Market Structure and Performance

Prepared for the International Encyclopedia of the Social & Behavioral Sciences
by
Paul B. Ellickson[*]
University of Rochester

Abstract

What determines structure and profitability across markets for goods and services? Understanding the relationship between market structure and performance is critical for determining effective economic policy governing anti-trust, intellectual property, industry regulation, and international trade. Economists have analyzed these issues for many decades. An initial descriptive literature has given way to a more formalized game-theoretic approach. One stream of literature paints on a large canvas, identifying strong mechanisms that can hold across the broad run of industries. The other focuses on specific industries, building detailed models tailored to key institutional details. Both have yielded important insights and raised additional questions and directions for future research. Recent work promises to unite both streams under a common methodological framework.

Key words: Competition, Structure, Profitability, Price Cost Margins, Bounds, Endogenous Sunk Costs, Fixed Costs, Variable Profits, Barriers to Entry, Entry Models, Game Theory

Introduction

Why are some markets dominated by a handful of firms, even at a global scale, while others remain consistently fragmented? Why are some firms able to charge prices that far exceed their costs of production, while other firms can barely break even? Does competition drive firms to be more productive? Questions like these have occupied industrial organization (IO) economists – micro-economists that study markets characterized by imperfect competition – for many decades. One way of framing these questions is to focus on the connection between market structure and market performance. Do more concentrated markets facilitate greater collusion, yield higher prices, and/or experience fewer productivity gains? Can firms make strategic investments in R&D or advertising that yield market structures that are more conducive to collusion? Because economic policy is often implemented at a level that impacts many industries at the same time, it is useful to develop answers and insights that apply to the broad run of industries. At the same time, more detailed or "local" predictions may be needed to understand, for example, the impact of a particular merger in a specific industry.

Perhaps not surprisingly, economists have approached these questions from both perspectives. Using the structure of game-theory, the *bounds approach* focuses on identifying strong mechanisms that apply under the most general of conditions, while the *structural method* uses the game-theoretic structure to focus on more narrow predictions tailored to specific industries. The following short survey provides a brief overview of the evolution of these literatures. I begin with a review of an earlier descriptive literature that identified many of the key stylized facts that

---

[*] Contact Information: Paul B. Ellickson, Simon School of Business, University of Rochester, Rochester, NY 14627. Email: paul.ellickson@simon.rochester.edu

are still being grappled with today. I then highlight the limitations of this early analysis as identified by a more formal game theoretic approach to modeling industrial structure. In response, two separate but complementary streams of literature have emerged. I first discuss the bounds approach, which emphasizes broad cross-sectional predictions that hold across many industries, and then the structural method, which focuses on detailed institutional settings in which stronger modeling assumptions are appropriate. I conclude with a brief summary of a recent literature that combines features of both approaches.

Definitions and Empirical Measures

To set the stage, it is useful to first provide some definitions, along with the empirical measures used to quantify both market structure and market performance. *Performance* generally refers to how competitive (or efficient) an industry is or, more broadly, how successful it is at delivering benefits to consumers. The main focus is on competition, and its connection to the profitability of firms. Competitiveness can be captured by the degree of market power, which refers to a firm's ability to set price above the cost of producing a single additional unit. A natural measure of market power is the price cost margin ($p - c$, or $(p - c)/p$). A share-weighted average of price-cost margins yields the more general Lerner index of market power for the industry as a whole, providing an overall measure of competitiveness. Although some authors prefer to calculate a rate of return on investment (to get more directly at profits), issues of measurement and interpretation can introduce significant noise into this calculation (Fisher and McGowan, 1983), leading many researchers to turn away from these and other accounting based constructs.

Market structure refers to the identity and relative sizes of the firms that comprise an industry (and, in the case of multi-product firms, the products they produce). Often, it can be fully captured by a list of each firm and their respective market shares (i.e. their share of total industry revenue). In industries with a large number of firms, or when relying on published summary statistics, employing a function of these shares to summarize market structure can be more practical. The most common such function is the k-firm concentration ratio, which reports the combined share of the largest k firms in the industry. Typical values of k are 4 and 8, though many such measures may be reported (unfortunately, the share of the largest firm is often masked for confidentiality reasons). The sum of the squared market shares, called the Herfindahl index, provides a single measure of the *overall* structure that is often employed in merger analysis. Critically, any measure of market structure requires a careful definition of the relevant market, a task that can be complex in certain cases (e.g. does a "smart" phone belong in the same market as a digital camera?) and plays an important role in anti-trust analysis.

