied theory that followed focused on simple models where different assumptions were used to generate an understanding of what could happen in imperfectly competitive markets. Though insightful, the different assumptions (on functional forms, the timing of decisions, ...) led to different conclusions, and it was not clear which assumptions were appropriate for any given situation. The empirical work that followed tried to build frameworks that let the data tie down the needed assumptions ².

What I will try to do here is provide an overview of the issues that we have a pretty good handle on, point out areas which are just developing, and mention some of the more glaring pieces we are still missing. When considering these points, keep in mind that the question is not whether our models are "correct" in any absolute sense of the word. For either policy, or for understanding historical events, the analysis should be judged by whether it improves upon the alternatives available to the decision maker (or historian). Similarly, methodological contributions should be judged in terms of their potential for enabling us to improve our analytical abilities.

I will divide the discussion into "static" and "dynamic" analysis. Static analysis considers competition among a fixed set of products given their cost functions. A typical question analyzed in static analysis is how will prices change in response to a change in market conditions (e.g. a merger, tariff, or regulatory change). Dynamics considers the implications of current actions (prices or investments in either physical assets or product development) on future profit opportunities. My discussion of dynamics will deal solely with what I view is the basic problem hampering empirical work on dynamics and recently developed ways of circumventing it. A more complete review is provided in another paper in this volume.

²There were earlier attempts to merge empirical work with theory for some of the primitives of economic models (e.g. Houthakker, 1955, or Gorman, unpublished circa 1960). However, their insights were difficult to generalize and generality was needed to enable the data to determine the appropriate assumptions. The computer revolution and the data, econometric techniques, and computational abilities that accompanied it, changed that.

1 Static Analysis of Demand and Equilibrium.

The most common form of static analysis has firms choosing prices. The assumption that makes the analysis static is that the choices of price today does not have an independent effect on future profit opportunities. This rules out many markets³, but the tools of static analysis are required to analyze markets where prices have dynamic effects, and they are easier to understand in isolation. Static analysis has three components.

1. A demand function which determines the quantities of the products marketed purchased at various prices.
2. A cost function which determines how much it cost to produce those products.
3. A behavioral assumption which determines how firms set prices⁴.

Given these three components, the static framework is able to analyze the impact of counterfactual changes in policy or the environment on prices, profits, and, if the demand system is based on integrating individual demands, on the distribution of consumer surplus.

Demand Functions. Early work on market demand was largely in “product space”. This consisted of estimating demand for each product as a function of its price and the prices of competing products. This had two problems in our context. First, if there were J competing products even a linear system would require on the order of J^2 coefficients and with markets with fifty or more products this was simply too many parameters to estimate with any precision. Second, a demand system estimated in this way could not predict demand for new product introductions, and this was necessary in order to predict changes in the products likely to be marketed in response to environmental changes.

Our solution to these problems was to move to an analysis in “characteristic space”. A product was defined by its characteristics and individual preferences over products were determined by the interaction of product characteristics with household attributes (large families could prefer large cars, poor families could be more sensitive to price, etc.). The empirical analysis estimated the importance of these interactions; see Berry, Levinsohn, and Pakes (1995, 2004) or BLP⁵. The number of parameters that needed to be estimated depended only on the joint distribution of the interaction terms over households. This did not depend on the number of products marketed. Moreover, at

³Static analysis ignores the impact of current prices (or quantities); on future demand (as would occur in markets for durable, experience, or network goods), future costs (as in learning by doing, or input adjustment costs), and/or on future equilibrium conditions (due to collusive incentives or signaling in models with asymmetric information).

⁴Given the demand function, prices determine quantities purchased. An alternative has firms setting quantities with prices adjusting so that those quantities are bought.

⁵The earliest I.O. paper in this vein was Bresnahan (1987), who used a single “quality” dimension to distinguish between automobiles in a telling analysis of collusive periods in the auto industry. Formal analysis of the non-parametric identification properties of these models is reviewed in Berry and Haile (2016). In addition there has been a large amount of subsequent work on improving its computational properties and the choice of instruments.

least in principal, we could predict demand for a product given prices before it was marketed. Of course we will not do as well predicting preferences for new products whose characteristics are quite different from those in our data (e.g. lie outside of the convex hull of the characteristics in our data).

The movement to characteristic space generated two computational problems. To compute market demand we had to sum up over the demands induced by the preferences of many households. The market demand was needed both to determine prices and, when household level demands were not available, to estimate the interactions needed for the demand system from a combination of market level data on quantities prices and product characteristics with publicly available data on the distribution of consumer attributes. Advances in computational abilities and the accompanying development of simulation techniques to approximate sums essentially solved this problem⁶. Second, at least for consumer goods, there are often too many potential characteristics to use. Our solution was to allow for an “unobserved” characteristic which picked up the effect of the residual characteristics that we did not condition on. Since firms priced on all their characteristics we had to allow for a correlation between price and the unobservable. This is the analogue of the “simultaneity” problem in product space demand models, but here the unobservable was inside a complicated non-linear aggregate. BLP introduced a contraction mapping that allowed the researcher to express the unobservable characteristics as a linear function of observables. This enabled us to use standard instrumental variable techniques to account for correlation between the unobserved and observed characteristics.

At least two limitations of the framework are worth pointing out. There are active research programs trying to accommodate each, but no commonly accepted framework for their analysis. First, the basic framework assumes that the researcher knows the products that are in the agents’ choice set, and that households know the characteristics of those products. Second, the framework has difficulty distinguishing between correlations in tastes and causal factors that lead to similar actions. Is the fact that individuals make similar choices over time, or that individuals in a network make similar choices, a result of similarity in tastes that our observables do not control for, or is it driven by causal relationships?

The demand framework outlined above is now part of the toolkit of most Industrial Organization economists, and has spread rapidly to other fields of economics (public finance, health and environmental economics, market design ...) and consultancies, and has provided useful insights for regulatory and data gathering agencies⁷.

Demand and Pricing in Retail Markets. How have we done with our demand and equilibrium assumptions in retail markets? In any given market this is easily checked. If we rely on our demand system and our equilibrium assumption then price

⁶The use of simulation to approximate integrals in economics dates at least to McFadden et al. (1979), who used it for prediction, and Pakes (1986) who embodied simulation techniques in an estimation algorithm. The econometric problems that arose in using these techniques were first analyzed in Mcfadden (1989) and Pakes and Pollard (1989).

⁷My own work has been directed at using the framework to develop ways of alleviating biases in the consumer price index; see Pakes (2003) and Erickson and Pakes (2011).

Table 1: Wollmann & 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.

should equal cost plus a markup. The markup can be obtained separately from the estimates of the demand system and should have a coefficient of one in the pricing equation. We also need a marginal cost function, but provided we are willing to specify marginal cost as a function of product characteristics and the prices of inputs (e.g. wages for labor) the pricing equation can estimate marginal costs.

