

Multiplicity of equilibria and information structures in empirical games: challenges and prospects

Session at the 9th Triennial Choice Symposium

**Ron N. Borkovsky · Paul B. Ellickson ·
Brett R. Gordon · Victor Aguirregabiria ·
Pedro Gardete · Paul Grieco · Todd Gureckis ·
Teck-Hua Ho · Laurent Mathevet · Andrew Sweeting**

Published online: 19 June 2014
© Springer Science+Business Media New York 2014

Abstract Empirical models of strategic games are central to much analysis in marketing and economics. However, two challenges in applying these models to real-world data are that such models often admit multiple equilibria and that they require strong

Ron N. Borkovsky, Paul B. Ellickson and Brett R. Gordon were the co-chairs of the session at the 9th Triennial Choice Symposium.

R. N. Borkovsky (✉) · V. Aguirregabiria
University of Toronto, Toronto, Canada
e-mail: ron.borkovsky@rotman.utoronto.ca

V. Aguirregabiria
e-mail: victor.aguirregabiria@utoronto.ca

P. B. Ellickson
University of Rochester, Rochester, NY, USA
e-mail: paul.ellickson@simon.rochester.edu

B. R. Gordon
Columbia Business School, New York, NY, USA
e-mail: brg2114@columbia.edu

P. Gardete
Stanford University, Stanford, CA, USA
e-mail: gardete@stanford.edu

P. Grieco
Pennsylvania State University, State College, PA, USA
e-mail: plg15@psu.edu

T. Gureckis · L. Mathevet
New York University, New York, NY, USA

T. Gureckis
e-mail: todd.gureckis@nyu.edu

L. Mathevet
e-mail: lmath@nyu.edu

informational assumptions. The first implies that the model does not make unique predictions about the data, and the second implies that results may be driven by strong a priori assumptions about the informational setup. This article summarizes recent work that seeks to address both issues and suggests some avenues for future research.

Keywords Static discrete games · Dynamic games · Structural estimation · Multiplicity of equilibria · Information structures · Learning

1 Introduction

Structural estimation of strategic games has allowed researchers to address a wide variety of important questions about competitive conduct in concentrated markets. The early literature—spawned by Bresnahan and Reiss (1991)—restricted attention to static discrete games, i.e., one-shot games in which each firm faces a finite set of possible actions. This framework has been used to study entry, pricing formats, product lines, product quality, and store format and location in oligopolistic settings. Ericson and Pakes (1995) introduced an estimable model of dynamic competition that built upon the static discrete games literature in two ways. First, it incorporated forward-looking firms that interact repeatedly, allowing researchers to examine richer industry behavior. Second, it allowed firms to make not only discrete but also continuous choices. Resulting models consider dynamic pricing, R&D investment, and advertising outlays, among other topics.

Despite tremendous progress over the past decade, these related literatures continue to face similar challenges. First, the possibility of multiple equilibria poses challenges for estimation and counterfactual analysis within the context of both static and dynamic games. Second, both frameworks have, thus, far been forced to make strong informational assumptions. In this paper, we review these challenges and the recent progress made in addressing them and discuss directions for future research.

Multiplicity of equilibria poses challenges for estimation of both static and dynamic games because it gives rise to a coherency problem. In the presence of multiplicity, a model does not yield unique predictions. This makes it impossible to write down probability statements for the basic outcomes of the model and, accordingly, a likelihood function to use in estimation. We discuss the main approaches that have been used to address the coherency problem in estimating static discrete games. Multiplicity poses far greater challenges to estimation of dynamic stochastic games because simply finding multiple equilibria can be extremely difficult. We discuss an approach for systematically searching for multiple equilibria in dynamic stochastic games and explain how this approach might help scholars address the problems that multiplicity poses for both estimation and counterfactual analysis.

