



Munich Personal RePEc Archive

Recent developments in empirical IO: dynamic demand and dynamic games

Aguirregabiria, Victor and Nevo, Aviv

University of Toronto, Northwestern University

29 December 2010

Online at <https://mpra.ub.uni-muenchen.de/27814/>

MPRA Paper No. 27814, posted 03 Jan 2011 19:48 UTC

Recent Developments in Empirical IO: Dynamic Demand and Dynamic Games*

Victor Aguirregabiria
University of Toronto

Aviv Nevo
Northwestern University

First version: July 8, 2010. This version: December 29, 2010.

Abstract

Empirically studying dynamic competition in oligopoly markets requires dealing with large states spaces and tackling difficult computational problems, while handling heterogeneity and multiple equilibria. In this paper, we discuss some of the ways recent work in Industrial Organization has dealt with these challenges. We illustrate problems and proposed solutions using as examples recent work on dynamic demand for differentiated products and on dynamic games of oligopoly competition. Our discussion of dynamic demand focuses on models for storable and durable goods and surveys how researchers have used the "inclusive value" to deal with dimensionality problems and reduce the computational burden. We clarify the assumptions needed for this approach to work, the implications for the treatment of heterogeneity and the different ways it has been used. In our discussion of the econometrics of dynamics games of oligopoly competition, we deal with challenges related to estimation and counterfactual experiments in models with multiple equilibria. We also examine methods for the estimation of models with persistent unobserved heterogeneity in product characteristics, firms' costs, or local market profitability. Finally, we discuss different approaches to deal with large state spaces in dynamic games.

Keywords: Industrial Organization; Oligopoly competition; Dynamic demand; Dynamic games; Estimation; Counterfactual experiments; Multiple equilibria; Inclusive values; Unobserved heterogeneity.

*We wish to thank Gautam Gowrisankaran, Igal Hendel, Pedro Mira, Martin Pesendorfer, Marc Rysman, Matt Shum, and Junichi Suzuki for thoughtful comments.

1 Introduction

Important aspects of competition in oligopoly markets are dynamic. Demand can be dynamic if products are storable or durable, or if utility from consumption is linked intertemporally. On the supply side, dynamics can be present as well. For example, investment and production decisions have dynamic implications if there is learning by doing or if there are sunk costs. Identifying the factors governing the dynamics is key to understanding competition and the evolution of market structure, and for the evaluation of public policy. Advances in econometric methods and modeling techniques, and the increased availability of data, have led to a large body of empirical papers that study the dynamics of demand and competition in oligopoly markets.

A key lesson learned early by most researchers is the complexity and challenges of modeling and estimating dynamic structural models. The complexity, and "curse of dimensionality" is present even in relatively simple models, but is especially problematic in oligopoly markets where firms produce differentiated products, or have heterogeneous costs. These sources of heterogeneity typically imply that the dimension of these models, and the computational cost of solving and estimating them, increases exponentially with the number of products and the number of firms. As a result, much of the recent work in structural econometrics in IO has focused on finding ways to make dynamic problems more tractable in terms of computation, and careful modeling to reduce the state space while properly accounting for rich heterogeneity, dynamics, and strategic interactions.

We cannot provide here a complete survey of the large body of recent work. Instead, we focus on three main challenges that have been discussed in the literature and that we consider particularly important for applied work. Two challenges are common to models of dynamic demand and dynamic games: (1) the dimensionality problem and ways to reduce the state space and the computational burden; (2) the treatment of heterogeneity in firm, consumer, and market characteristics. The empirical application of dynamic games has to also deal with (3) the challenge of multiplicity of equilibria in estimation and prediction.

