

# Estimation of (Dynamic) Games: A Discussion

Martin Pesendorfer\*

London School of Economics and Political Sciences

December 2010

## 1 Introduction

The last ten years saw a number of advances in empirical industrial organizations. Exciting new methods for estimating dynamic models and especially dynamic games have been developed. These new methods enable researchers to use a wider set of empirical tools. The new methods benefitted the well received literature on estimating static games. At the same time there have been improvements in single agent dynamic model estimation.

The two papers provide excellent surveys covering selected setups in detail. The setups include models with strategic interaction in static and dynamic settings and dynamic single agent models with a focus on storable and durable goods demand models. The single and multi-agent dynamic frameworks enable researchers to explicitly model dynamic linkages in state variables or in payoffs. The setups also nests the static games (or single agent) setting which arises as a special case when agents are myopic and assign no value to the future.

Three types of dynamic linkages can be distinguished: First a strategic linkage arises when a firm's (consumer's) strategy is allowed to depend on past choices of (all) firms (consumers). Second, an informational linkage arises when informational uncertainty is gradually resolved over time. The third type of linkage arises when the

---

\*I wish to thank the seminar audience at the World Congress for feedback and to Philipp Schmidt-Dengler for comments on the draft.

current payoff depends on state variables influenced by prior actions. The surveyed literature studies the third type of linkage.

Example of payoff linkages in industrial organization are numerous. In the highway paving industry committed capacity from prior won projects may affect pricing decisions for current projects. A second example are storable goods which may affect consumers demand. Ketchup or laundry detergents can be stored at home and thus purchased when store prices are low. Durable goods are another example. Computer printers or automobiles are durable and thus the decision to buy a new product depends on the use-value of the currently held product. A consumer may decide when to replace the existing product with a new one. Finally, entry (and exit) incurs a sunk cost (scrap value) which in turn implies a dynamic linkage as already active firms will face different costs of being active.

The papers contain a vast amount of material, making it impossible to comment in detail on every section. My plan is to re-iterate some basic elements the papers share and to place the papers within a broader context of the field. Doing so I will restate the conceptual framework that is shared among the papers.

I will give critical remarks on some contentious issues that have caused some difficulties in the empirical IO literature over the years. These includes equilibrium multiplicity, unobserved heterogeneity, the iterative methods based on the best reply mapping and the characteristics space demand model.

The next section illustrates some basic features of the framework and analysis which are shared in the two papers. Section 3 comments and concludes on limitations and issues that are open for further research.

## 2 Empirical Framework

This section sketches some key features of the empirical model that are shared in the two papers. I highlight the equilibrium properties and describe the basic estimation technique.

### 2.1 Model Sketch

The framework shared in both surveys is a  $N$  player dynamic game with discrete time,  $t = 1, 2, \dots, \infty$ . Every period the following events take place:

*States.* Every player observes the state variables which consist of two elements: First a *public* state variable  $s_i^t \in \mathbf{S}_i = \{1, \dots, L\}$  which is observed by all players. Second, a *private* payoff shock  $\varepsilon_i^t \in \mathbb{R}^K$  drawn independently from the strict monotone and continuous conditional distribution function  $F(\cdot | s_i^t, \mathbf{s}_{-i}^t)$ . The private state variable is observed only by the player and not observed by other players or the econometrician.

*Actions.* Each player chooses an action or makes a choice  $a_i^t \in A_i = \{0, 1, \dots, K\}$ . All  $N$  players make their decisions simultaneously. The actions are taken after players observe the state and their idiosyncratic payoff shocks. An action profile  $\mathbf{a}^t$  denotes the vector of joint actions in period  $t$ ,  $\mathbf{a}^t = (a_1^t, \dots, a_N^t)$ .