In the literature on empirical games, strong informational assumptions have typically been required in order to ensure tractability. We discuss three of the most prominent

T.-H. Ho
University of California-Berkeley, Berkeley, CA, USA
e-mail: hoteck@haas.berkeley.edu

A. Sweeting
University of Maryland, College Park, MD, USA
e-mail: sweeting@econ.umd.edu

assumptions and explain how recent methodological contributions can help relax them. First, the standard approach in the literature on static discrete games entails assuming that firms' payoffs are either publicly observable (complete information) or privately observable by individual firms (incomplete information). However, it stands to reason that in many markets, there are both public *and* private components to a firm's payoff function. We discuss some recent research that demonstrate how these different frameworks can lead to qualitatively different conclusions for the same empirical application. This motivates a model with a flexible information structure that includes both public and private components and allows one to estimate the extent to which firms' payoffs are publicly known. Second, in games of incomplete information, one typically assumes that a player's beliefs about the behavior of other players are *in equilibrium*, i.e., they are consistent with the information available to the player and its rivals' equilibrium strategy profiles. However, in reality firms often face significant uncertainty about their rivals' strategies. In fact, firms are often deliberately secretive about their strategies and may even try to conceal them. We discuss some recent research that allow for biased beliefs about rivals' strategies and shows that robust inference is possible in this context. Third, empirical dynamic games have typically made a symmetric information assumption, i.e., firms' perfectly observe each others' payoff relevant states (e.g., product quality, capacity, and knowledge stock) and their actions. However, the IO theory literature makes a convincing case for the role of asymmetric information in games, suggesting that these should also be explored empirically. We discuss two recent approaches that incorporate persistent asymmetric information into empirical dynamic games.

Finally, we discuss possible directions for future research related to both multiplicity of equilibria and information structures—highlighting the role that firm learning can play in facilitating progress in both areas. The theoretical literature on learning in games (Fudenberg and Levine 1998) is motivated by the notion that firms may engage in an adaptive process through which they “learn” how to play an equilibrium. Formally incorporating firm learning of this sort into IO models could help address the challenges posed by multiplicity by providing an equilibrium selection mechanism. This might be a particularly attractive solution to the coherency problem because the selection mechanism would be rooted in assumptions regarding firm behavior. Incorporating learning into games would also allow for models with much richer information assumptions that might exploit data to reveal how the information environment evolves endogenously over time.

The remainder of the paper is organized as follows. In Section 2, we discuss the challenges posed by multiplicity of equilibria and some recent progress. In Section 3, we discuss some frontier research on information structures that allow one to relax the strong assumptions traditionally made in the literature. Section 4 discusses directions for future research.

2 Multiplicity

In section, we examine why multiplicity poses a challenge for estimation and inference in economic models of games. The fundamental issue is that in the presence of multiplicity, a model predicts multiple possible outcomes for a given input without providing any guidance on their relative likelihood. In such a situation, it is impossible

to derive probability statements for the basic outcomes of the model (i.e., $P(Y|X, \theta)$ may not be defined). Without these probability statements, an econometrician cannot write down a likelihood function to use in estimation. In more complex models, even computing the set of equilibrium outcomes over which to compute the likelihood poses a challenge.

2.1 Inference

Several strategies can be used to conduct inference in the presence of multiplicity. Here, we provide only a brief overview to motivate our later discussion. The interested reader is referred to Ellickson and Misra (2011) for additional details.

First, as originally proposed by Bresnahan and Reiss (1991), one can aggregate non-unique predictions into sets of outcomes about which probability statements can be defined. For example, in an entry game, if a small market can accommodate only one of two potential entrants, but the model is not detailed enough to predict the entrant's identity, one can build a likelihood around predicting the number of firms that serve the market rather than the identities of the firms that choose to enter. While aggregation relies on the particulars of the model under consideration, it has been used with success in several empirical applications. Of course, aggregating events does result in some loss of information, as shown by Tamer (2003).

A second approach, first proposed by Ciliberto and Tamer (2009), is to use bounds on probability statements to infer the model parameters, thus making greater use of the information in the data. A notable feature of this approach is that inference is valid whether or not the model itself is point identified. A final approach is to complete the model by directly imposing an equilibrium selection mechanism. The simplest example would be to assume that a unique outcome is chosen based on a selection rule that is independent of all observable and unobservable constructs. This is the approach taken by Sweeting (2009). Bajari et al. (2010b) extend this approach to allow the selection criterion to depend parametrically on observable characteristics and the equilibrium "type" (e.g., maximizing joint payoffs, employing mixed strategies, etc.). Grieco (2014) extends the selection mechanism even further to allow it to be a non-parametric function of model observables and unobservables.