I asked Tom Wollmann to use the demand system for trucks he estimated (Wollmann *AER*, 2018) to do that analysis. The results were reported in Pakes (2017) and are reviewed in table 1 below. The R^2 is a measure of the fraction of the variance in price explained by the model; it is remarkably high (as good as I have seen for a choice variable in the social sciences) and the coefficient of the markup is well within a standard error of one. We also considered changes in price over time. The only explanatory variable that exhibits noticeable changes over time for a given product is the markup as input price movements are negligible. I.e., a product's price changes are almost solely responses to changes in the set of competing products. Changes in the markup picked up fifty to sixty percent of the variance in product prices over time. Though this is also an unusually large number for analyzing “within” changes in a choice variable in social science settings, it shows that we do less well in analyzing changes in prices over time than analyzing the differences across products at a given point in time, an issue I turn to now.

The use of equilibrium assumptions. Conditional on the characteristics of the goods marketed, their cost functions, and the behavioral assumption setting price, the demand and pricing framework above enables one to solve for the change in prices, profits, and distribution of consumer welfare in retail markets. This generates an ability to analyze the impact of counterfactuals on these same variables. The rest of this section comments on extensions that generate the same capabilities in markets where the assumptions underlying the retail analysis are inappropriate. However, before doing so I want to note a feature of the equilibrium analysis that we have not adequately dealt with and arises in all settings.

Nobody believes that markets react to an environmental change by instantaneously moving to a rest point. Unfortunately we know very little about the adjustment process. An appropriate adjustment model would enable us to investigate “transition dynamics”. However, the need for a model of adjustment is deeper than that. Since models of markets have interacting agents they generate the possibility of multiple equilibria; multiple sets of actions wherein each agent is doing the best it can given the actions of other agents. Multiplicity becomes a near certainty when we allow for the more complex models that endogenize the choice of products to be marketed (sections 2.1 and 3).

Though we may be able to let the data determine which equilibria were played in the past, the multiplicity issue must be resolved to analyze counterfactuals, and counterfactuals underlie most policy implications. An accurate empirical model of how firms learn to adapt to change would be one way of choosing among alternative equilibria. Theorists have written extensively on learning (see Fudenberg and Levine, 1998, and Hart and Mas-Colell, 2013), but there has been less empirical work. A start on the empirical analysis of adjustment in markets was made in Doraszelski, Lewis, and Pakes (2018). They find support for simple learning models in situations where firms have been interacting for some time, but note that when major changes occur firms seem to experiment and we lack a framework for analyzing those situations.

1.1 Extensions of the Static Framework

There has been a significant amount of work extending the analysis to account for additional features that are central to policy. I turn to these before coming back to a deeper look at the analysis of costs and productivity.

Asymmetric Information, Adverse Selection, and Moral Hazard. Implicit in the analysis above was the assumption that the cost of supplying the marketed product was not a function either of; who purchased the product, or of details of the contract enabling the sale. However, different purchasers of health (or car) insurance have different probabilities of illness (or accidents) and will cost the insurer different amounts. Similarly, loans to different people (or firms) have different default rates. In these markets actions by the purchaser after the purchase can cause costs that are not internalized by the purchaser but rather are allocated to the service provider,

generating the problem of “moral hazard”. In addition if the service provider does not price all the purchaser’s characteristics that determine costs “adverse selection” can occur. A limited ability to price cost characteristics can result from the provider not observing them or because regulations prohibit pricing on them. I illustrate issues that these phenomena generate with examples from health insurance markets.

Adverse selection obtains its name from the fact that if insurance is offered at a higher price those who continue to purchase it will be disproportionately sick, since insurance is more valuable to them (for a simple exposition see Einav, Finkelstein, and Cullen, 2010) . In the extreme adverse selection can cause market unraveling (Akerlof, 1970); price increases that lead to cost increases that lead to price increases, ... and eventual market disappearance. Even if the market does exist the fact that cost becomes a function of price generates computational issues as equilibrium prices must now take account of the change in costs induced by price changes, and can imply that the Nash in prices equilibrium described above might not exist (the analysis of the ACA exchanges in Handel et al. 2015, considers alternatives).

There are at least two ways in which moral hazard effects the health insurance market. First, the fact that health insurers pick up most of the costs of consumers accessing health services generates an incentives for consumers to access services when benefits are less than costs. Second, the costs of the treatments health care providers prescribes are also mostly paid by an insurer. The fact that the relationship between symptoms and diseases is uncertain and the provider can be blamed for mistaken diagnosis, accentuates the moral hazard problem in prescribing and there is general agreement that doctor’s choices have large impacts on healthcare costs. Recent empirical work has attempted to evaluate how different reimbursement mechanisms mitigate the resulting moral hazard problems.

Early empirical work evaluated the impact of co-pays, co-insurance, and out of pocket maximums on consumer behavior (see Cutler, McClellan, and Newhouse, 2000). More recent work has investigated various impacts of provider reimbursement mechanisms on the choice among treatment options. The worry that fee for service contracts with providers might induce over-provision has lead to research on the effect of payment mechanisms that are based on ex-ante patient risk (“capitation”), or on a categorization of illnesses (“bundled payments”). The question was whether providing monetary incentives to providers would induce skimping on care. The early results were encouraging. For example, Ho and Pakes (2014) study of women giving birth finds that capitation causes doctors to allocate to less convenient hospitals, but does not cause a decrease in the quality of care. On the other hand Einav et al. (2018) and Eliason, et. al. (2018) both show that the incentives leading to the growth of long term care facilities has not increased the quality of care. So different institutional details generate differentially good outcomes and need to be studied individually.

The studies need to keep in mind that the objective of health care policy is different than that most frequently used in economics (the maximization of consumer, or consumer plus producer, surplus). Our society is “contractarian”; we have instituted laws that guarantee that a citizen who fulfills society’s duties has a right to a minimal

level of goods and services; including a minimal level of healthcare. So we are faced with a market design issue: how do we insure a basic level of health care at minimal cost?

Two final points are worth noting. The contractarian argument for a minimal level of health care when applied to care for communicable diseases (e.g. COVID-19), requires us to extend care to those who are not formally members of the society but interact with society's members. Finally, in both health insurance and loan markets issues related to moral hazard and asymmetric information generate complex contracts. The complexity of these contracts has led to the establishment of public institutions to regulate them, and a need for research into preferred regulatory rules.

Vertical Markets. Many markets are “vertically” connected to other markets. In a typical vertical structure an intermediary contracts with a number of goods or service providers in an “upstream” market and then markets a bundle of contracted goods to consumers in the “downstream” market. Examples include: the cable television market where networks contract with content providers in the upstream market and then re-market a “tier” of content to consumers, the health care industry where insurers contract with health service providers and re-market access to a network of providers as a health insurance policy, and large retailers who purchase goods from competing manufacturers and re-market them in their outlets.


The difficulty in analyzing the impacts of changes in these settings is that there are often a small number of buyers and sellers in the upstream market and the allocation of profits between them is determined by bargained contracts rather than by Nash pricing. Moreover, any change in one side of the market will affect the other, so a solution to the contracting game is needed to analyze changes in the retail market. Crawford and Yurukoglu (2012)'s analysis of the impact of a la carte (rather than tier) pricing of television content paved the way for empirical models incorporating upstream bargaining.