A key focus in dynamic structural models is on ways to reduce the state space. A problem that is tractable in an example used to illustrate a method might quickly become intractable when applied to answering questions in real markets. For instance, even static models of demand for differentiated products face a significant dimensionality problem due to the large number of products. The dimensionality problem becomes a difficult issue when we try to extend the methods of Rust (1987), originally applied to a durable good decision – the replacement of a bus engine – to demand for durable differentiated products. The dimension of the state space increases (exponentially)

with the number of products. A concept that has proved very useful in reducing the state space in the modeling both of dynamic demand and dynamic games is the inclusive-value (McFadden, 1974). We show different examples of how the inclusive value has been used to reduce the state space and the assumptions needed to justify these approaches. We also show how we can, by the right conditioning, estimate many of the model parameters without the need to solve a dynamic programming problem.

We cannot over-emphasize the importance of allowing for heterogeneity, across consumers, firms, products and markets, in order to explain micro data. Not accounting for this heterogeneity can generate significant biases in parameter estimates and in our understanding of competition between firms. For instance, in the estimation of dynamic games of oligopoly competition, ignoring unobserved market heterogeneity when present can lead to serious biases in our estimates of the degree of strategic interaction between firms. Unfortunately, some of the methods used to reduce the state space and ease the computational burden limit the ability to estimate observed and unobserved heterogeneity. This at times creates a trade-off between estimation methods that are faster, and potentially allow for the estimation of models that are richer in observed variables and have more flexible parametric forms, and methods that can handle only simpler models but can allow for richer unobserved heterogeneity. Interestingly, the two literatures we survey have taken somewhat different approaches in handling this trade-off. Using our examples, we highlight the trade-offs and ways they have been addressed.

Multiple equilibria is a prevalent feature in dynamic games. We focus on several of the practical problems this multiplicity introduces in estimation and prediction. A key way the literature has dealt with multiple equilibria is to assume that a unique game is being played in the data. One potential issue is whether this equilibrium is stable, in the sense we define below. As we show below, it turns out this has important implications for the performance of many common methods. In addition we review recent methods for the estimation of dynamic games that can deal with multiple equilibria and unobserved heterogeneity. We also examine the implementation of counterfactual experiments in models with multiple equilibria.

We have organized the paper in two parts. Section 2 deals with dynamic models of demand of differentiated products, and Section 3 with dynamic games of oligopoly competition.

2 Dynamic Demand for Differentiated Products

2.1 Overview

Over the last thirty or so years, demand estimation has been a key part of studies in empirical Industrial Organization. The key idea is to estimate demand and use the estimates to recover unobserved costs by inverting a pricing decision. Once cost has been recovered, the estimated demand and cost can be used to study the form of competition, understand firm behavior, generate counterfactuals (e.g., the likely effect of a merger) or quantify welfare gains (for example, from the introduction of new products.)¹

Much of the literature has relied on static demand models for this type of exercise. However, in many markets demand is dynamic in the sense that (a) consumers current decisions affect their future utility, or (b) consumers' current decisions depend on expectations about the evolution of future states. The exact effect of dynamics differs depending on the circumstances, and can be generated for different reasons. The literature has focused on several cases including storable products, durable products, habit formation, switching costs and learning. As our goal is to demonstrate key challenges faced by empirical researchers, and not to provide a complete survey, we focus on the first two cases: storable and durable products.²

In the case of storable products, if storage costs are not too large and current price is low relative to future prices (i.e., the product is on sale), there is an incentive for consumers to store the product and consume it in the future. Dynamics arise because consumers' past purchases and consumption decisions impact their current inventory and therefore could impact both the costs and benefits of purchasing today. Furthermore, consumers expectations about future prices, and availability of products, also impact the perceived trade-offs of buying today versus in the future.

In the case of durable products, dynamics arise due to similar trade-offs. The existence of transaction costs in the resale market of durable goods (for example, because of adverse selection, Akerlof, 1973) implies that a consumer's decision today of whether or not to buy a durable good, and which product to buy, is costly to change in the future and, for that reason, it will impact her future utility. Therefore, when a consumer makes a purchase, she is influenced by her current holdings of the good and by her expectations about future prices and attributes of available products. For

¹See, for example, Bresnahan (1981, 1987), Porter (1983), Hausman, Leonard and Zona (1994), Berry, Levinsohn and Pakes (1995), Goldberg (1995), Nevo (2001).