The Structure-Conduct-Performance Paradigm

Systematic efforts to describe the relationship between market structure and market performance began with the development of the Structure-Conduct-

Performance (SCP) paradigm, a framework most closely associated with Edward Mason and Joseph Bain. Although it has been sharply critiqued in the ensuing decades, the SCP literature served to systematically identify and highlight many of the stylized facts and empirical regularities that motivate the modern, game-theoretic approach. While too vast to adequately summarize here, Schmalensee (1989) provides a comprehensive review of the SCP literature and the stylized facts it uncovered.

The SCP framework posited a one-way chain of causation running from industry structure (firm concentration) to firm conduct (pricing) to market performance (profitability, innovation). The central claim was that high concentration would lead to high prices and thus high profits. Recognizing that free entry by competing firms ought to quickly erode these profits, significant barriers to entry would then be required to maintain a concentrated structure. In his earliest work, Bain associated these entry barriers with fixed economies of scale in production (i.e. falling average costs), which could arguably be taken as exogenous to the firm's decision-making problem (i.e. pre-determined). The validity of this exogeneity assumption will play a critical role in what follows. Analyzing a cross section of US manufacturing industries, Bain (1956) documented a significant correlation between scale economies and concentration, consistent with the framework's predictions. The rapid increase in available datasets facilitated the production of many such studies throughout the 1960s and 1970s.

At the same time, Bain and later contributors to the SCP literature identified other industries that did not feature these production economies yet still remained concentrated due to intensive investment in advertising and/or research and development (R&D). For example, the soft drink industry has remained remarkably concentrated for many decades despite little required investment in fixed costs of production. Instead, there is an intense rivalry in advertising outlays required to effectively build and maintain a national brand. The shift of focus to these new sources of scale economies led to two problems. First, the empirical results here were more mixed, particularly in the case of R&D, which exhibited little systematic correlation with market structure. More critically, many economists argued that advertising and R&D should more appropriately be thought of as decision variables of the firm, not as pre-determined features of the business environment. Addressing this concern is a key focus of the modern, game-theoretic literature.

The SCP literature was also successful in identifying a second strong empirical regularity: the impact of increased competition on prices (and price-cost margins). The hypothesis that a fall in market concentration leads to a fall in prices and price-cost margins found strong empirical support, as well as solid theoretical underpinnings. This conjecture provides another important jumping off point for the more recent literature. Unfortunately, the relationship between market structure and *profitability* proved much more vexing and the empirical evidence decidedly mixed. Measurement and interpretation issues plagued the accounting numbers used to quantify various measures of profit or rate of return (Fisher and McGowan, 1983). Perhaps more damning were issues of interpretation.

The SCP approach takes structure as given, implying that markets that are both concentrated and profitable must be in some way collusive or non-competitive.

An alternative viewpoint, most closely associated with economists from the Chicago school (Harold Demsetz and Sam Peltzman), reverses the order. To understand their concern, consider two industries, one in which all firms share the same capabilities and a second where firms exhibit sharp differences in efficiency. The industry with heterogeneous firms will likely end up more concentrated – with the efficient firms garnering a larger market share – and more profitable (at least for the leading firms, since free entry will only drive the profits of the marginal firm(s) to zero). However, in this case, the interpretation is clearly pro-competitive. Moreover, the relevant empirical evidence has tended to support the efficiency argument.

The Demsetz critique was emblematic of a deeper conceptual problem with the SCP framework. The assumption of a one-way chain of causation from structure to performance ignores the possibility that, for example, perceptions of conduct and performance might instead drive the entry decisions that determine structure. Consider a market in which firms perceive a fierce degree of price competition upon entry. Anticipating a limited scope for sustained profitability, a smaller number of firms might enter than would have had they perceived a more collusive economic environment (e.g. accommodation by a well-functioning cartel). Thus, policies that foster competition (e.g. tougher anti-trust policy) might actually lead to *more* concentrated market structures. This surprising conclusion is a central point of the bounds approach discussed below. Broadly speaking, while the SCP approach was successful in identifying a large number of empirical regularities, it failed to address the complex role of strategic interaction in determining firm conduct. These interactions play a key role in the game-theoretic approach discussed next.