In the Nash (1953) bargaining model each side's leverage is determined by the difference in what it would earn with and without the contract. The profits, i.e. the difference between the downstream firm's revenue from marketing the good and the upstream supplier's costs, are split by a rule which maximizes a product of these differences. So to calculate the split we need to calculate a counterfactual; i.e. if two firms are currently contracting we need an idea of what would happen were the contract broken. Would the firm change its other contracts, and would the change in the contract between these two firms induce changes in the contracts of its (upstream and downstream) competitors? Note that the way the profits are split will determine investment incentives, a topic both largely absent from detailed empirical modeling of vertical markets and thought to be central to determining the desirability of vertical merger activity (see chapter 4 of Whinston, 2006). Though the appropriate solution to the problem of constructing counterfactuals depends on the institutional setting, an example will show that some way of accounting for vertical interactions is crucial when considering policy.

Regional television rights for Major League Baseball (MLB) are allocated to the teams (figure 1 has the regions), but to watch a team from out of your region you must purchase the out of market bundle (the OMB) which gives purchasers access to all games. So a New Yorker living in Florida could not watch the Yankees without purchasing the OMB. A class action law suit was brought against MLB characterizing this agreement to be a conspiracy to divide the market. The claim was that consumers would be better off if teams were allowed to market their telecasts in all regions. Figure 2 illustrates the two ways of analyzing this issue that were argued in litigation.

The first treats the market as a standard retail market. The team's costs are its telecast cost and each team markets their product in all regions, choosing each region's price to maximize their profits in that region. The league markets the OMB. This violates the institutions that govern this market. In reality the teams would have to bargain contracts with a cable or satellite network that controls consumers' access to the telecasts. The costs to these distributors would include the markup over costs that results from the bargained contracts and keeps the teams afloat. The distributor would then re-market both the individual team channels and the OMB to consumers, internalizing the fact that when they increase a price of a given team some of the lost viewers would go to the OMB or to other teams. This plus the increase in cost would increase prices consumers would have to pay for either the teams' channels or the OMB bundle above what it would be were each team able to set prices individually. Moreover, some of the teams would now make more money by opting out of the OMB, which would harm consumers further. That is the predictions from a standard retail market analysis were unlikely to be accurate (figure 3 summarizes), and would likely have lead to the wrong outcome.

MLB Home Television Territories



- ❖ League defines each team's "home television territory."
- ❖ In-market telecast rights are assigned to the teams.
- ❖ Out-of-market ("OOM") telecast rights are held by the League.
- ❖ OOM telecasts are sold as part of an OOM Bundle ("OMB").



	Plaintiff's Assumptions	Market Reality
Cost to Marketer	Cost = telecast costs	RSN/MVPD negotiation sets cost Cost = telecast cost + RSN markup
Who Sets Prices	Each RSN sets the price to consumers for its own product	MVPDs set prices to consumers for all products it markets (includes all RSN's)
Pricing Incentives	Each RSN sets its price independently to maximize its own revenue minus the telecast cost	An MVPD sets prices to maximize revenue from all products it markets minus the sum of their fees

Assumptions vs. Likely Reality

❖ Plaintiff's BFW pricing assumptions.

- MVPDs face higher cost than Plaintiff assumes.
- MVPDs prices multiple products, but Plaintiff assumes RSNs set prices “independently.”

Implication: Consumers face higher prices than Plaintiff assumes.

❖ Plaintiff did not check equilibrium conditions.

- Plaintiff assumes there still exists an OMB including all teams. However, the Yankees would make more money if they drop out and increase price.

Implication: Both Yankee and OMB consumers lose.

Markets that Do Not Use Price to Allocate. In markets where prices are used, allocations are determined by tastes, the prices of the goods marketed, and endowments. However, there are many circumstances in which society (or a relevant subset like a sports league) does not want to allow monetary endowments to overly impact allocations. Examples include the allocation of: students to public schools, of kidneys to patients, of doctors to residency programs⁸, and the drafts in professional sport leagues. To analyze the impact of environmental changes in these settings we still need a demand system and a cost or production function, but the pricing equation is replaced by an allocation rule.

The example I briefly consider is our society’s notion that students should have equal opportunity to access their preferred schools (at least for the schools in the public school system). Students submit a preference ordering for schools and a centralized mechanism allocates. The analysis of a counterfactual, be it a change in the allocation rule or in the number or types of schools, would require a demand system estimated from the submitted preferences, the allocation rule, and an outcome equation which maps the student allocation and the school infrastructure into outcomes of interest.

Economists have successfully convinced many school systems to use allocation mechanisms which induce students to report their true preferences (that are “incentive compatible”) and do not generate “justifiable envy” (given schools’ priorities, there is not a student who would prefer another school to their assignment when a lower priority student has a place at the desired school); the leading example being the Gale Shapley (1962) deferred acceptance algorithm. The new allocation rules are generally viewed as leveling the playing field between more and less informed parents (Pathak and Sonmez, 2008).

Abdulkadiroglu et al. (2017 AER) and Agarwal and Somaini (2018, ECMA) compare the allocations that the new and old mechanisms generate. This requires estimation of the utility functions necessary to form demand systems. Assuming the submitted preference orderings were truthful, the data available for demand analysis was exceptional; BLP (2004) showed that having second, as well as first, choice data greatly improved the ability to determine differences in preferences that are not related to observables, and the submitted preference lists are an improvement on that. To see why consider a family who has no parent at home when school is let out, and therefore wants to send the child to a school near a relative’s house. The fact that the first few submitted preferences were to schools in the relative’s neighborhood would allow both the algorithm and empirical work to take account of that preference even though it is not explicitly conveyed.

The demand system and the fact that there is an algorithm to calculate allocations conditional on demand generate exceptionally accurate estimates of two of the three elements needed to analyze the impact of environmental change. The missing element is a production function. To the extent that we want to analyze either the impact of the new allocation rules or how investments in the school system affect traditional measures of performance (graduation rates, subsequent wages,...) we will also need the

⁸In this case the worry was misaligned timing incentives, see Roth (1984).

production function. More generally detailed empirical analysis of markets that do not use prices to allocate has, perhaps understandably, begun by focusing on comparisons of the allocations from different allocation rules. A next step would be to integrate an analysis of production enabling a deeper look at the outcomes those mechanisms generate.

2 Production and Cost Functions.

In addition to being a primitive used in the analysis of market allocations, production functions (or their “dual” cost functions⁹) are used to analyze two specific (and related) issues in empirical I.O.: the analysis of the efficiency of the distribution of output among firms, and the analysis of productivity. Productivity is the ratio of output to an index of inputs where the index is formed from estimates of production function coefficients. It is central to the analysis of infrastructure projects and other activities that might generate benefits or costs to society that are not internalized by the firms performing the activity (research and development being an important example). For reasons I discuss below, estimates of both production functions and of fixed costs are now also being used to enhance the analysis markups over costs considered in section (1).