*State transition.* The probability density function  $g : A \times \mathbf{S} \times \mathbf{S} \rightarrow [0, 1]$  describes the state transition. A typical element  $g(\mathbf{a}^t, \mathbf{s}^t, \mathbf{s}^{t+1})$  equals the probability that state  $\mathbf{s}^{t+1}$  is reached when the current action profile and state are given by  $(\mathbf{a}^t, \mathbf{s}^t)$ . The symbol  $\mathbf{G}$  denotes the transition matrix in which column  $s' \in \mathbf{S}$  consists of the vector of probabilities  $\left[ g(\mathbf{a}, \mathbf{s}, \mathbf{s}')_{\mathbf{a} \in \mathbf{A}, \mathbf{s} \in \mathbf{S}} \right]$ .

The *period payoff of player  $i$*  is collected at the end of the period after all actions are observed and consists of two components: (i) the action and state dependent period payoff, and (ii) the action dependent payoff shock realization,

$$\pi_i(\mathbf{a}^t, \mathbf{s}^t) + \sum_{k=1}^K \varepsilon_i^{tk} \cdot \mathbf{1}(a_i^t = k)$$

where  $\mathbf{1}(a)$  equals one if event  $a$  occurs and zero otherwise. The symbol  $\mathbf{\Pi}_i$  denotes the period payoff vector defined by  $\mathbf{\Pi}_i = \left[ \pi_i(\mathbf{a}, \mathbf{s})_{\mathbf{a} \in \mathbf{A}, \mathbf{s} \in \mathbf{S}} \right]$ . The payoff  $\pi_i$  may be parametrized in applications. For example, the demand specification in AN(2010) assumes that the utility of choice  $k$  in state  $s$  can be written as  $\pi_i(k, \mathbf{s}) = u_i - C_i(s^t) + a_{kt}\beta_i - \alpha_i p_{kt} + \xi_k$ , which includes a consumption benefit  $u_i$ , a storage cost  $C_i$ , and an additive term measuring the indirect utility from purchasing product  $k$ .

*Game payoff.* Players discount future payoffs with discount factor  $\beta < 1$ . The game payoff of player  $i$  equals the sum of discounted period payoffs.

The surveys restrict attention to pure Markovian strategies  $a_i(\varepsilon_i^t, s^t)$  in which an action for player  $i$  is a function of the player specific payoff shock and the state variables. The Markovian assumption abstracts from calendar time and the time superscript may be omitted.

*Definition (Markov perfect equilibrium):* A collection  $(a, \sigma) = (a_1, \dots, a_N, \sigma_1, \dots, \sigma_N)$  is a Markov perfect equilibrium if

- (i) for all  $i$ ,  $a_i$  is a best response to  $a_{-i}$  given the beliefs  $\sigma_i$  at all states  $\mathbf{s} \in \mathbf{S}$ ;
- (ii) all players use Markovian strategies;
- (iii) for all  $i$  the beliefs  $\sigma_i$  are consistent with the strategies  $a$ .

The ex ante expected value function (or ex ante expected game payoff) is the discounted sum of future payoffs for player  $i$  before player-specific shocks are observed and actions are taken. Let  $\sigma_i(\mathbf{a}|\mathbf{s})$  denote the conditional ex ante (before  $\varepsilon_i$  is observed) belief of player  $i$  that action profile  $\mathbf{a}$  will be chosen conditional on state  $\mathbf{s}$ . The ex ante expected value function can be written as,

$$V_i(\mathbf{s}; \sigma_i) = \sum_{\mathbf{a} \in \mathbf{A}} \sigma_i(\mathbf{a}|\mathbf{s}) \left[ \pi_i(\mathbf{a}, \mathbf{s}) + \beta \sum_{\mathbf{s}' \in \mathbf{S}} g(\mathbf{a}, \mathbf{s}, \mathbf{s}') V_i(\mathbf{s}'; \sigma_i) \right] + \sum_{k=1}^K E_\varepsilon [\varepsilon_i^k | a_i = k] \cdot \sigma_i(a_i = k | \mathbf{s}), \quad (1)$$

where  $E_\varepsilon$  denotes the expectation operator with respect to the player specific productivity shock. Equation (1) is satisfied at every state vector  $\mathbf{s} \in \mathbf{S}$ .<sup>1</sup>