²As we noted in the Introduction, our goal is not to provide a complete survey, so we will not offer a comprehensive discussion of this wide literature. For some examples, in addition to the papers we discuss below, see Hartmann (2006), Carranza (2006), Esteban and Shum (2007), Nair (2007), Rossi (2007), Shcherbakov (2008), Sweeting (2008), Lou, Prentice and Yin (2008), Lee (2009), Osborne (2009), Perrone (2009), Schiraldi (2010), as well as many others.

instance, a consumer that currently owns a one year old car is likely to make a different purchasing decision than an identical consumer who owns a ten year old car. The dynamics are most important in industries where prices and available products are changing rapidly over time, such as many consumer goods, or where there are policies that have dynamic effects, such as scrapping subsidies in the automobile industry.

Ignoring the dynamics and using the data to estimate a static demand model generates biased and inconsistent estimates. Besides the econometric bias it is important to realize that in many cases static estimation does not recover desired quantities and thus fails to address many interesting questions. For example, in many applications it is important to separate between a short run price elasticity, in response to a temporary price change, and a long run elasticity in response to a permanent price change. In general, due to econometric bias, static estimation does not recover short run responses, but even if it does, in some very special cases, it cannot separately recover the long run response.

Computing price responses is obviously important to fields like IO and marketing but the possible uses of the models discussed below are much wider and include many fields in Economics. Here are a few examples. Recently, macro economists have looked at micro level price data to study price rigidities.³ A central issue in this literature is how to treat temporary price reductions, or "sales". A key to understanding sales, and why they exist, is to understand consumer response. Similarly, a key issue in trade is the pass through of exchange rates. Here again separating between short run and long run price responses is critical. In another example, obesity and unhealthy eating habits are plaguing many countries and have led to suggestions of taxing unhealthy high fat and high calories foods. To evaluate the effectiveness of these policies, it is crucial to estimate the heterogeneity in price response: if a tax was to be imposed who responds and by how much. Furthermore, it is probably important to estimate the degree of habit persistence in the consumption of these unhealthy food products. Adoption of energy efficient cars and appliances is an important aspect of environmental economics. To the extent that demand is dynamic, as discussed above, modeling the dynamics is crucial. Modeling the dynamics of durable goods purchases has important implications for evaluating scrapping policies and computing price indices.

The dynamic factors impacting demand have long been recognized⁴ and indeed many different

³See, for example, Kehoe and Midrigan (2008), Eichenbaum, Jaimovich and Rebelo (2008), and Nakamura and Steinsson (2009).

⁴For an early contribution see General Motors Corporation (1939), a volume developed out of papers sponsored by General Motors and presented in a joint session of the American Statistical Association and the Econometric Society in 1938

models to capture these dynamics have been offered in the literature ranging from models where the dynamics decisions are modeled explicitly to modeling approaches where the dynamics are handled by including lags and leads of variables (e.g., prices). The IO literature has mostly taken the approach of explicit modeling, often referred to as a "structural" approach.

In order to implement these approaches in markets with differentiated products and address important applied questions researchers have had to deal with several issues including large state spaces, unobserved (endogenous) state variables and heterogeneity. In this section we survey the approaches taken to deal with these issues.

2.1.1 Background: Static Demand for Differentiated Products

Key lessons to learn from static demand estimation is the importance of allowing for heterogeneity and the difficulty of dealing with the dimensionality of the problem, while still allowing for flexible enough functional forms. Consider a classical (static) demand system for J products, $\mathbf{q} = D(\mathbf{p}; \mathbf{r})$, where \mathbf{q} is a J -dimensional vector of quantities demanded, \mathbf{p} is a J -dimensional vector of prices, and \mathbf{r} is a vector of exogenous variables. A key problem in estimating this system is the dimensionality – due to the large number of products the number of parameters is too large to estimate. Several solutions have been offered in the literature,⁵ but the most common solution in the IO literature is to rely on a discrete choice model (McFadden 1974, Berry, Levinsohn, and Pakes, 1995).

