Cambridge University Press
978-1-107-01606-4 - Advances in Economics and Econometrics: Tenth World Congress, Volume III: Econometrics
Edited by Daron Acemoglu, Manuel Arellano and Eddie Dekel
Excerpt
More information

# PART I

# ECONOMETRICS OF INDUSTRIAL ORGANIZATION

ONE

# Game Theory and Econometrics: A Survey of Some Recent Research

Patrick Bajari, Han Hong, and Denis Nekipelov

## 1.0 Introduction

In this chapter, we survey an emerging literature at the intersection of industrial organization (IO), game theory, and econometrics. In theoretical IO models, game theory is by far the most common tool used to model industries. In such models, a researcher specifies a set of players and their strategies, information, and payoffs. Based on these choices, the researcher can use equilibrium concepts to derive positive and normative economic predictions. The application of game theory to IO has spawned a large and influential theoretical literature (see Tirole [1988] for a survey). Game theory can be used to model a broad set of economic problems; however, this flexibility sometimes has proved problematic for researchers. The predictions of game-theoretic models often delicately depend on the specification of the game. Researchers may not be able to agree, a priori, on which specification is most reasonable, and theory often provides little guidance on how to choose among multiple equilibria generated by a particular game.

The literature that we survey attempts to address these problems by letting the data tell us the payoffs that best explain observed behavior. In the literature that we survey, the econometrician is assumed to observe data from plays of a game and exogenous covariates that influence payoffs or constraints faced by the agent. The payoffs are specified as a parametric or nonparametric function of the actions of other agents and exogenous covariates. The estimators that we discuss "reverse-engineer" payoffs to explain the observed behavior. In the past decade, researchers proposed methods to

This survey article draws heavily from a related paper that is also coauthored with Victor Chernozhukov. We acknowledge insightful discussion and comments by Martin Pesendorfer. The manuscript has been substantially revised to account for these comments.

estimate games for a diverse set of problems, including static games in which agents choose among a finite number of alternatives (Seim (2006); Sweeting (2009); Bajari, Hong, Krainer, and Nekipelov (2010), dynamic games in which the choice set of agents is discrete (Aguirregabiria and Mira (2007); Pesendorfer, Schmidt-Dengler, and Street (2008); Bajari, Chernozhukov, Hong, and Nekipelov (2009); Pesendorfer and Schmidt-Dengler (2010)); and dynamic games with possibly continuous strategies (Bajari, Benkard, and Levin (2007)).

Although these methods studied a diverse set of problems, most relied on a common insight. Games can be complicated objects, and it may take weeks of computational time to compute even a single equilibrium to dynamic games in particular (Benkard (2004); Ericson and Pakes (1995); Doraszelski and Judd (2007)). A brute-force approach that repeatedly computes the equilibrium for alternative parameter values will not be computationally feasible for many models of interest. Researchers realized that the numerical burden of estimation would be lessened if estimation were divided into two steps. In the first step, the reduced form to the game is estimated using a flexible method. The reduced form is an econometric model of how the choice of an agent depends on exogenous or predetermined variables. In many models of interest, this boils down to estimating canonical models from applied econometrics. The second step attempts to recover the structural parameters of the model, that is, how payoffs depend on actions and control variables. As we discuss in this chapter, if we condition on the reduced forms, estimation of the structural parameters can be viewed as a single agent problem. In the simplest static case, the second stage will be no more complicated than estimating McFadden's conditional logit. In the dynamic case, estimation can be performed using well-understood methods from single agent dynamic discrete choice (Rust (1994)).

An advantage of this approach is that it allows an economist to ground the specification of the game in the data rather than on prior beliefs about what is reasonable. To be clear, we are not claiming that these methods are a substitute for traditional approaches in IO theory; rather, we view these approaches as a complement that can be highly valuable in applied research. For instance, in a merger, regulation, or litigation application, an economist often is interested in using game theory to analyze a very specific market. The estimates that we describe can allow an economist to build "crash-test dummies" of what might happen under alternative policies or allow him to assist firms within those industries in decision making. In many applied-policy settings, there is little or no evidence from quasi-experimental sources to inform decision making. Also, the policy changes may be too expensive or complex to engage in small-scale experiments that gauge their effects.

In what follows, our goal is to introduce key insights in the literature to the broadest audience possible by restricting attention to models and methods that are particularly easy to understand. We do not strive for the most elaborate or econometrically sophisticated estimation strategy. Instead, we introduce simple strategies, so that after reading this chapter a nonspecialist with a working knowledge of econometrics can program the estimators using statistical packages without much difficulty. Where appropriate, we direct readers to more advanced papers that discuss more refined estimation procedures, which typically rely on insights that make our simple estimators possible. Therefore, we hope that our survey is useful to more sophisticated readers by focusing attention on what we believe to be the key principles required for estimation.