Productivity is not directly observed but rather is constructed using the production function coefficients, i.e. the parameters which determine the inputs used to produce given amounts of output. These parameters need to be estimated and estimation is complicated by the fact that payoff maximizing firms’ decisions are determined, in part, by the firm’s productivity. The solution, whose modern version was initiated by Olley and Pakes (1996)¹⁰, was to model the relationship between these decisions and productivity. Olley and Pakes (henceforth OP) dealt with establishment level data and endogenized input and exit choices. In other contexts the analysis needs to account for the relationship between productivity and merger decisions, and/or the impact of productivity on the decision to privatize state owned companies.

OP analyzed the impact of the breakup of AT&T on the telecommunications equipment industry. AT&T had monopolized the provision of telephone services as a result of their control over the telephone lines and purchased almost all of its equipment from its wholly owned subsidiary, Western Electric. In 1982 Judge Green split AT&T into seven regional operating companies and made it illegal for any of them to own an equipment manufacturer. This induced entry by a set of foreign producers (Ericson, Northern Telecom, Hitachi,...) and a major retrenchment in Western Electric.

Table (2) provides an indication of why it is important to account for input and entry/exit decisions when analyzing establishment level production data. The first two columns present estimates from two traditional “balanced panel” estimation algo-

⁹Cost functions are derived by minimizing costs subject to the production function and a requirement that a given quantity of output is produced. Recent research has focused on estimating production functions rather than cost functions for several reasons; not least among them is that cost data is typically harder to access. The exception is regulated industries in which the regulator can require firms to submit costs.

¹⁰Marschak and Andrews’ (1944) stellar paper predates the modern literature

rithms: these only use the establishments that are active throughout the years covered by the panel. The within estimates assume that the productivity of a plant is constant over time which, ignoring problems induced by exit, implies that unbiased coefficient estimates can be obtained by regressing the changes in output on the changes in inputs. The total column is ordinary least squares. The third and fourth columns present the within and total estimates from a data set which keeps observations on establishments in all years they are active. The fifth column presents the estimates from using a model of input and exit decisions to control for the relationship between productivity and those decisions.

The balanced panel uses about one third of the observations available; only 40% of the establishments that were active in the initial year of the study (1972) were still active in the final year (1987), and 80% of the those active in 1987 were not active in 1972. Research is often interested in the impacts of major changes in the environment; such changes are often accompanied by a lot of entry and exit. Those who exit tended to be plants whose productivity fell as a result of the breakup, and the larger the plant the more the productivity had to fall before they exited. This induces a negative correlation between productivity and the capital of surviving firms and explains the low capital coefficients in columns (1) and (2); a finding at odds with the fact that this is a very capital intensive industry. Relatedly the firms whose productivity did grow increased their labor force causing a conflation of the impact of labor on output and the impact of productivity on labor demand, and upwardly biasing the labor coefficients. The last column uses all the available data and a model of input and exit choices to correct for the “selection” and “endogeneity” problems. Compared to the within balanced panel estimates the labor coefficient falls by about 30% and the capital coefficient almost doubles. The comparison of column (5) to the simpler full sample estimates in (3) and (4) reveals a labor coefficient which is similar to the within coefficient but a capital coefficient that is much closer to the total estimates.

Given the production function coefficients we can analyze the impact of the changes on the efficiency of the output allocation. Industry productivity is a share weighted average of the productivity of the plants within the industry (formally if $s_{i,t}$ is the share of output and $p_{i,t}$ is the productivity of plant i , then the productivity of the industry is $p_t = \sum_i s_{i,t} p_{i,t}$). Olley and Pakes introduce a decomposition of industry wide productivity growth into a part attributable to an increase in the average productivity of the establishments ($\bar{p}_t$) and a part due to the reallocation of output across establishment $\sum s_{i,t}(p_{i,t} - \bar{p}_t)$. The results are in the next table.

Alternative Estimates of Production Function Parameters^a
In The Telecommunications Equipment Industry.
(Standard Errors in Parentheses)

Sample:	Balanced Panel		Full Sample		
	(1)	(2)	(3)	(4)	(5)
Estimation Procedure	Total	Within	Total	Within	Model Based
Labor	.851 (.039)	.728 (.049)	.693 (.019)	.629 (.026)	.608 (.027)
Capital	.173 (.034)	.067 (.049)	.304 (.018)	.150 (.026)	.355 (.058)
Age	.002 (.003)	-.006 (.016)	-.0046 (.0026)	-.008 (.017)	.010 (.013)
Time	.024 (.006)	.042 (.017)	.016 (.004)	.026 (.017)	.020 (.046)
Other Variables	—	—	—	—	Kernel in P and h
Observations ^a	896	896	2592	2592	1758

a. The dependent variable is the log of value added. Value added production function assume a constant ratio of materials per unit of output. The number of observations is smaller in column 5 because the model requires the researcher to drop the first observation on each plant.

A “registration and certification” program which started in 1977 allowed competing firms to attach their devices to the public network provided they satisfied safety concerns. The table shows that the program was associated with an increase in the efficiency of the output allocation between 1978 and 1982. The disruption from the breakup seems to have caused a fall in both the average productivity and the efficiency of the allocation between 1982 and 1984. However, the allocation improves again thereafter. OP go on to show that the productivity gains were largely a result of a reallocation of capital to more efficient establishments (see the correlation between productivity and capital in column 4); capital was misallocated in the regulated period.

Decomposition of Productivity Growth^a 1975-87.

Year	p_t	$\bar{p}_t$	$\Sigma_i s_{it}(p_{it} - \bar{p}_t)$	$\rho(p_{i,t}, k_{i,t})$
1975	0.72	0.66	0.06	-0.11
1976	0.77	0.69	0.07	-0.12
1977	0.75	0.72	0.03	-0.09
1978	0.92	0.80	0.12	-0.05
1979	0.95	0.84	0.12	-0.05
1980	1.12	0.84	0.28	-0.02
1981	1.11	0.76	0.35	0.02
1982	1.08	0.77	0.31	-0.01
1983	0.84	0.76	0.08	-0.07
1984	0.90	0.83	0.07	-0.09
1985	0.99	0.72	0.26	0.02
1986	0.92	0.72	0.20	0.03
1987	0.97	0.66	0.32	0.10

^a $\rho(p_{i,t}, k_{i,t})$ is the correlation between estimated productivity and capital. The rest of the variables are defined in the preceding text.

Revenue versus Quantity Functions. The last table does contain an apparent anomaly. The equipment industry was an industry with rapidly advancing technology, but the table indicates no trend in average productivity. OP mention that their output measure was based on revenue, so their productivity findings should be interpreted as changes in revenue per unit of inputs rather than in output per unit of inputs. The incentives underlying the corrections discussed above are similar when we interpret their procedure as one of estimating a revenue generating function, and the determinants of revenue per unit input is of obvious importance in determining likely responses to change. However, the difference between revenue and quantity generating functions becomes important when we want to distinguish between the implications of the change on quantity and price, as we often do when we want to analyze welfare. In the AT&T case a fall in prices generated by increased competition after the breakup could easily generate flat revenue productivity despite technological advances.