2.2 Computation

While solving for multiple equilibria of static discrete games is often straightforward, doing so for dynamic games (Ericson and Pakes 1995) is challenging. In this section, we discuss the homotopy method, a systematic approach to searching for multiple equilibria in such dynamic games.¹

Ericson and Pakes (1995) provide a canonical model of dynamic competition in an oligopolistic industry with investment, entry, and exit. Their framework facilitates empirical and numerical analyses of a wide variety of phenomena that are too complex to be explored in analytically tractable models. Methods for computing equilibria are therefore a key part of this research stream. Although equilibrium existence—in

¹ See Doraszelski and Pakes (2007) for a broad review.

particular, pure strategy Markov-perfect equilibria—is guaranteed (Doraszelski and Satterthwaite 2010), the potential for multiplicity is widely recognized.

To date, the Pakes and McGuire (1994) algorithm is most often used to compute equilibria of these dynamic games. To identify multiple equilibria (for a given parameterization of the model), the algorithm must be restarted from different initial guesses. However, different initial guesses do not necessarily lead to different equilibria and this trial-and-error approach is sure to miss a substantial fraction of them (Besanko et al. 2010a). It is important, therefore, to consider alternative algorithms that can identify multiple equilibria.

The homotopy method achieves this by tracing out an entire path in the equilibrium correspondence by varying one or more selected parameters of the model. If this path bends back on itself, then the homotopy method has identified multiple equilibria.² The homotopy method is guaranteed to find all equilibria on a path it traverses and, therefore, to find all multiple equilibria that arise in this manner. However, since multiple equilibria for a given parameterization do not necessarily lie on the same path, the homotopy method is not guaranteed to find all equilibria.

In the first paper, to apply the homotopy method to dynamic stochastic games, Besanko et al. (2010a) explore how learning-by-doing and organizational forgetting affect firms' pricing strategies and the evolution of industry structure. They find that multiple equilibria give rise to qualitatively different behaviors and short- and long-run industry structures. Several papers have since employed the homotopy method and all have found multiplicity.³

3 Information structures

In this section, we discuss recent contributions that make it possible to relax three strong and pervasive informational assumptions made in the empirical games literature. First, we discuss relaxing the assumption of either complete or incomplete information in static discrete games by incorporating public and private components into firms' payoff functions. Second, we explore relaxing the assumption that players' beliefs are *in equilibrium*, allowing for the possibility that players' beliefs about rivals' behavior may be biased. Third, we discuss two recent approaches to relaxing the symmetric information assumption typically made in empirical dynamic games. These contributions could produce new and exciting branches of the literature on empirical games for two reasons. First, by relaxing these assumptions, researchers can devise more flexible models that impose fewer restrictions on how data are interpreted. Second, these contributions allow scholars to empirically explore a variety of applied problems that were previously out of reach; Section 3.3 provides one example.

3.1 Flexible information structures

Informational assumptions about the game being played can have a substantial impact on what can be inferred from data. For example, models of strategic interaction often

² See Figure 1 in Besanko et al. (2010a) and Figure 1 in Borkovsky et al. (2012) for examples.

³ See Besanko et al. (2010b, c), Borkovsky et al. (2010, 2012), and Besanko et al. (2013).

presume either an incomplete or complete information framework. Under complete information, all players perfectly observe the payoff functions of their opponents, so the only potential source of uncertainty over rival actions arises via mixed strategies. Under incomplete information, each player knows her own type. This means she must make her decision in the face of uncertainty about her opponent.⁴

Grieco (2014) has shown that different informational assumptions can lead researchers to draw different conclusions from the same data. He advocates the use of a flexible structure that allows for both complete and incomplete information components. While the flexible structure may not be point identified, partial identification inference procedures can produce confidence intervals that are economically meaningful and robust to different combinations of complete and incomplete information. The key advantage of this approach is that results will not be driven by an a priori assumption about the informational setup.