Many of the papers discussed are less than a decade old; this literature clearly is in its infancy. However, we hope to demonstrate both the generality and computational simplicity of many of these methods. The domain of applicability is not limited to IO; these methods could be applied to estimating seemingly complicated structural models in other fields, such as macro, labor, public finance, and marketing. The conceptual framework discussed in this survey offers a viable alternative to nested-swapping-fixed-point algorithms. By treating the ex-ante value as a parameter that can be identified nonparametrically, rather than as a nuisance function that is difficult to compute, the efficient and flexible estimators can be developed in the context of conditional-moment models.

## 2.0 Motivating Example

In this section, we describe a simple econometric model of a game and discuss the heuristics of estimating this model. Our example is not intended to be particularly general or realistic; rather our intention is to exposit the key principles in formulating and estimating these models in the simplest possible manner. In the following sections, we demonstrate that these principles extend to more general settings.

We first consider the static decision by a firm to enter a market, similar to the static-entry models of Bresnahan and Reiss (1991), Berry (1992), Ishii (2005), Seim (2006), Jia (2008), among others. For concreteness, suppose that we observe the decisions of two "big-box" retailers, such as Walmart and Target. For each retailer $i = 1, 2$, we observe whether it enters geographically separated markets $m = 1, \ldots, M$. This set of potential markets frequently is defined as spatially separated population centers from census or related data. We let $a_{i,m} = 1$ denote a decision by retailer $i$ to enter and $a_{i,m} = 0$ to not enter.

In the model, firms simultaneously choose their entry decisions. The model is static, and firms receive a single-period flow of profits. Economic theory suggests that a firm will enter a market $m$ if the expected profit from entry is greater than the profit from not entering. Oligopoly models suggest that a firm's profits should depend on three factors. The first factor is consumer demand, which is commonly measured by market size and represented by $POP_m$, the population of market $m$. Other variables that would proxy for demand include demographic features such as income per capita. However, to keep our notation simple, we use only a single control for demand and ignore possibly richer specifications.

A second factor that enters profit is cost. Holmes (2011) argued that distribution is a key factor in the success of big-box retailers. Walmart was founded in Bentonville, Arkansas, and its subsequent entry decisions display a high degree of spatial autocorrelation. Specifically, Walmart followed a pattern of opening new stores in proximity to existing stores and gradually fanned out from the central United States. Holmes argued that proximity to existing distribution centers explains this entry pattern. We let $DIST_{im}$ denote the closest distribution center of firm $i$ to market $m$. We use this as a measure of costs; a fully developed empirical model would use a richer specification.

A third factor that enters profits is competitive interactions. Suppose that Target is considering entering a market with 5,000 people. If Target believes that Walmart also will enter this market, there will be 2,500 customers per store. This is unlikely to be an adequate number of customers for both stores to be profitable. More generally, most oligopoly models predict that entry by a competitor will depress profits through increased competition. Therefore, $a_{-i}$, the entry decision by $i$'s competitor, should be an argument in the profit function. Oligopoly models suggest that other actions of competitors, such as pricing decisions and product choice, also may matter. However, we keep our specification parsimonious for illustration.

From the previous discussion, we specify the profits of firm $i$ as follows:

$$\begin{cases} u_{im} = \alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot a_{-i,m} + \varepsilon_{im} \text{ if } a_{im} = 1 \\ u_{im} = 0 \text{ if } a_{im} = 0 \end{cases} \tag{1}$$

Here, we normalize the profits of not entering to zero. The profit from entering depends on the exogenous covariates discussed previously and parameters $\alpha$, $\beta$, and $\delta$. These parameters index the contribution to profits of demand, cost, and competitive factors, respectively. We let $\varepsilon_{im}$ denote an independently and identically distributed shock to the profitability of firm $i$'s entry decision in market $m$. Inclusion of such shocks is standard in econometric models in which agents make discrete choices. Failure to include such shocks leads to degenerate models that make deterministic predictions and, hence, will be rejected trivially by the data.

In our model, we assume that $\varepsilon_{im}$ is private information to firm $i$. In practice, this means that firm $-i$ is unable to perfectly forecast firm $i$'s profits. In practice, this is a realistic assumption in many markets. Some researchers considered the case in which $\varepsilon_{im}$ is common knowledge; see, for example, Tamer (2003); Ciliberto and Tamer (2009); and Bajari, Hong, and Ryan (2010). The choice to model preference shocks as private information has two practical advantages. First, estimation is much easier. Many properties of the model can be studied in closed form, which then leads to clean identification. Second, flexible versions of the model can be estimated almost trivially in standard software packages (e.g. STATA and Matlab).[1]

We suppose that an econometrician has access to data on entry decisions and exogenous covariates from a cross section of markets $(a_{1,m}, a_{2,m}, POP_m, DIST_{im})$ for $m = 1, \ldots, M$. The goal of estimation will be to learn the parameters $\alpha$, $\beta$, and $\delta$ of the game. That is, we attempt to recover the game being played from the observed behavior of firms in the marketplace. Economic theory generally starts by specifying payoffs and then solving for equilibrium behavior. However, in econometrics, we study the inverse problem of recovering the game from observed actions rather than deriving the actions from the specification of the game.