The Game-Theoretic Approach to Market Structure

The formalism of game theory provided IO economists with a rigorous framework with which to examine the subtle interconnections between firm strategies, market structure, and market outcomes. Additional motivation for differences in firm shares and fortunes were found in the strategic asymmetries arising from the opportunity to move first, the ability to "learn by doing", or the incentive to exploit "network effects" across interconnected markets. However, more than a decade of careful research brought the sobering conclusion that robust results that could apply to the broad range of industries were few and far between. Seemingly arbitrary features of the game setup often lead to vastly different equilibrium outcomes. Critically, many results necessarily turn on such difficult to proxy factors as a firm's beliefs about its rivals' actions or the exact timing of investment decisions, as well as other inherently unobservable constructs such as the number and identities of potential entrants. On the one hand, the flexibility of the game-theoretic framework allows it to rationalize a wide spectrum of outcomes. On the other, a wide range of models can explain the same specific outcome, undermining a key assumption in most statistical models (a one to one mapping from parameters to outcomes) and thereby complicating statistical inference. IO economists faced a tradeoff between constructing highly specialized models tailored to a particular setting, or opting for less sharp predictions that might hold across a broad range of settings.

It is important to note that this ambiguity is a natural feature of the economic setting, not a failure of the game-theoretic approach. Realistic games often have multiple equilibria, arising from different beliefs or the vagaries of implicit coordination. More generally, the outcomes of formal models often depend delicately on assumptions that are difficult to observe or verify (e.g. Do players move simultaneously or sequentially? Is competition best modeled as occurring in prices or quantities?). This lack of information is an intrinsic feature of the economic environment. Many of the constructs that we would like to condition on (e.g. costs) are either unobserved or challenging to control for. In response, the modern game-theoretic empirical literature has moved in two separate, yet complementary directions. One embraces the ambiguity; the other places structure upon it.

The first, which is now known as the bounds approach, continues the SCP tradition in seeking to identify predictions that hold across a broad range of industries. To do so, it favors mechanisms that are robust to a wide range of modeling choices and depend only on observable features of the economic environment. The second approach, typically referred to as the structural method, focuses instead on single industries in isolation, where more specific modeling assumptions can be brought to bear, allowing unobserved elements to either be recovered from the assumptions of the model or controlled for in the statistical design. The most recent work combines elements of both.

The Bounds Approach

The bounds approach is most closely associated with the seminal work of John Sutton (1991, 1998), his students, and co-authors. This section draws heavily on Sutton's exposition. The goal of the bounds literature is to identify a number of "strong mechanisms" that characterize equilibrium outcomes across a broad range of industry settings. The idea is to invoke as few assumptions as possible, building only from the most strongly confirmed empirical regularities, and deriving testable implications that hold irrespective of arbitrary or difficult to verify modeling choices. In practice, this involves placing restrictions on outcomes, rather than strategies. By constructing, painting on so large a canvas will limit the sharpness of predictions.

Sutton starts from a general game-theoretic set-up in which a large number of firms (beginning life as potential entrants to an industry) compete in a series of stages, culminating in a final, product-market competition stage that determines industry and firm profits. In the simplest version, there are just two stages: an entry/investment stage and a price competition stage. Firms are strategic and forward-looking in their choices and investments, anticipating how their actions (and those of their rivals) will influence market outcomes. Unlike the SCP framework, entry is endogenous. For example, firms that do not foresee earning non-negative returns in the product market will choose not to enter, reflecting a key component of the Demsetz viewpoint. In this setting, the outcomes that may survive as equilibria must satisfy two fundamental properties: *viability* and *stability*. Viability ensures that all firms that choose to enter the industry expect to earn non-negative profits. This condition eliminates cases in which all firms choose to enter.

Stability requires that any firm that chose not to enter expected to lose money (i.e. earn negative profits). Note that these are standard properties that any pure strategy, Nash equilibrium must satisfy within this class of models.