3.2 Biased beliefs

Aguirregabiria and Magesan (2012) study identification of dynamic games when players' beliefs about rivals' actions may be biased. In empirical games with players unbiased or equilibrium beliefs, a standard assumption to obtain identification of payoff functions consists of the following exclusion restriction: Each player has an observable state variable that enters in his payoff function, but it does not have a direct effect in the payoff of the other players (see Bajari et al. 2010a, among others). Dynamic games of oligopoly competition provide multiple examples of this type of exclusion restriction. For instance, in a standard model of market entry-exit, the incumbent status of a firm in the market affects the firm's profit by determining whether it must pay an entry cost to be active in the market, but it does not have any direct effect on the profits of the other firms. The profits of the other firms are only affected indirectly through this firm's current choice to be active in the market.

Aguirregabiria and Magesan show that this exclusion restriction can be used to detect biases in players' beliefs about other players' behavior and to test the null hypothesis of unbiased beliefs. To describe the main idea of this test, consider a two-player game where each player $i \in \{1, 2\}$ chooses a binary action $Y_i \in \{0, 1\}$ to maximize his expected payoff. Let S_1 be the "special" state variable that enters into the payoff function of player 1 but not into that of player 2. Under this condition, the choice probability of player 2 (i.e., the probability of $Y_2=1$ conditional on common knowledge state variables) depends on the state variable S_1 only because player 2 believes that this state variable affects the behavior of player 1. The dependence of the choice probability of player 2 with respect to S_1 reveals information about player 2's beliefs. Aguirregabiria and Magesan show that this information identifies nonparametrically a player's beliefs as a function of the special state variable up to an intercept and a scale constant. Using this identification result, they construct a test for the null hypothesis of unbiased beliefs.

⁴ See Narayanan (2013) and Misra (2013) for Bayesian approaches to estimate complete or incomplete information games, respectively, that mix over equilibrium.

3.3 Asymmetric information

Empirical dynamic games have typically made a symmetric information assumption because it yields tremendous benefits in terms of both analytical and numerical tractability. However, in the IO theory literature, important phenomena such as limit pricing and predation have been modeled as two period asymmetric information models. For example, Milgrom and Roberts (1982) show how an efficient incumbent monopolist may set a low price to deter an uninformed potential entrant from entering. Extending these models to more realistic dynamic settings with repeated interactions is complicated by the nature of the perfect Bayesian Nash equilibrium (PBNE) concept, which requires the specification of each player's beliefs about the unobserved state variables given the entire history of the game. This is likely to be computationally intractable. Moreover, the ability to construct many different possible beliefs may lead to a severe multiplicity of equilibria, complicating both estimation and analysis of counterfactuals.

Two recent papers have taken different approaches to these problems. Fershtman and Pakes (2012) propose replacing the PBNE concept with an alternative, experience-based equilibrium (EBE), where only players' beliefs about their expected payoffs from different actions are specified rather than their beliefs about other players' types. Under additional assumptions about the maximum number of past periods on which beliefs can be conditioned, this simplifies the computational problem substantially. However, it may make the multiplicity problem worse, as all PBNEs should have associated EBEs that generate similar equilibrium actions and there may be additional EBEs too.

Gege et al. (2013) take an alternative approach. They devise a dynamic version of Milgrom and Roberts' (1982) limit pricing model with persistent asymmetric information and maintain the PBNE concept. They address the aforementioned challenges by identifying a set of assumptions on model primitives under which a combination of PBNE and standard refinements applied recursively yield a unique equilibrium. In this equilibrium, a player's action perfectly reveals his current type so that equilibrium beliefs have a particularly simple form, i.e., they depend only on the player's last action. In this way, both of the problems identified above are circumvented. They also show that the equilibrium model can explain some pricing patterns from the airline industry, in particular, the substantial decreases in price that occur when Southwest Airlines becomes a potential entrant, but not an actual entrant, on a route.