Continuation values can also be written after a choice has been made. Player  $i$ 's continuation value net of the payoff shocks under action  $a_i$  with beliefs  $\sigma_i$  is given by:

$$u_i(a_i; \sigma_i, \boldsymbol{\theta}) = \sum_{\mathbf{a}_{-i} \in \mathbf{A}_{-i}} \sigma_i(\mathbf{a}_{-i} | \mathbf{s}) \cdot \left[ \pi_i(\mathbf{a}_{-i}, a_i, \mathbf{s}) + \beta \sum_{\mathbf{s}' \in \mathbf{S}} g(\mathbf{a}_{-i}, a_i, \mathbf{s}, \mathbf{s}') V_i(\mathbf{s}'; \sigma_i) \right] \quad (3)$$

where  $\sigma_i(\mathbf{a}_{-i} | \mathbf{s})$  denotes the conditional beliefs of player  $i$  that players  $-i$  choose an action profile  $\mathbf{a}_{-i}$  conditional on state  $\mathbf{s}$ .

Next, the properties of the equilibrium decision rules are summarized.

## 2.2 Equilibrium Properties

Let  $\boldsymbol{\theta}$  denote the parameter vector summarizing the model elements  $(\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_N, F, \beta, g)$ .

It is optimal to choose action  $a_i$  given the beliefs  $\sigma_i$  if the continuation value under action  $a_i$  exceeds the continuation value under all alternative actions:

$$u_i(a_i; \sigma_i, \boldsymbol{\theta}) + \varepsilon_i^{a_i} \geq u_i(a'_i; \sigma_i, \boldsymbol{\theta}) + \varepsilon_i^{a'_i} \quad \text{for all } a'_i \in \mathbf{A}_i \quad (4)$$

---

<sup>1</sup>The state space is finite and the ex ante value function can be re-written as an explicit matrix equation:

$$\mathbf{V}_i(\sigma_i) = [\mathbf{I}_s - \beta \boldsymbol{\sigma}_i \mathbf{G}]^{-1} [\boldsymbol{\sigma}_i \boldsymbol{\Pi}_i + \mathbf{D}_i(\sigma_i)]. \quad (2)$$

where  $\mathbf{D}_i(\sigma_i) = [D_i(\mathbf{s}, \sigma_i)]_{\mathbf{s} \in \mathbf{S}}$  is the vector of expected payoff shocks with element  $D_i(\mathbf{s}, \sigma_i) = \sum_{k=1}^K E_\varepsilon [\varepsilon_i^k | a_i = k; \sigma_i] \cdot \sigma_i(a_i = k | \mathbf{s})$  and  $\mathbf{I}_s$  denotes the identity matrix.

where  $\varepsilon_i^0 = 0$ . We may evaluate condition (4) before the payoff shock  $\varepsilon_i$  is observed. Doing so, gives an expression for the ex ante optimal choice probability of player  $i$  given the beliefs  $\sigma_i$ :

$$\begin{aligned} p(a_i|\mathbf{s}, \sigma_i) &= \Psi_i(a_i, \mathbf{s}, \sigma_i; \boldsymbol{\theta}) \\ &= \int \prod_{k \in \mathbf{A}_i, k \neq a_i} \mathbf{1}(u_i(a_i; \sigma_i, \boldsymbol{\theta}) - u_i(k; \sigma_i, \boldsymbol{\theta}) \geq \varepsilon_i^k - \varepsilon_i^{a_i}) dF \end{aligned} \quad (5)$$

In matrix notation, equation (5) can be written as:

$$\mathbf{p} = \Psi(\boldsymbol{\sigma}; \boldsymbol{\theta}) \quad (6)$$

where  $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_N)$  denotes the vector of the optimal players' choice probabilities for all states, players and actions other than action 0; and  $\boldsymbol{\sigma}$  denotes the vector of players' beliefs. We may omit the choice probability for action 0 in equation (6) as it is already determined by the remaining choice probabilities,  $p_i(0|\mathbf{s}) = 1 - \sum_{k=1}^K p_i(k|\mathbf{s})$ .