We let $\sigma(a_{i,m} = 1)$ denote the probability that a firm $i$ enters market $m$. Firm $i$ will make a best response to its equilibrium beliefs about $-i$'s equilibrium entry decision. Therefore, $i$'s decision rule is:

$$a_{i,m} = 1 \Longleftrightarrow \alpha \cdot Pop_m + \beta \cdot DIST_{im} + \delta \cdot \sigma(a_{-im} = 1) + \varepsilon_{im} > 0 \qquad (2)$$

[1] However, the assumption that shocks are private information has substantive implications. Bajari, Hong, Krainer, and Nekipelov (2010) showed that the number of equilibria to the model typically will be much smaller. Indeed, in the examples they studied, the average number of equilibria appears to decrease with the number of players in the game. A game with complete information would have an increasing average number of equilibria in the number of players or other measures of the complexity of the game (McLennan 2005). Therefore, the predictions of the model are changed in ways that are not completely understood from a theoretical viewpoint. It is obvious that this is an important arena for future research. Navarro and Takahashi (2010) proposed a specification test for the private-information model.

That is, firm $i$ will enter if its expected profit from doing so is greater than zero. We note that firm $i$ does not know firm $-i$'s profit exactly because it does not observe $\varepsilon_{-im}$. Therefore, we are studying the Bayes-Nash equilibrium to the game.

In practice, it is common for researchers to assume that $\varepsilon_{im}$ has an extreme value distribution as in the conditional-logit model. If we make this distributional assumption, then the standard result from discrete-choice equation (2) implies that:

$$\sigma(a_{im} = 1) = \frac{\exp(\alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \sigma(a_{-im} = 1))}{1 + \exp(\alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \sigma(a_{-im} = 1))} \tag{3}$$

$$\sigma(a_{im} = 0) = \frac{1}{1 + \exp(\alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \sigma(a_{-im} = 1))} \tag{4}$$

We note that this closely resembles the standard binary logit model, in which choice probabilities can be expressed using the exponential function in closed form. The formula depends on exogenous covariates and parameters in a closed-form manner. The main difference is that the logit probabilities for $i$ depend on the decisions of $-i$ through $\sigma_{-i}(a_{-i,m} = 1)$. Therefore, instead of being a single-agent problem, the decisions of the agents are determined simultaneously.

We note that the equilibrium probabilities add up to one. Therefore, one of the equations in (3) and (4) is collinear. As a result, we can express the Bayes-Nash equilibrium to this model as a system of two equations in two unknowns:

$$\sigma_1(a_{1m} = 1) = \frac{\exp(\alpha \cdot POP_m + \beta \cdot DIST_{1m} + \delta \cdot \sigma_2(a_{2m} = 1))}{1 + \exp(\alpha \cdot POP_m + \beta \cdot DIST_{1m} + \delta \cdot \sigma_2(a_{2m} = 1))} \tag{5}$$

$$\sigma_2(a_{2m} = 1) = \frac{\exp(\alpha \cdot POP_m + \beta \cdot DIST_{2m} + \delta \cdot \sigma_1(a_{1m} = 1))}{1 + \exp(\alpha \cdot POP_m + \beta \cdot DIST_{2m} + \delta \cdot \sigma_1(a_{1m} = 1))} \tag{6}$$

We note that, in general, this defines a system of equations for each different market $m$, each of which admits a distinct choice-probability model. Pesendorfer and Schmidt-Dengler (2010) showed that the assumption of uniqueness of equilibrium in the data is more convincing when they are defined across different time periods for a given geographical location than when markets are defined across different geographical locations. Typically, the unique-equilibrium assumption is unlikely to hold in spatially heterogeneous markets, except in special cases such as the herding model of Bajari, Hong, Krainer, and Nekipelov (2010), in which an equilibrium shifter is clearly identifiable.