4 Suggestions for future research

In this section, we discuss possible directions for future research related to both multiplicity of equilibria and information structures—highlighting the role that firm learning can play in facilitating progress in both areas.

Multiplicity Multiplicity poses particularly serious problems for empirical dynamic games. First, as explained above, if one cannot rule out multiplicity, then one cannot write down a likelihood function. To address this problem, most estimation methods used thus far have assumed that the same equilibrium is played in all geographic markets and/or time periods. While this assumption is trivially satisfied if the

equilibrium is unique, it is potentially restrictive in the presence of multiplicity. The homotopy method could potentially help relax this assumption. If one were to devise an estimation method that accounted for the possibility that multiple equilibria are played in the data, then one would require a method for systematically searching for multiple equilibria. One could then use the homotopy method to precompute equilibria in the first stage of such an estimation method.⁵

Second, it is difficult to draw conclusions from policy experiments if there are multiple equilibria, as one cannot determine which one arises after a change in policy. By systematically searching for multiple equilibria using the homotopy method, one could thoroughly characterize the set of counterfactual equilibria and, in doing so, bound the set of possible outcomes that might arise.⁶ Echenique and Komunjer (2009) take an alternative approach within the context of models with strategic complementarities. Because such models satisfy a monotone comparative statics condition, they are able to derive testable implications for how a change in parameterization affects the “smallest” and “largest” equilibria that a model admits. Applying this approach to the analysis of Berry et al. (1999) of voluntary export restraints (VER), they show that the presence of VER results in an increase in prices when both the smallest and largest equilibria are played.

Information In many empirical settings, it is common to find that two observations (e.g., two markets) with very similar observable exogenous characteristics have very different observed endogenous outcomes (e.g., market structure or prices). In practice, the most common (almost universal) explanation for this empirical finding is that these markets are different in the values of some exogenous fundamentals that are unobservable to the researcher. However, there are situations in economics where it is plausible to argue that a substantial part of the observed variation in endogenous outcomes might be explained by variation in agents’ beliefs and not by variation in fundamentals. The separate identification of the contribution of variation in unobserved fundamentals versus variation in beliefs is an important problem, though admittedly a challenging one, in the empirical analysis of games or equilibrium models in general. There is very little work on this topic, but it provides a promising area for future research.

Related to the previous point, most of the empirical works on the importance of beliefs in explaining variation in outcomes have been concentrated in Experimental Economics and more specifically in empirical work in laboratory experiments. Extending this work to field experiments could be extremely valuable, though an important restriction in lab experiments is that preferences are assumed to be elicited and completely known to the researcher (up to some private information shock). This assumption can be relaxed, and payoffs and beliefs can be separately identified, with the help of randomized field experiments. The implementation of randomized field experiments, where the researcher can control for part of the players’ payoffs (e.g., a randomized subsidy), can be designed to generate the type of exclusion restrictions

⁵ This would build on Grieco (2014), which allows for the possibility that multiple equilibria are played in the data within the context of a static discrete game.

⁶ All the solution homotopy methods (Judd, Renner and Schmedders 2012) have the potential to be very useful in this respect.

needed to separately identify payoffs and beliefs, and they can be useful tools to bring together ideas in Experimental Economics, Quantitative Marketing, and Empirical IO.

Incorporating learning Formally incorporating learning into estimable IO models has the potential to address the challenges posed by multiplicity and to extend the literature on information structures. The literature on the theory of learning in games (Fudenberg and Levine 1998) is motivated by the notion that firms may need to engage in an adaptive process through which they learn how to play an equilibrium.⁷ Incorporating learning of this sort into IO models would allow scholars to derive equilibrium selection mechanisms.⁸ There are several reasons why this might be a particularly attractive solution to the coherency problem that multiplicity poses. First, the selection mechanisms would be rooted in basic assumptions about how firms learn to play an equilibrium. Second, such models would not only predict which equilibrium is selected but also shed light on the process through which firms arrive at the selected equilibrium over time. A recent paper by Doraszelski et al. (2014) is one of the first to start addressing such questions.