There have been attempts to separate out price from quantity movements using only data on revenue and inputs. This would have the advantage that it does not require separate demand estimates for each market that the analysis of section (1)

uses to separate markups from costs. As a result, it could be done on a broad sample of industries using either publicly available data on traded firms, or data from Census Bureaus. De Loecker and Warzynski (2012) were the first to provide a set of assumptions that allowed for markup estimation from production data. They require an estimate of the output elasticity of a variable factor of production. This is not what we estimate when analyzing the relationship between revenue and total inputs in the multi-product firms that dominate the data. Still their and subsequent research developed increasingly rich ways of analyzing the relationship between inputs and revenue productivity and characterized how that relationship changed over time and across industries (see Raval, 2019 and Demirer, 2020). To get back to markups we have to incorporate the fact that multi-product firms do not have a production function, but rather a “production possibility frontier”; i.e. for a given set of inputs different combinations of output can be produced.

There have been attempts to analyze production possibility frontiers but they have been hampered by the lack of data on the separate outputs the firm produced and the inputs they used. The good news is the recent availability of disaggregated firm level output data (see the references in Dhyne et al. 2020). Inputs disaggregated by use are typically not available and may well be ill-defined as often many inputs are simultaneously used in the production of several outputs. It is possible to estimate the production possibility frontier from disaggregated product level quantities and aggregated input quantities (see Pakes, 2020), though to do so in a reliable way will require modifying current estimation algorithms¹¹.

This would facilitate the analysis of the factors underlying the observed secular fall in labor share, or equivalently the increase in returns to capital and intangibles, and its relationship to changes in markups. De Loecker, et. al. (2020) investigate the trends in these returns over time, and show that the observed increase is largely a result of its increase in the largest ten percent of firms, and suggest that it might be related to increases in markups. If so there remains the question of what generated the markup increase. Possible reasons range from changes in the nature of antitrust enforcement (Wollmann, 2019, notes specific rule changes, while Philippon, 2019, focuses more generally on the culture of the antitrust agencies), to an increase in the markups needed to support the fixed and sunk costs that go into the development of new goods and services. If the latter it may well be required for the large increases in consumer surplus that result from those investments. Both possibilities require further investigation; i.e., we need an analysis of the relationship between product specific markups and antitrust activity, and a deeper understanding of how new products and processes are developed and positioned in the economy.

¹¹In particular we will need a way of controlling for the endogeneity of inputs when multiple outputs are produced by the firm, and, in addition, we need functional forms that differ in their convexity for inputs and outputs.

2.1 Fixed Costs and Product Repositioning.

We have not considered the estimation of fixed costs and those estimates are needed in order to analyze product repositioning. Product repositioning refers to a given set of competitors changing the characteristics of the products they market. In industries where this can be done relatively quickly, say as quickly as prices can be changed, an analysis of product repositioning is required to accurately assess the impacts of a policy or environmental change on prices, even in the very short run. Typically repositioning is analyzed in two stages, a stage where product characteristics are chosen followed by a stage where prices are chosen. Consistency between the stages is insured by the requirements of backward induction (or “subgame perfection”)¹².

A good illustration of just how quickly products can be repositioned is given in Chris Nosko’s thesis (2014). He looks at the response of both Intel and its major competitor, AMD, to Intel’s introduction of the Core 2 duo chip in July 2006. Figure 4 shows that in the month prior to the introduction, there was intense competition for high performance chips, with AMD selling the highest priced product at just over \$1000, and seven chips selling at prices between \$600 and \$1000. Figure 5 shows that in October, three months after the introduction of the Core 2 Duo, AMD no longer markets any high priced chips, there are no chips between \$600 and \$1000 marketed, and Intel markets all chips priced above \$1000. Nosko goes on to explain that 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.

The analytic framework for analyzing product repositioning is relatively recent but has been used to analyze the impacts of the truck bailout in the great recession (cabs can be interchanged with trailers easily; Wollmann, 2018), the impacts of withdrawing low-end computers (Eisenberg, 2014), and the cost of negotiating contracts (Ho, 2009, Crawford and Yurukoglu, 2012). In industries where repositioning is relatively easy it should be considered by the regulatory agencies in their central role of evaluating the likely impacts of mergers on prices. Given an ability to analyze demand, prices, and profits (see section 1), the analysis only requires estimates of the fixed costs associated with introducing/withdrawing a product (or a contract), and a method of choosing among counterfactual equilibria.


The ability to estimate fixed costs relies on being able to compute counterfactual profits for a set products. Then we can compute the profits that would have been earned were a product that was introduced not marketed. The increment in profits from introducing the product should be expected to be greater than the fixed cost

¹²In a formal discrete time model product repositioning can only be analyzed in this way if the characteristics chosen in the current period do not impact the profits earned from the products that might be chosen in future periods. When the cost of changing characteristics is small, empirical models often ignore the impacts of repositioning on future profits. Two-period models have also been used in empirical work that analyzes choices where longer term considerations are clearly central; the leading example being two-period entry models. The classic two-period empirical entry models assume there was no history before the initial period and there will be no future after the second. As a result, it should be viewed as a reduced form which organizes historical data in telling ways; their relationship to dynamic models is clarified in Pakes (2014). This limits the usefulness of two-period models for counterfactual analysis, but at least until the recent developments in dynamic modeling, there were few other tools available for analyzing entry and exit.


associated with the introduction conditional on the information the firm had at its disposal when the decision to introduce was made. Similarly, if a product that could have been marketed was not, the increment in profits that could have been earned should be expected to be less than the fixed costs. So the average of these increments should provide upper and lower bounds to the average fixed costs (and a slightly more complicated algorithm could estimate these bounds conditional on measured product characteristics).

Pakes et al. (2012) provide conditions on information sets and the profit estimates which insure that these inequalities can be used for both estimation and inference. However the use of bounds in modern economic analysis dates to the work of Manski (see his 2003 book) and the literature on the distribution of estimated bounds (more accurately of the set of acceptable values) dates to Imbens and Manski (2004) and Chernozhukov et al. (2007). The needed distributions do need to be simulated, but modern computers and recent computational suggestions make this feasible (especially if the fixed costs are linear in observables, see Andrews et al. (2020), and the literature cited there).


The remaining issue is how to specify the counterfactual equilibrium given a value for fixed costs. Though multiple equilibria is not generally thought to be a problem when analyzing prices, it is typically unavoidable when analyzing repositioning, since where one agent places its products has a discontinuous effect on the profitability of the placements of its competitors. The counterfactuals analyzed in applied work have used learning algorithms (Wollmann, 2018) or enumeration of outcomes that seem plausible (Eisenberg, 2014), but as noted in section (1) and discussed further below, more empirical work on equilibrium selection is needed.




Price/Performance – June 2006



Price/Performance – July 2006



Price/Performance – Oct 2006



Price/Performance – January 2008

3 Dynamics.

The applied work on dynamics started in a manner analogous to the way we started empirical work on static models. We took the Markov Perfect paradigm used by our theory colleagues, in particular Maskin and Tirole (1987a and b) who showed how it could be used to understand I.O. phenomena in simple settings, and provided a framework which had the ability to incorporate the details that seemed essential to understanding individual markets (see Ericson and Pakes, 1995). That framework made assumptions that insured that; (i) the state variables that determined agents' choice of controls evolved as a Markov process, and (ii) the equilibrium was some form of Markov perfection (i.e., controls are chosen so that no agent has an incentive to deviate given their perceptions of profit opportunities, and those perceptions are consistent with the actions of their competitors and exogenous events).