Based on these two properties, Sutton establishes a series of strong results or "robust mechanisms" characterizing equilibrium market structure. The first applies to industries in which advertising and R&D play no major role in competition (i.e. the types of industries with which Bain was originally concerned). These are referred to as exogenous sunk cost markets or Type 1 industries. What happens to market structure as market size expands in these settings? Sutton shows that, in these exogenous sunk cost markets, as market size increases, the minimum level of concentration that can be supported as an equilibrium configuration decreases asymptotically to zero. Phrased differently, the *lower bound* to concentration decreases to zero as market size expands. Note that this does not imply that every market will experience a monotonic decrease in concentration as market size expands – this is a bound not a functional relationship. In some cases, markets may remain concentrated indefinitely. This lack of precision is a product of the theory's breadth: although we can construct examples in which the profitable opportunities that arise as market size expands are systematically capture by the largest firms, we cannot *rule out* equilbria in which markets become more fragmented as markets grow in size. For this reason, this is sometimes called a "fragmentation" result. Note that we will find the opposite conclusion when we shift attention to advertising and R&D intensive industries below.

Sutton's second "robust mechanism" concerns the relationship between the shape and position of this lower bound and the degree or "toughness" of competition that prevails in a given market. The degree of competition determines the shape of the function that aggregates the decisions made by all firms in the investment stage(s) and produces the vector of profits arising in the competition stage. This shape is influenced by decisions such as whether to model competition as taking place in quantities (so called "Cournot" competition, which is considered somewhat collusive) or prices (so called "Bertrand" competition, which is viewed as more competitive) and environmental factors like the degree of anti-trust enforcement and the ease of transportation. What happens to concentration or, more precisely, the lower bound to concentration, as competition becomes tougher? Consider a shift from a more collusive setting to one that is more competitive. For example, suppose the regulating authority increases the penalty (or strengthens enforcement) associated with anti-cartel laws. This change would imply an increase in the toughness of competition, as firms would presumably now find it more difficult to collude. Sutton's second result establishes that an increase in the toughness of price competition will lead to an upward shift in the lower bound to concentration, although the lower bound will still tend towards zero as market size expands. This shift is driven by the viability condition noted above: a return to equilibrium will require some firms to exit or consolidate, thereby raising the equilibrium level of concentration for a given size market. Note that this has important policy implications: efforts to increase competition can lead to a more concentrated market structure. This is likely not what one might first expect and illustrates the power of the game-theoretic perspective.

Both results have garnered strong empirical support. Sutton (1991) provides a wide-ranging, multi-industry, cross-national analysis that includes both formal statistical tests and compelling case studies. More recently, Sutton (2007) provides a comprehensive overview of the related literature, which includes many empirical studies that complement the bounds approach.

Perhaps the most provocative of Sutton's results concerns those industries in which advertising or R&D play a critical role. Owing to the strategic nature of these investments, these "Type 2" industries are referred to as *endogenous* sunk cost markets. In these settings, Sutton instead finds that concentration is bounded below, breaking the fragmentation result described above. As market size expands, the share of the largest firm remains strictly positive. The formal result can be summarized as follows. Let $S\pi$ denote the profit earned in the final stage of competition, where S represents the size of the market and π a function capturing all previous investment and pricing decisions. Assuming marginal costs are flat and markets increase in size through successive replication (meaning larger markets are no more heterogeneous than their smaller counterparts), suppose that for some constants $a > 0$ and $k > 0$, a firm that spends k times as much on fixed costs as any rival can earn at least aS in the final stage, there then exists a lower bound to concentration that is independent of the size of the market. Thus, unlike the previous exogenous fixed cost case, the importance of R&D or advertising places a sharp limit on the structure of markets that can survive in equilibrium. Fragmentation cannot occur here. Larger markets are still more profitable than small, but as market size expands the incentive to increase the level of fixed investments rises accordingly. In particular, larger markets will induce greater investment by at least some firms, eliminating the indefinite fall in concentration observed in the Type 1 markets. Sutton therefore provides a compelling reason for why some industries are consistently dominated by a handful of firms: the escalating nature of their required sunk investments. Sutton (1991) provides strong empirical support for this non-convergence result across a broad range of industries.