In the literature on the theory of learning in games, assumptions regarding how firms learn to play an equilibrium often relate—either directly or implicitly—to the game’s information structure and, in particular, how it evolves over time as the game is played. Incorporating learning into IO models would therefore make it possible to extend the literature on information structures in natural directions. For example, biased beliefs could endogenously evolve to equilibrium beliefs, and an incomplete information environment could endogenously evolve to a complete information environment.

The foundational research on the theory of learning in games helped inspire more recent research in behavioral economics on reinforcement learning (Roth and Erev 1995), adaptive learning (Camerer and Ho 1999), and sophisticated learning (Camerer et al. 2002) in games.⁹ Incorporating these richer models of learning into IO models would allow scholars to explore how firm behavior is affected by past successes and the history of how rivals played, as well as the extent to which firms are sophisticated/forward-looking. One appealing aspect of the experience-weighted attraction (EWA) model of Camerer and Ho (1999) is that it nests some of the more important learning processes from the theory of learning in games. Formally incorporating EWA into IO models would allow scholars to test these learning algorithms against one another or, more generally, to assess the extent to which each process is reflective of the learning that occurs in the data. One challenge is that the aforementioned frameworks are typically applied to games with known payoffs that are fixed across periods, whereas IO researchers are often interested in stochastic environments.

Rather than specifying in detail the mechanisms through which agents learn (as done in EWA), one could pursue equilibrium selection using an axiomatic approach rooted in the theory of learning in games. This would allow for equilibrium selection that is consistent with selected learning processes, but would not require formally incorporating learning into IO models. Mathevet (2013) advocates this route and proposes an axiomatic approach to study repeated interaction between two boundedly rational

⁷ An early example of empirical research that incorporates learning into games is provided by Gardete and Pedro (2013).

⁸ Lee and Pakes (2009) explore the implications for equilibrium selection of best reply and fictitious play learning processes within the context of a static game of bank ATM allocation.

⁹ See Gureckis and Love (2013) for a review of reinforcement learning.

players. After theoretically characterizing the axioms' solution, he evaluates its experimental performance in eight different repeated games and finds that it performs at least as well as the EWA learning model of Camerer and Ho (1999) and the reinforcement learning model of Erev and Roth (1998).