Equation system (6) determines the optimal choice probabilities as a function of players' beliefs. In equilibrium, beliefs are consistent leading to a fixed point problem in ex ante choice probabilities  $\mathbf{p}$ :

$$\mathbf{p} = \Psi(\mathbf{p}; \boldsymbol{\theta}). \quad (7)$$

Equation system (7) characterizes the equilibrium choice probabilities  $p$ . Note that equation system (7) is not the only equation system doing so. For example, any monotone transformation of the left hand side and right hand side variables in (7) can also be used as an equation system characterizing the equilibrium. Static models, as in BHN (2010), are a special case with  $\beta = 0$ . The single agent dynamic demand model studied in AN (2010) arises as a special case when there is no strategic interaction, or when the number of players equals one,  $N = 1$ .

*Limitations.* The modelling framework and approach has limitations: First, it is better suited for situations in which there is a small number of players and states. In situations with many players, or states, the transition matrix  $G$  will be large and solving the equations in (7) may become computationally involved. Second, it may be difficult to simulate the equilibria even if estimation is possible as there can be multiple equilibria. Finding the set of fixed points of equation (7) may be computationally infeasible.

BHN (2010) nicely illustrate restrictions that are commonly used in the literature to ensure that there exists a unique set of model primitives  $(\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_N, F, \beta, g)$

that can be inferred from a sufficiently rich data set characterizing choice and state transition probabilities.

## 2.3 Estimation Method

Let  $\boldsymbol{\theta} = (\boldsymbol{\theta}^\pi, \boldsymbol{\theta}^F, \beta, \boldsymbol{\theta}^g) \in \Theta \subset \mathbb{R}^q$  denote an identified parameter vector.<sup>2</sup> The parameter vector  $\boldsymbol{\theta}$  may include parameters entering the period payoffs  $\pi_i(\mathbf{a}, \mathbf{s}; \boldsymbol{\theta}^\pi)$  for all  $i$ , the transition probabilities  $g(\mathbf{a}, \mathbf{s}, \mathbf{s}'; \boldsymbol{\theta}^g)$ , the distribution of payoff shocks  $F(\varepsilon; \boldsymbol{\theta}^F)$  and the discount factor  $\beta$ .

The surveys discuss two types estimation approaches commonly used in this literature. The direct likelihood approach and the indirect two step estimator.

The direct likelihood approach is frequently used for single agent problems. It is based on the idea that for a given parameter value the implied optimal choices and the value function can be calculated by solving fixed point problems numerically. The likelihood is maximized by solving the fixed point problems repeatedly for each set of parameters. As AN (2010) nicely illustrate the direct approach can be extended to incorporate unobserved state variables by adding an additional layer in which the distribution of these unobserved states is estimated.

The second estimation approach is indirect building on the idea of two step estimators. In the first step the choice probabilities  $\widehat{\mathbf{p}}_T$  and transition probabilities  $\widehat{\mathbf{g}}_T$  are consistently estimated directly from the data by using the sample frequencies, a Kernel smoother or some other consistent method. In the second step the parameter vector  $\boldsymbol{\theta}$  is estimated using the least squares criterion function

$$\min_{\boldsymbol{\theta} \in \Theta} [\widehat{\mathbf{p}}_T - \Psi(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T, \boldsymbol{\theta})]' \mathbf{W}_T [\widehat{\mathbf{p}}_T - \Psi(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T, \boldsymbol{\theta})]. \quad (8)$$

where  $\mathbf{W}_T$  is a symmetric positive definite weight matrix. The criterion function is based on the equilibrium relationship evaluated at the first step probability estimates,  $\widehat{\mathbf{p}}_T - \Psi(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T, \boldsymbol{\theta})$ . The least squares estimator  $\widetilde{\boldsymbol{\theta}}_T(\mathbf{W}_T)$  is the one that sets the equilibrium relationship closest to zero in the metric associated with the scalar product defined by  $\mathbf{W}_T$ . Under suitable regularity conditions the least squares estimator  $\widetilde{\boldsymbol{\theta}}_T(\mathbf{W}_T)$  is consistent and asymptotically normally distributed.