Applications and Extensions of the Bounds Approach

Several authors have applied Sutton's framework in a variety of institutional settings, often extending the scope of what are considered endogenous sunk investments. For example, Ellickson (2007) uses the bounds approach to analyze competition in the US supermarket industry. While relatively un-concentrated at the national level, the grocery industry is dominated at the local (city) level by 3 to 4 chain firms, an empirical regularity that holds for geographic markets both large and small. While R&D and advertising play almost no role in this industry, Ellickson argues that supermarkets invest competitively in information technology aimed at increasing the number of products offered in each store. These investments limit the number of firms that can profitably compete. A key challenge is identifying independent markets of varying size with which to test this non-convergence result. Ellickson argues that, on the supply side, relatively independent grocery markets can be defined at the level of the distribution center, and uses variation in the size of these “distribution markets” to identify the lower bound to concentration. He finds that the share of the largest firm is bounded below at a limiting level of over 12% while, in a comparison industry in which costs are exogenous (barber shops and hair salons), that level falls to under .5%.

Sutton recognized the role of independent submarkets in establishing robust predictions. The challenge is particularly acute in the case of R&D, where intense spending along a narrow research trajectory can lead to both high levels of investment and deceptively low levels of market concentration in industries where markets are defined too broadly. In Sutton (1998), he expands his earlier analysis to include settings in which R&D leads to multiple, independent trajectories. Once again taking a bounds approach, he provides a method for identifying the types of R&D intensive industries that will remain persistently concentrated. Finally, he also develops a bound prediction that applies to the long-standing debate over the size distribution of firms, providing a game-theoretic rational for what was historically a purely statistical literature.

Structural Methods

Reacting to both the inherent ambiguities embedded in the game-theoretic approach and the limitations on what we can clearly observe and quantify as empirical researchers, Sutton’s bounds approach is based on making as few assumptions as possible. In contrast, the more dominant structural approach to empirical analysis adopts a more focused perspective, emphasizing single industry settings in which stronger assumptions can be employed and (ideally) validated. In so doing, it can deliver far more precise predictions. It also makes explicit the assumptions underlying identification of causal effects, allowing the reader to judge for themselves the appropriateness of various modeling choices. Furthermore, by uncovering the economic primitives that drive agents’ strategic choices, alternative (counterfactual) scenarios can be evaluated, including large changes to economic policy that may be well outside what is observed in the data. Of course, the validity of these predictions relies on the realism of the proposed model, which must be carefully tailored to the specific institutional setting. Consequently, the structural approach focuses expressly on specific industries, where detailed data and institutional knowledge can mitigate some of the concerns regarding unobservable, or poorly measured, covariates and justify the stronger assumptions needed to sharpen our predictions. When carefully executed, structural methods can yield insights that are not available from simple descriptive analysis.

While the earliest structural work attempted to estimate a “conduct parameter” characterizing firm behavior, the modern approach assumes a form of conduct and uses detailed modeling assumptions to recover difficult to measure constructs like the variable and fixed costs of production. For example, to evaluate a merger in the soft drink industry, a detailed demand system would be estimated. By assuming a form particular of conduct (e.g. Nash in prices competition amongst differentiated firms), estimates of the variable cost of production can be recovered (eliminating the need to rely on the accounting measures criticized earlier). With these estimates in hand, a fully specified model of competition can then be used to simulate outcomes under different ownership structures, quantifying the impact on both consumers and producers (e.g. Nevo, 2000). Thus, the impact of the merger can be evaluated prior to its consummation, thereby guiding optimal policy decisions. The entry of new products or changes in other regulatory policies can be analyzed similarly. These methods now play a key role in merger analysis.

While structural methods have been used to address a broad range of questions in industrial organization, the two areas closest in spirit to the market structure and performance agenda are the "static entry" literature and the recent work on dynamic games. Both streams of literature include a substantial methodological component, as researchers work through the implications that the aforementioned informational incompleteness has on statistical inference, while tackling the daunting computational burden of estimating multi-agent models. In keeping with the applied focus of this overview, I will emphasize the more policy related issues addressed.