References

- Aguirregabiria, V. & MagesanA. (2012). Identification and estimation of dynamics games when players' beliefs are not in equilibrium. Working paper, University of Toronto.
- Bajari, P., Hong, H., Krainer, J., & Nekipelov, D. (2010a). Estimating static models of strategic interactions. *Journal of Business and Economic Statistics*, 28(4).
- Bajari, P., Hong, H., & Ryan, S. (2010b). Identification and estimation of discrete games of complete information. *Econometrica*, 78(5), 1529–1568.
- Berry, S., Levinsohn, J., & Pakes, A. (1999). Voluntary export restraints on automobiles: evaluating a trade policy. *American Economic Review*, 89(3), 400–430.
- Besanko, D., Doraszelski, U., Kryukov, Y., & Satterthwaite, M. (2010a). Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica*, 78(2), 453–508.
- Besanko, D., Doraszelski, U., Lauren, L., & Satterthwaite, M. (2010b). On the role of demand and strategic uncertainty in capacity investment and disinvestment dynamics. *International Journal of Industrial Organization*, 28(4), 383–389.
- Besanko, D., Doraszelski, U., Lauren, L., & Satterthwaite, M. (2010c). Lumpy capacity investment and disinvestment dynamics. *Operations Research*, 58(4), 1178–1193.
- Besanko, D., Doraszelski, U., & Kryukov, Y. (2013). The economics of predation: what drives pricing when there is learning-by-doing? *American Economic Review*, 104(3), 868–897.
- Borkovsky, R., Doraszelski, U., & Kryukov, S. (2010). A user's guide to solving dynamic stochastic games using the homotopy method. *Operations Research*, 58(4), 1116–1132.
- Borkovsky, R., Doraszelski, U., & Kryukov, S. (2012). A dynamic quality ladder duopoly with entry and exit: exploring the equilibrium correspondence using the homotopy method. *Quantitative Marketing and Economics*, 10(2), 197–229.
- Bresnahan, T. F., & Reiss, P. C. (1991). Empirical models of discrete games. *Journal of Econometrics*, 48(1–2), 57–82.
- Camerer, C., & Ho, T.-H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, 67, 837–874.
- Camerer, C., Ho, T.-H., & Chong, J.-K. (2002). Sophisticated learning and strategic teaching. *Journal of Economic Theory*, 104, 137–188.
- Ciliberto, F., & Tamer, E. (2009). Market structure and multiple equilibria in airline markets. *Econometrica*, 77(6), 1791–1828.
- Doraszelski, U. & Pakes, A. (2007). A framework for applied dynamic analysis in IO. In: J. J. Heckman and E. Leamer (eds.), *Handbook of Industrial Organization* (Ch. 60), North Holland.
- Doraszelski, U., & Satterthwaite, M. (2010). Computable Markov-perfect industry dynamics. *Rand Journal of Economics*, 41(2), 215–243.
- Doraszelski, U., Lewis, G., Pakes, A. (2014). Just starting out: learning and price competition in a new market. Working paper, Harvard University.
- Echenique, F., & Komunjer, I. (2009). Testing models with multiple equilibria by quantile methods. *Econometrica*, 77(4), 1281–1298.
- Ellickson, P. B., & Misra, S. (2011). Estimating discrete games. *Marketing Science*, 30(6), 997–1010.
- Erev, I., & Roth, A. (1998). Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria. *The American Economic Review*, 88(4), 848–881.
- Ericson, R., & Pakes, A. (1995). Markov-perfect industry dynamics: a framework for empirical work. *Review of Economic Studies*, 62, 53–82.
- Fershtman, C., & Pakes, A. (2012). Dynamic games with asymmetric information: a framework for empirical work. *The Quarterly Journal of Economics*, 127(4), 1611–1661.
- Fudenberg, D., & Levine, D. (1998). *The Theory of Learning in Games*. Cambridge: MIT.

- Gardete, Pedro M. (2013). Peer effects in buying behavior: evidence from in-flight purchases. Working paper, Stanford University.
- Gedge, C., Roberts, J., & Sweeting, A. (2013). An empirical model of dynamic limit pricing: the airline industry. Working paper, University of Maryland and Duke University.
- Grieco, P. (2014). Discrete games with flexible information structures: an application to local grocery markets. *RAND Journal of Economics*, 45(2), 303–340.
- Gureckis, T.M. and Love, B.C. (2013) Reinforcement learning: a computational perspective. working paper, New York University.
- Judd, K., Renner, P., & Schmedders, K. (2012). Finding all pure strategy equilibria in games with continuous strategies. *Quantitative Economics*, 3(2), 289–331.
- Lee, R. & Pakes, A. (2009). Multiple equilibria and selection by learning in an applied setting. *Economic Letters*.
- Mathevet, L. (2013). An axiomatic approach to bounded rationality in repeated interactions: theory and experiments. Working paper, New York University.
- Milgrom, P., & Roberts, J. (1982). Limit pricing and entry under incomplete information: an equilibrium analysis. *Econometrica*, 50(2), 443–459.
- Misra, S. (2013). Markov Chain Monte Carlo for incomplete information discrete games. *Quantitative Marketing and Economics*, 11(1), 117–153.
- Narayanan, S. (2013). Bayesian estimation of discrete games of complete information. *Quantitative Marketing and Economics*, 11(1), 39–81.
- Pakes, A., & McGuire, P. (1994). Computing Markov-perfect Nash equilibria: numerical implications of a dynamic differentiated product model. *Rand Journal of Economics*, 25(4), 555–589.
- Roth, A., & Erev, I. (1995). Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, 8, 164–212.
- Sweeting, A. (2009). The strategic timing of radio commercials: an empirical analysis using multiple equilibria. *RAND Journal of Economics*, 40(4).
- Tamer, E. (2003). Incomplete simultaneous discrete response model with multiple equilibria. *Review of Economic Studies*, 147–165.