A number of two step estimators proposed in the literature correspond to different choices of the weight matrix  $\mathbf{W}_T$ . These estimators include the pseudo maximum

---

<sup>2</sup>We include  $\boldsymbol{\theta}^g$  in the parameter vector  $\boldsymbol{\theta}$  to permit the possibility that a parametric restriction is placed on the state transition  $g$ .

likelihood estimator discussed extensively in AN (2010) and methods of moments estimators. As explained earlier, the equilibrium relationship  $\hat{\mathbf{p}}_T = \Psi(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta})$ , is not the only relationship characterizing equilibria. For example, the logarithmic transformation can be applied on both sides of the equation. This transformation is illustrated in BHN (2010). The optimal weight matrix  $\mathbf{W}_T$  exists and yields the efficient least squares estimators. It is characterized in AN (2010). The estimator with the optimal weight matrix is equivalent to the maximum likelihood estimator.

The two survey papers discuss a number of issues associated with the estimation method and give some applications. Next, I will comment on some issues and explanations given in the two surveys which I found not entirely convincing.

### 3 Discussion and Comments

This section discusses some limitations and comments on potential problems with the described framework. I will comment on (i) multiplicity of equilibria, (ii) unobserved state variables, (iii) iterative schemes, (iv) the demand framework, and (v) data sets.

#### 3.1 Multiplicity of Equilibria

The static games literature on entry during the late 1980s and early 1990s struggled with the difficulties arising from equilibrium multiplicity inherent to games. The aim of the static literature was to restrict the payoff space or to make additional assumption in order to find a unique prediction that enables estimation. A key advantage of the recent dynamic games literature is that the uniqueness concern can be avoided by making use of the time dimension in the data.

Jofre-Bonet and Pesendorfer (2003) introduce the idea that two-step estimators can be used to estimate a Markovian industry equilibrium. The California highway paving industry is examined as one regional market in which the same players interact over time. The observed outcomes, or data, are generated from play from one equilibrium of the infinitely repeated game. The Markovian framework enables the use of repeated observations stemming from one market to obtain consistent estimates. Thus, asymptotics is in the time dimension. By using the time dimension equilibrium multiplicity concerns are entirely avoided.

Industry data frequently contain repeated observations of the same players at

different time periods. The interpretation of the industry time series in terms of the earlier described framework is a single realization from the equilibrium path. This is a powerful new estimation technique especially well suited for industry data.

Some recent papers have extended the approach to estimate using data with both, a time and cross sectional dimension. Cross sectional data bring along concerns about multiplicity of equilibria, and unless the equilibrium is unique, the resulting estimates may be inconsistent as is nicely illustrated in BHN (2010). The inconsistency arises as the equilibrium is chosen endogenously in every market. The equilibrium choice is made by market participants. The equilibrium choice is not observed and cannot be ignored (or assumed away) by the econometrician. Therefore, ‘assuming that the same equilibrium is played in distinct markets’ does not address the problem. To the contrary, proceeding on the basis of this ‘assumption’ may yield inconsistent estimates.

Let me illustrate the inconsistency concern in more detail. The assumption of an unbounded error support implies that any outcome is possible in any equilibrium. Equilibria differ only in the probability weight assigned to individual outcomes. Now, pooling choices across distinct markets may not reflect an equilibrium any longer as the mixture of equilibria may not be an equilibrium in itself. So, parameter estimates may become inconsistent if cross sectional data are pooled. The approach described in AN (2010), which is based on ‘the assumption that the same equilibrium is played in distinct markets,’ suffers from this difficulty. The assumption is about endogenous outcomes and not about model primitives. The assumption may not hold. The approach may lead to inconsistent estimates.

Cross-sectional data can be used to test the null hypothesis that the same equilibrium is played in distinct cross-sectional markets. The test approach is nicely illustrated in BHN (2010). If the time dimension for each market is sufficiently long, then the null hypothesis of identical conditional choice probabilities in distinct markets can be tested with suitable statistical tests. If the null is not rejected, then the data may be pooled to obtain more precise estimates. Yet, it needs to be emphasized that this approach yields consistent estimates as the time dimension gets large but not as the cross sectional dimension gets large.