The static entry literature began with the seminal work of Bresnahan and Reiss (1990), which focused on the strategic entry decisions of small retail firms in isolated rural markets (the emphasis on isolated markets helps address the issue of market definition noted earlier). Bresnahan and Reiss examined the relationship between market size and the number of entrants to the market, attempting to isolate the impact of price competition on firm profits. They reasoned that if total industry profits decrease in the number of entrants and firms must pay a fixed and sunk cost to enter the market, then the market size that could accommodate two firms would need to be more than twice as large as a market that could only accommodate one (holding other drivers of profitability fixed). In particular, if competition drives down profits, the number of firms that enter a market should grow less than proportionately with the size of that market. Note that this provides an indirect test of the price-competition mechanism identified under the SCP paradigm and taken as a primitive in the bounds approach.

By focusing on the number of entrants (as opposed to their identities), they were also able to finesse some issues of incompleteness associated with multiple equilibria. In a series of papers, Bresnahan and Reiss documented a strong price-competition effect that tapered off once 3 to 5 firms were present in a market. This tapering off also provides support for the intuition that collusive behavior may be harder to sustain with a larger number of participants, a common assumption in many treatments of tacit collusion. Recognizing that the less than proportionate increase in firms might instead reflect a version of the Demsetz critique, Berry (1992) extended the Bresnahan and Reiss approach to accommodate heterogeneity in firms. Berry reasoned that if the more efficient firms enter first, it would then make sense that later entrants would require a proportionately larger market (greater demand) over which to spread their relatively higher cost outlays. Focusing on the airline industry, Berry continued to find a strong price-competition effect despite these differences in efficiency. However, in keeping with a central theme of the bounds approach, this sharper prediction required a stronger assumption: a specific mechanism for selecting amongst equilibria (with sharp differences across firms, the number of entrants may no longer be unique as it was with Bresnahan and Reiss). Subsequent authors have extended the framework to include additional sources of product differentiation (Mazzeo, 2002), network economies (Jia, 2008), and alternative information structures (Seim, 2006).

Up to now, the focus has been entirely on cross sectional implications of what are mostly static models of firm competition and industry structure. Clearly, most industries are dynamic, with structures that continue to evolve over time. Concepts like "learning by doing" and price skimming are inherently dynamic, while the payoffs to research and development are often realized several years in the future. Fortunately, in the 1970s and 1980s, detailed firm-level census data increasingly became available, allowing researchers to document industry evolution at an extremely disaggregated level. The nature of this evolution has important implications for issues like first mover advantages, industry shakeouts, learning, and innovation. What emerged was clear evidence of heterogeneity across both firms and industries, with entry and exit often occurring simultaneously. To begin to address these issues, Ariel Pakes and his co-authors have developed a flexible framework for analyzing firm dynamics from both a theoretical and empirical perspective. This framework, originally formulated by Ericson and Pakes (1995) but extended in many directions by subsequent authors, can accommodate simultaneous entry and exit, strategic investment, imperfect competition, and both firm and market level heterogeneity. While this research agenda is still in its early stages, its application has blossomed in the past few years due to recent advances in estimation and computational techniques (Doraszelski and Pakes, 2007).

Goettler and Gordon (2011) provide a nice illustration of how this framework can be applied to examine the relationship between market structure and performance. Up to now, we have emphasized the impact of market structure on firm profits, but structure should also influence other outcomes like innovation and productivity. In particular, we might ask if competition spurs productivity? Or does a monopolist, who is in an ideal position to fully appropriate the fruits of their labor, have greater incentives to innovate? Goettler and Gordon develop a structural model of competition between Intel and AMD in the production of microchips (CPUs). Entering a long-standing debate between adherents of the viewpoints offered by Kenneth Arrow and Joseph Schumpeter, the authors ask what would happen to innovation (and prices) if AMD were to cease to exist. Consistent with the Schumpeter argument, they find that the rate of innovation would be 4.2 percent higher in AMD's absence. However, consumer surplus would also fall by 4.2 percent due to the higher prices that would be charged by Intel. Whether these insights might apply in other settings is an open question, but a similar approach could be used to analyze other industries.