The framework generates a Markov process on the underlying states of the system. So provided the number of those states is finite, the state of a system will in finite time wander into a recurrent class of states, and once in that class will stay within it forever. The recurrent class consists of a subset of the states that the economics of the problem imply can be visited repeatedly; other states are transient, they may be visited while approaching the recurrent class, but the market can not support them in any lasting way.

The framework led to computational explorations of theoretical issues which were difficult to study analytically, but only a limited amount of empirical work¹³. The basic problem was that once the researcher included the variables that seemed because of their competitive impact on demand, supply, and/or investment conditions, an assumption of full information generated a state space that was too large to be manageable. An alternative was to assume that each firm only had access to a subset of those variables and look for a Bayesian perfect equilibrium, but then the complexity of the calculations required to obtain optimal policies increased dramatically (posteriors needed to be calculated, and they had to be consistent with optimal policies).

More generally once we took into account the cognitive and memory requirements for computing the policies generated by the standard models it was both; (i) difficult for the researcher to use them, and more importantly, (ii) it was hard to believe that these models were the best way to approximate how management made decisions. The first attempt in Industrial Organization to attack this problem was by Benkard et

¹³Computational examples include Besanko et al. (2010) who study learning by doing with forgetting, Fershtman and Pakes (2000) who study collusion in a dynamic game with investment and entry, and Gowrisankaran (1999) and Mermelstein et al. (2020), both of whom endogenize merger decisions. There had been earlier empirical work using different frameworks; for example Porter (1983) studied collusion using the influential repeated game framework in Green and Porter (1984). For an early empirical piece which uses the Ericson Pakes framework and incorporates the institutional detail that seems necessary for policy analysis see Benkard (2004). There has also been work on formulating market demand functions when consumers choices are explicitly dynamic, see Hendel and Nevo (2006) and the literature they cite. Somewhat dated reviews of dynamic empirical work on markets is available in Akerberg et al. (2007), and of computational work is in Doraszelski and Pakes (2007).

al. (2008) who introduced the notion of “Oblivious Equilibrium” as a computational approximation to full information Markov perfect models (much as “mean field theory” in physics approximates the behavior of high dimensional objects by simple averages). The approximation assumed that no matter the current state, the distribution of future states used in calculating continuation values was the long-run average of the probabilities of those states appearing. They show that this can greatly simplify the computation of policies and provide examples where it approximates Markov perfect policies well¹⁴.

Fershtman and Pakes (2012) take a behavioral approach to the problem and introduce the notion of experience based equilibrium (EBE). EBE allows different firms to rely on different information sets when making their decisions. This either because of asymmetric information, or because cognitive constraints limit management’s ability to associate different policies with too fine a partition of the state space. The state space that governs the evolution for the market as a whole can still be quite large, consisting of the set of information sets of the different participants, but individual decisions only depend on their own information.

The equilibrium has two requirements; (i) agents choose actions which maximizes their perceptions of their discounted returns, (ii) those perceptions are at least consistent with what the agent observes. These conditions imply that for states that are visited repeatedly, that is for those in the recurrent class, perceptions will eventually be consistent with observed outcomes, and perceived discounted values from the actions taken in those states will equal the actual expected discounted values resulting from those actions.

Fershtman and Pakes also provide an algorithm for computing equilibria. The algorithm has an interpretation as a learning algorithm; that is it shows how to use past outcomes to update the agent’s perceptions of the expected discounted values from making different choices, assuming only that the agent always chooses the action that maximizes its perceived discounted value. When the algorithm’s learning process is an adequate approximation to the way management learns, it can plot likely transitions from one environment to another. That is, it provides a transition path *and* selects out an equilibrium (or if the transition path is simulated repeatedly it generates a distribution of likely equilibria).

The learning process is “asynchronous”; an iteration starts at a particular state of the market and proceeds by moving to another state and then updating its estimates of discounted values from feasible actions at the initial state. The updating is similar to reinforcement learning; it keeps in memory an estimate of the discounted values from each possible action at each of the individual’s states, simulates individual outcomes from the algorithm’s current state, and then updates the initial discounted values for the particular state the algorithm is in by a weighted average of the simulated outcome evaluated at the initial estimates of the discounted values and the initial estimates. The simulated outcomes are also used to update the state of the algorithm, so the

¹⁴For more on the underlying problem, see Pakes (2016), and for extensions to the Oblivious Equilibrium concept see Ifrach and Weintraub, (2017).

algorithm mimics reality in that it only learns from the outcomes from one state in each “period” (i.e. iteration).

EBE is similar in spirit to the theoretical literature on self-confirming equilibrium (Fudenberg and Levine, 1993). As in that literature “off the equilibrium path” behavior is unconstrained. In the Markov environment considered in EBE this implies the possibility of a different source of multiplicity than those already apparent in Markov or Bayesian perfect equilibria of dynamic games (both of which satisfy the conditions of an EBE). We classify the points in the recurrent class as interior or boundary points. At an interior point all feasible actions always result in outcomes within the recurrent class. So in equilibrium each agent’s perceptions of the values generated by all possible actions are objectively correct. Boundary points have feasible (though perceived to be non-optimal) actions which take them out of the recurrent class. Beliefs about outcomes outside of the recurrent class are not constrained by the equilibrium notion, and different sets of beliefs about the value of outcomes not in the recurrent class can generate different recurrent classes. Asker et al. (forthcoming) provide a testable refinement (“boundary consistency”) which restricts the equilibria by insuring that beliefs on values outside of the recurrent class are consistent with the history prior to entering the recurrent class and initial beliefs.

The computational structure of the algorithm builds on Pakes and McGuire (2001). The update of continuation values only requires adding two numbers; it does not require computing an integral over possible future states¹⁵. Also, the algorithm wanders into a recurrent subset of all possible states and stays within it. The cardinality of that subset depends on the economics of the problem, and need not increase in any particular way with the number of state variables. These two properties of the algorithm, together with a stochastic test of convergence introduced in the Fershtman Pakes article, imply that the algorithm’s computational burden does not necessarily increase exponentially (or even geometrically) with the number of state variables (the algorithm does not have a “curse of dimensionality”). In our calculations the number of states in the recurrent class tends to grow linearly in the dimension of the state space, though the number of iterations before convergence can be large and ways of mitigating this would be helpful.