Under the assumption that only one equilibrium is played out in the data, this system is extremely convenient to work with econometrically. First, we note that the equilibrium conditions can be written in a closed form, which is much more convenient than complete-information games, in which the equilibrium set cannot be characterized and often is quite complex when the number of players is large. A second advantage is that the equilibrium will be smooth locally at all but a measure-zero set of covariates and parameters; this is because the equilibrium will inherent the smoothness of the logit model. This facilitates the econometric analysis of the model. Third, our model is a generalization of the standard binary logit model, which is one of the best-studied models in econometrics; as a result, many of the well-worn tools from this literature are applicable to our model. This is compelling for an applied researcher because the econometric analysis of discrete games, to a large extent, turns out to be a reasonably straightforward extension of discrete choice models.

## 2.1 Two-Step Estimators

Much of the literature on the empirical analysis of games relies on multistep estimators. In what follows, we describe an approach that has the greatest computational simplicity rather than focus on estimators that are more efficient or have other desirable econometric properties at the cost of being more difficult to estimate. The estimator that we describe works in two steps: (1) the economist estimates the reduced form of the model; and (2) the economist estimates the structural parameters taking the reduced form as given. We heuristically sketch this estimator for intuition. We formalize the econometric details more precisely in the next section.

### 2.1.1 Reduced Form

The reduced form is the distribution of the dependent variable given the exogenous variables in our model. Formally, the reduced form can be viewed as the solution to Equations (5) and (6). We can view this system as two equations in the two unknown entry probabilities. In general, the solution to this equation cannot be expressed in closed form and will depend on the exogenous variables $POP_m$, $DIST_{1m}$, and $DIST_{2m}$. We let $\sigma_1(a_{1m} = 1 | POP_m, DIST_{1m}, DIST_{2m})$ and $\sigma_2(a_{2m} = 1 | POP_m, DIST_{1m}, DIST_{2m})$ denote the solution.

The reduced form is a "flexible" estimate of $\sigma_1(a_{1m} = 1 | POP_m, DIST_{1m}, DIST_{2m})$ and $\sigma_2(a_{2m} = 1 | POP_m, DIST_{1m}, DIST_{2m})$. We form this using the underlying data $(a_{1,m}, a_{2,m}, POP_m, DIST_{im})$ for $m = 1, \ldots, M$ and a suitably flexible estimation method, for example, a sieve logit (Newey and Powell [2003] and Ai and Chen [2003]). An important observation is that provided that the private shocks are independent across players, we do not need to estimate the joint probability of actions of all players. Estimating the choice probability for each player one at a time is sufficient to provide consistent estimates in the second stage. Here, we use $M$ to denote the number of markets that can be pooled because the same equilibrium is played in these markets. Typically, $M$ denotes the number of time periods in a single geographical location, as in Pesendorfer and Schmidt-Dengler (2010). Occasionally, in rare cases, it also might refer to the number of spatially heterogeneous markets. We let $z_k(POP_m, DIST_{1m}, DIST_{2m})$ denote the vector of terms in a $k^{th}$-order polynomial in $POP_m, DIST_{1m}, DIST_{2m}$. We could model the reduced-form choice probabilities as follows:

$$\sigma_i(a_i = 1|s, \beta) = \frac{\exp(z_k(s)'\theta)}{1 + \exp(z_k(s)'\theta)}$$

where we let $k \to \infty$ as the number of markets $M \to \infty$, but not too fast (i.e. $\frac{k}{M} \to 0$). For any finite sample size, we can estimate the sieve logit using a standard software package (e.g. STATA). This is simply a method to model choice probabilities in a flexible way that exhausts information in the data. Other flexible methods also are possible. In the next section, we demonstrate that in many cases, the choice of method to estimate the first stage typically does not matter for the asymptotic distribution of structural parameters. In our applied work, we found that with large sample sizes, results are reasonably robust to the specification of the first stage as long as it is sensibly and flexibly specified. We let $\widehat{\sigma}_1(a_{1m} = 1|POP_m, DIST_{1m}, DIST_{2m})$ and $\widehat{\sigma}_2(a_{2m} = 1|POP_m, DIST_{1m}, DIST_{2m})$ denote these first stage estimates.

If our problem is well behaved and the data across multiple markets are generated from a unique equilibrium, it is the case that $\widehat{\sigma}_i(a_{im} = 1|POP_m, DIST_{1m}, DIST_{2m})$ will converge to:

$$\sigma_i(a_{im} = 1|POP_m, DIST_{1m}, DIST_{2m})$$

as the sample size becomes large. When multiple equilibria are present in different markets, such estimates potentially can diverge. We suppose that we replace $\sigma_i$ with our consistent estimate $\widehat{\sigma}_i$ in Equation (2). Then, the agent's decision rule can be rewritten as follows:

$$a_i = 1 \Longleftrightarrow \alpha \cdot Pop_m + \beta \cdot DIST_{im} + \delta \cdot \widehat{\sigma}_{-i}(a_{-im} = 1) + \varepsilon_{im} > 0$$