A Structural Bounds Literature

Both the bounds and structural IO literature are motivated by the twin limitations of our ability to predict outcomes (via formal game-theoretic models) and the degree to which we observe all the relevant institutional details and statistical covariates that drive observed structure. This model "incompleteness" arises from limited information on the part of the researcher regarding the correct modeling assumptions, the relevant covariates to include in a regression, the underlying equilibrium selection mechanism, and the key institutional details that govern firm behavior. As structural researchers seek to limit the number of assumptions they must invoke to conduct analysis, the two strands of literature start to merge. In particular, the pioneering work of Charles Manski and Elie Tamer regarding the econometric treatment of such incomplete models is now being used to tackle questions of market structure and performance. However, rather than starting from the most general setting (a la Sutton), the approach here is to begin from the particular and move to the more general. For example, in an extension of the Berry model of airline competition, Ciliberto and Tamer (2009) are able to relax the requirement of an explicit equilibrium selection mechanism in a model with heterogeneous firms, delivering on the promise of more detailed insights with fewer assumptions. Similar methods might relax the requirements on assumed conduct even further (to allow bargaining or matching models to be considered as well).

Conclusions

The relationship between market structure and market performance has occupied IO economists for several decades. Providing concrete insights is important for determining effective policy toward anti-trust, intellectual property, regulation, and international trade. An initial descriptive literature has given way to a more formalized game-theoretic approach to tackling these questions. One stream of literature paints on a large canvas, identifying strong mechanisms that can hold across the broad run of industries. The other focuses on the more narrow picture, building detailed models tailored to specific industrial contexts. Both have yielded important insights and raised additional questions and directions for future research. Recent work promises to unite both streams under a common methodological framework.

References

Bain, J.S. (1956) "Barriers to New Competition" Harvard University Press. Cambridge, MA.

Bresnahan, T.F. and P.C. Reiss (1990) "Entry in Monopoly Markets" *Review of Economic Studies* 57, 531-553.

Berry, S. (1992) "Estimation of a Model of Entry in the Airline Industry", *Econometrica* 60, 889-917.

Ciliberto, F. and E. Tamer (2009) "Market Structure and Multiple Equilibria in Airline Markets" *Econometrica* 77, 1791-1828.

Doraszelski, U. and A. Pakes (2007) "A Framework for Applied Dynamic Analysis in IO" In: Armstrong, M., Porter R. (Eds.) Handbook of Industrial Organization, vol. 3. North-Holland, Amsterdam, 1887-1966.

Ellickson, P.B. (2007) "Does Sutton Apply to Supermarkets?" *Rand Journal of Economics* 38, 43-59.

Ericson, R. and A. Pakes (1995) "Markov Perfect Industry Dynamics: A Framework for Empirical Work" *Review of Economic Studies* 62, 53-82.

Fisher, F.M. and J.J. McGowan (1983) "On the Misuse of Accounting Rate of Return to Infer Monopoly Profits" *American Economic Review* 73, 82-97.

Jia, P. (2006) "What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retail Industry", *Econometrica* 76, 1263-1316.

Goettler, R.L. and B.R. Gordon (2011) "Does AMD Spur Intel to Innovate More?" *Journal of Political Economy* 119, 1141-1200.

Mazzeo, M. (2002) "Product Choice and Oligopoly Market Structure" *Rand Journal of Economics* 33, 221-242.

Nevo, A. (2000) "Mergers with Differentiated Products: The Case of the Ready-to-eat Cereal Industry" *Rand Journal of Economics* 31, 395-421.

Schmalensee, R. (1989) "Inter-industry differences of structure and performance". In: Schmalensee, R., Willig R. (Eds.) Handbook of Industrial Organization, vol. 2. North-Holland, Amsterdam, 951-1009.

Seim, K. (2006) "An Empirical Model of Firm Entry with Endogenous Product-Type Choices" *Rand Journal of Economics* 37, 619-640.

Sutton, J. (1991) “Sunk Cost and Market Structure” MIT Press, Cambridge MA.

Sutton, J. (1998) “Technology and Market Structure” MIT Press, Cambridge, MA.

Sutton (2007) “Market Structure: Theory and Evidence” In: Armstrong, M., Porter R. (Eds.) Handbook of Industrial Organization, vol. 3. North-Holland, Amsterdam, 2301-2368.