## **3.2 Unobserved State Variables**

Unobserved state variables may arise when the researcher has access to less information than was available to the market participants. We shall consider single and multi-agent situations.

Unobserved state variables for single agent models can be built into the described framework by creating an outer loop in the likelihood procedure. This approach is described in AN (2010).

Unobserved state variables for games cause trickier problems as now it matters who observes what. Two cases can be distinguished: First, suppose the state variables are known to all players but not to the econometrician. Second, the state variables are not known to all players but instead private information. The first case fits nicely into the described framework of estimating dynamic games as is illustrated in BHN (2010). As long as there is a way to estimate the choice probabilities consistently, payoff parameters can be inferred using the earlier described method.

The second case, when unobserved states are persistent and not known to other players but private information, is very tricky as in this case the Markovian restriction has no ‘bite’. In this case players try to infer other players state based on past play which gives rise to a model of learning going beyond the Markovian framework analyzed in the literature and discussed in the two surveys.

### 3.3 Recursive Schemes

The recursive scheme discussed in section 3 of AN (2010) has serious shortcomings. First, it does not work well in applications. Second, it may yield inconsistent estimates. Finally, the recursive scheme is not efficient.

Lyapunov stability of the best response mapping is not a convincing equilibrium selection rule in private information games. Iterative application of the best response mapping is reasonable for complete information games, such as a Cournot game, in which strategies are observable. Here players make their choices based on private information and strategies are not observable. Thus, the concept of ‘learning’ from observing other players’ play is not describable in a simple way and goes well beyond the Markovian framework.

Some numerical methods to calculate Markovian equilibria are based on updating the best response mapping. These numerical papers suffer from the same difficulty as they do not find the set of all equilibria but only those equilibria that correspond to a stable fixed point in the best response mapping.

### 3.4 Dynamic Demand

The survey AN (2010) discusses a dynamic demand model which builds on the classic characteristic space model for static discrete choices. The dynamic element is incorporated by adding-on a dynamic equation in some way.

Strong functional form assumptions are the building block for the static characteristic-space discrete-choice demand model. The functional form assumptions determine the indirect utility and the nature in which the econometric error term enters the framework. These functional form assumptions impose restrictions on the underlying preferences of agents. Yet, the link between these functional form assumption and neoclassical demand theory is not well understood, not even in a static context. In a sense the static discrete choice demand model is a black box that is difficult to understand already.

Incorporating dynamics into this framework by adding a dynamic storable goods relationship or some other dynamic element is a challenge. Recent research incorporates a parametric dynamic equation into the static model in some way without being well grounded in dynamic choice theory. The model parameters of the microeconomic model are then estimated by comparing the moments of the simulated micro model to the observed data sometimes at the aggregate level. Yet, it is unclear what drives the parameter estimates as the model framework tends to lack transparency.

### 3.5 Data

The two surveys describe a number of new empirical methods for single-industry time-series data. A natural question to ask is whether these methods provide a good fit. Are some methods better than others? How well do these methods perform with field data? A typical paper in this literature makes some simulations or studies an application to show that the method is feasible. Individual applications are based on new data sets or on new simulations making a comparison across papers very difficult and intractable.

My suggestion, or rather an appeal, is to create a few standardized industry data sets and to test these methods on these standardized data sets. Other applied fields in economics, such as labor economics, agree on a few basic data sets. Methods have to prove themselves on these data sets and can be compared directly. The dynamic games literature and IO in general would become much more tractable if researchers

studied a few key industries with publicly shared data. The features of the individual methods are then readily comparable.

## References

- [1] Aguirregabiria, V., and A. Nevo (2010), Recent Developments in Empirical Dynamic Models of Demand and Competition in Oligopoly Markets.
- [2] Bajari, P., H. Hong and D. Nekipelov (2010), Game Theory and Econometrics: A Survey of Some Recent Research.
- [3] Jofre-Bonet, M., and M. Pesendorfer (2003), Estimation of a Dynamic Auction Game, *Econometrica*, 71, 1443-1489.