With appropriate choice of agents’ information sets several previous empirical papers can be characterized as being based on an EBE. A recent and important example is Brancaccio et al. (2020). They endogenize transportation costs in a model of international trade of products transported by bulk shipping, and then show how various changes in the environment are likely to effect trade costs and hence trade. Boats arriving in a given location are either matched with a shipper or left to chose whether and where to travel empty to. The boat’s decision is modeled as a dynamic program where the boat owner is assumed to only know the average value of each location over a long time period and the associated transport costs; i.e. the owners do not know the distribution of values consistent with the current state of the system. They solve for the “steady state” equilibrium distribution of values and transport costs under these

¹⁵This dates to Robbins and Monro’s (1956) use of stochastic integration.

assumptions and alternative values of the parameter vector, average the data at the different ports over time, and estimate the values of the parameter from the observed cross-sectional distribution of boats and their movements.

In the Brancaccio et al. paper there was a fairly clear notion of what reasonable information sets might be, and they did not vary among participants. For many problems neither of these conditions will be met. Then information sets will have to be determined empirically by investigating the relationship between the controls the firm chooses and available data. Moreover that analysis is likely to generate serially correlated residuals. I.e., the researcher may not have data on all the determinants of management's decisions and those not observed, like most economic variables, are likely to be serially correlated. So empirical work will likely require estimation and computation algorithms that allow for serially correlated unobservables. Berry and Comapaini (2019) provide a method for estimating the parameters of a dynamic game when there are serially correlated errors. They assume "full information"¹⁶, but adapting their method to accommodate asymmetric information should be possible.

After a long hiatus methodological work on dynamics has been able to provide procedures which circumvent many of the problems that plagued earlier empirical work. The question that remains and is being explored in current research is whether these developments will lead to analogues of the rich and informative set of empirical studies that followed the developments in static models.

4 Conclusion.

The testimony to the accomplishment of empirical Industrial Organization, and its value to both policy and the understanding of historical events, is embodied in the wealth of empirical work on separate industries increasingly available. I have not tried to review those accomplishments here, focusing instead on methodology with a few cites to early work and periodic illustrations of usefulness. However, the depth of understanding that the field has been able to achieve cannot be truly appreciated without accessing the individual industry studies.

My goal here was to give colleagues from other fields an overview of what empirical I.O. can currently do and the problems still facing us. To the extent that we have made progress, it was a result of thoughtful attempts to understand the way markets work, and realizing that to embody the required detail into our models required further methodology; methodology often adapted from our theory or econometric colleagues. So the individual empirical studies also underlie virtually all of the developments in the field.

¹⁶That is each agent knows all determinants of all other agents' decisions except for a serially uncorrelated error, it is just the researcher who does not know the state of the system.

References

1. Abdulkadiroglu, A., N. Agarwal, and P. Pathak. 2017. “The Welfare Effects of Coordinated Assignment: Evidence from the NYC HS Match. *American Economic Review* 107(12): 3635-3689.
2. Akerberg D., L. Benkard, S. Berry, and A. Pakes. 2007. “Econometric Tools for Analyzing Market Outcomes.” In J. Heckman and E. Leamer (Eds.) *The Handbook of Econometrics* (pp. 4171-4276). Amsterdam: North-Holland.
3. Agarwal, N. and P. Somaini. 2018. “Demand Analysis using Strategic Reports: An Application to a School Choice Mechanism.” *Econometrica* 86(2): 391-444.
4. Akerlof, George A. (1970). ”The Market for ’Lemons’: Quality Uncertainty and the Market Mechanism.” *Quarterly Journal of Economics* 84 (3): 488-500.
5. Andrews I., J. Roth, and A. Pakes. 2019. “Inference for Linear Conditional Moment Inequalities.” Harvard working paper.
6. Asker, J., C. Fershtman, J. Jeon, and A. Pakes. Forthcoming. “A Computational Framework for Analyzing Dynamic Procurement Auctions: The Market Impact of Information Sharing.” *RAND Journal of Economics*.
7. Benkard, L., 2004. “A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft.” *Review of Economic Studies* 71(3): 581-611.
8. Benkard, L., B. Van Roy, and G. Weintraub. 2008. “Markov Perfect Industry Dynamics with Many Firms.” *Econometrica* 76(6): 1375-1411.
9. Berry, S., J. Levinsohn, and A. Pakes. 1995. “Automobile Prices in Market Equilibrium.” *Econometrica* 63(4): 841-890.
10. 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* 112(1): 68-105.
11. Berry, S. and P. Haile. 2016. “Identification in Differentiated Product Models.” *Annual Review of Economics*, 8: 27-52.
12. Berry, S. and G. Compiani. 2018. “An Instrumental Variables Approach to Dynamic Models.” Yale working paper.
13. Besanko, D., U. Doraszelski, and Y. Kryukov. 2014. “The Economics of Predation: What Drives Pricing When There is Learning-By-Doing?” *American Economic Review* 104(3): 868-897.
14. Brancaccio, G., M. Kalouptsi, and T. Papageorgiou. Forthcoming. “Geography, Transportation, and Endogenous Trade Costs. *Econometrica*.

15. Bresnahan T. 1987. "Competition and Collusion in the American Automobile Market: The 1955 Price War." *Journal of Industrial Economics* XXXV(4): 457-482.
16. Chernozhukov, V., H. Hong, and E. Tamer. 2007. "Estimation and Inference on Parameter Sets in Econometric Models." *Econometrica* 75(5): 1243-1284.
17. Crawford, G. and A. Yurukoglu. 2012. "The Welfare Effects of Bundling in Multichannel Television Markets." *American Economic Review* 102(2): 643-85.
18. Cutler, D., M. McClellan, and J. Newhouse. 2000. "How Does Managed Care Do It?" *RAND Journal of Economics* 31(3): 526-548.
19. De Loecker, J., J. Eeckhout, and G. Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *The Quarterly Journal of Economics* 135(2): 561-644.
20. De Loecker, J. and F. Warzynski. 2012. "Markups and Firm-Level Export Status." *American Economic Review* 102(6): 2437-2471.
21. Demirer, M. 2020. "Production Function Estimation with Factor-Augmenting Technology: An Application to Markups." MIT working paper.
22. Doraszelski, U., and A. Pakes. 2007. "A Framework for Applied Dynamic Analysis in I.O., Vol. 3, Chapter 30, 1889-1966, in M. Armstrong and R. Porter ed.s. The Handbook of Industrial Organization.
23. Doraszelski, U., G. Lewis, and A. Pakes. 2018. "Just Starting Out: Learning and Equilibrium in a New Market." *American Economic Review* 108(3): 565-615.
24. Dhyne, E., A. Petrin, V. Smeets, and Frederic Warzynski. 2020. "Theory for Extending Single-Product Production Function Estimation to Multi-Product Settings." Aarhus University working paper.
25. Einav, L., A. Finkelstein, and M. Cullen. 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices." *Quarterly Journal of Economics* 125(3): 877-921.
26. Einav, L., A. Finkelstein, and N. Mahoney. 2018. "Provider Incentives and Healthcare Costs: Evidence from Long-Term Care Hospitals." *Econometrica* 86(6): 2161-2219.
27. Eizenberg, A. 2014. "Upstream Innovation and Product Variety in the United States Home PC Market." *Review of Economic Studies* 81(3): 1003-1045.
28. Eliason, P., P. Grieco, R. McDevitt, and J. Roberts. 2018. "Strategic Patient Discharge: The Case of Long-Term Care Hospitals." *American Economic Review* 108(11): 3232-3265.

29. Ericson, R. and A. Pakes. 1995. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies* 62(1): 53-82.
30. Erickson, T. and Ariel Pakes. 2011. "An Experimental Component Index for the CPI: From Annual Computer Data to Monthly Data on Other Goods." *American Economic Review* 101(5): 1707-1738.
31. Fershtman, C. and A. Pakes. 2000. "A Dynamic Oligopoly with Collusion and Price Wars." *RAND Journal of Economics* 31(2): 207-236.
32. Fershtman, C. and A. Pakes. 2012. "Dynamic Games with Asymmetric Information: A Framework for Empirical Work." *Quarterly Journal of Economics* 127(4): 1611-1661.
33. Fudenberg, D. and Levine, D.K. 1993. "Self Confirming Equilibrium." *Econometrica* 61(3): 523-545.
34. Gale, D. and L.S. Shapley. 1962. "College Admissions and the Stability of Marriage." *American Mathematical Monthly* 69(1): 915.
35. Gorman, W.M. 1980. "The Demand for Related Goods: A Possible Procedure for Analysing Quality Differentials in the Egg Market." *Review of Economic Studies* 47: 843- 856. First circulated as Journal Paper No. 2319, Iowa Agricultural Experiment Station, November 1956.
36. Gowrisankaran, G. 1999. "A Dynamic Model of Endogenous Horizontal Mergers." *RAND Journal of Economics* 30(1): 56 83.
37. Green, E., and R. Porter. 1984. "Noncooperative Collusion under Imperfect Price Information." *Econometrica* 5(1): 87-100.
38. Handel, B., I. Hendel, and M. Whinston. 2015. "Equilibria in Health Exchanges: Adverse Selection vs. Reclassification Risk." *Econometrica* 83(4): 1261-1313.
39. Hendel, I. and A. Nevo. 2006. "Measuring the Implications of Sales and Consumer Inventory Behavior." *Econometrica* 74(6): 1637-1673.
40. Ho, K. 2009. "Insurer-Provider Networks in the Medical Care Market." *American Economic Review* 99(1): 393-430.
41. Ho, K. and A. Pakes. 2014. "Hospital Choices, Hospital Prices, and Financial Incentives to Physicians." *American Economic Review* 104(12): 3841-3884.
42. Houthakker, H.S. 1955. "The Pareto Distribution and the Cobb-Douglas Production Function in Activity Analysis." *The Review of Economic Studies* 23(1): 27-31.
43. Ifrach, B., and G.Y. Weintraub. 2017. "A Framework for Dynamic Oligopoly in Concentrated Industries." *The Review of Economic Studies* 84(3): 11061150.

44. Imbens, G. and C. Manski. 2004. "Confidence Intervals for Partially Identified Parameters." *Econometrica* 72(6): 1845-1857.
45. Manski, C. F. 2003. Partial Identification of Probability Distributions. Springer-Verlag, New York.
46. Marschak, J. and W.H. Andrews. 1944. "Random Simultaneous Equations and the Theory of Production." *Econometrica* 12(3/4): 143-220.
47. Mas-Colell, A. and S. Hart. 2013. "Simple Adaptive Strategies from Regret-Matching to Uncoupled Dynamics." World Scientific Series in Economic Theory, Vol. 4.
48. Maskin, E. and Tirole, J. 1988a. "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs." *Econometrica* 56(3): 549-569.
49. Maskin, E. and Tirole, J. 1988b. "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles." *Econometrica* 56(3): 571-599.
50. McFadden, D., F. Reid, A. Talvitie, M. Johnson, and Associates. 1979. "The Urban Travel Demand Forecasting Project." Final Report, Vol. I, Institute of Transportation Studies, University of California, Berkeley. <http://emlab.berkeley.edu/wp/uto>
51. McFadden, D. 1989. "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration." *Econometrica* 57(5): 995-1026.
52. Mermelstein, B., V. Nocke, M. Satterthwaite, and M. Whinston. 2020. "Internal versus External Growth in Industries with Scale Economies: A Computational Model of Optimal Merger Policy." *Journal of Political Economy* 128(1): 301-341.
53. Nash, J. 1951. "Non-Cooperative Games." *The Annals of Mathematics* 54(2): 286-295.
54. Nash, J. 1953. "Two-person Cooperative Games." *Econometrica* 21(1): 128-140.
55. Nosko C. 2014. "Competition and Quality Choice in the CPU Market." Chicago Booth working paper.
56. Olley, S. and Pakes, A. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64(6): 1263-1298.
57. Pakes, A., 1986. "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks." *Econometrica* 54(4): 755-784.
58. Pakes, A., 2003. "A Reconsideration of Hedonic Price Indexes with an Application to PCs." *American Economic Review* 93(5): 1578-1596.

59. Pakes A. 2014. "Behavioral and Descriptive Forms of Choice Models." *International Economic Review* 55(3): 603-624.
60. Pakes A. 2016. "Methodological Issues in Analyzing Market Dynamics." In *Advances in Dynamic and Evolutionary Games: Theory, Applications, and Numerical* F. Thuijsman & F. Wagener (eds.) . Switzerland: Springer International Publishing, pp. 43-75.
61. Pakes, A. 2017. "Empirical Tools and Competition Analysis: Past Progress and Current Problems. *International Journal of Industrial Organization* 53: 241-266.
62. Pakes, A. 2020. Comment on "Theory for Extending Single-Product Production Function Estimation to Multi-Product Settings" available at <https://scholar.harvard.edu>
63. Pakes, A. and P. McGuire. 2001. "Stochastic Algorithms, Symmetric Markov Perfect Equilibria, and the 'Curse of Dimensionality. *Econometrica* 69(5): 1261-1281.
64. Pakes, A. and David Pollard. 1989. "Simulation and the Asymptotics of Optimization Estimators. *Econometrica* 57(5): 1027-1057.
65. Pathak, P. and T. Sonmez. 2008. "Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism. *American Economic Review* 98(4): 1636-1652.
66. Philippon, T. 2019. "The Great Reversal: How America Gave Up on Free Markets?" The Belknap Press of Harvard University.
67. Porter, R. 1983. "A Study of Cartel Stability: The Joint Executive Committee, 1880-1886." *The Bell Journal of Economics* 14(2): 301-314.
68. Raval, D. 2019 "Testing the Production Approach to Markup Estimation." FTC working paper.
69. Robbins, H. and S. Monro. 1951. "A Stochastic Approximation Technique." *Annals of Mathematics and Statistics*, 22(3): 400-407.
70. Roth, A.E. 1984. "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory." *Journal of Political Economy* 92(6): 991-1016.
71. Whinston, M. 2006. Lectures on Antitrust Economics, MIT Press.
72. Wollmann, T. 2018. "Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles." *American Economic Review* 108(6): 1364-1406.
73. Wollmann, T. 2019; "Stealth Consolidation: Evidence from an Amendment to the Hart-Scott-Rodino Act" *American Economic Review: Insights*. 1(1), pp. 77-94.