The work-horse discrete choice model used in IO has a consumer i choosing option j from one of $J + 1$ options (J brands and a no purchase option). The (conditional indirect) utility the consumer gets from option j at time t is given by

$$u_{ijt} = a_{jt} \beta_i - \alpha_i p_{jt} + \xi_{jt} + \varepsilon_{ijt} \quad (1)$$

where p_{jt} is the price of option j at time t , a_{jt} is a $1 \times K$ vector of observable attributes of product j , ξ_{jt} is an unobserved (by the econometrician) product characteristic, ε_{ijt} is a stochastic term, α_i represents the consumer's marginal utility of income, and β_i is a $K \times 1$ vector of individual-specific marginal utilities associated to the attributes in the vector a_{jt} . In this model a product is viewed as a bundle of characteristics and therefore the relevant dimension is the number of characteristics, K , and not the number of products. Flexible substitution patterns are achieved by allowing for consumer heterogeneity in the willingness to pay and in the valuation of characteristics.

⁵For example, a common approach in the trade and macro literature is to use the Constant Elasticity of Substitution (CES) demand system, which is very economical on parameters. This model, however, is not flexible enough to explain micro level data. An alternative approach is to use the multi-level demand system developed by Hausman, Leonard and Zona (1994). See Nevo (2011) for a discussion.

The model can be estimated using consumer-level data. However, the wider availability of market level data and the development of appropriate econometric techniques have made estimation using market level data the more popular choice. The estimates from aggregate level data are generally considered more credible if the data come from many different markets with variation in the observed attributes of consumers, or the aggregate data is supplemented with so called micro-moments, basically the purchasing patterns of different demographic groups.

2.1.2 Dynamic Demand: Key Ingredients

In building dynamic demand models, the IO literature has continued to rely heavily on the discrete choice model.

Reducing the dimensionality If demand is dynamic the dimensionality problem is even worse. The basic idea of a discrete choice model – to project the products onto a characteristics space – that essentially solved the problem in the static context is not sufficient in the dynamic context. For example, consider the problem of a forward looking consumer trying to form expectations about future price and characteristics of products. In principle, this consumer needs to form expectations about the future $K + 1$ attributes of all products, the number of which could be changing, using the information of the current, and past, values of these attributes for all the products. Even if we assume that variables follow a first order Markov process, and that the number of products is fixed, the size of the state space is $(K + 1) * J$.

A very useful concept, which is used in the examples below is the *inclusive value*. McFadden (1978) defines the inclusive value (or *social surplus*) as the expected utility of a consumer, from several discrete options, prior to observing $(\varepsilon_{i0t}, \dots, \varepsilon_{iJt})$, knowing that the choice will be made to maximize utility. When the idiosyncratic shocks ε_{ijt} are distributed i.i.d. extreme value, the inclusive value from a subset $A \subseteq \{1, 2, \dots, J\}$ of the choice alternatives is defined as:

$$\omega_{it}^A = \ln \left(\sum_{j \in A} \exp \{ a_{jt} \beta_i - \alpha_i p_{jt} + \xi_{jt} \} \right) \quad (2)$$

When $\beta_i = \beta$ and $\alpha_i = \alpha$ the inclusive value captures the average utility in the population, up to a constant, averaging over the individual draws of ε , hence the term *social surplus*.

The inclusive value plays a key role in reducing the state space. In forming expectations the consumer just has to form expectations about the future inclusive value, or in some cases a low number of inclusive values for subsets of products, rather than expectations about the realizations of all attributes of all products. In order to reduce significantly the state space, this property

is coupled with a behavioral assumption on the information that consumers use to form these expectations.

Heterogeneity Just like static models, allowing for heterogeneity is key to explain the data and retain flexible demand systems. In some cases, however, some degree of unobserved heterogeneity needs to be sacrificed in order to deal with the dimensionality problem. As we show below, the trade-off in some cases is between a richer model that includes more observed heterogeneity and a model that relies on unobserved heterogeneity.

Data Just like the static model, the dynamic model can be estimated using consumer level or market level data. The advantages of consumer level data seem more obvious in the dynamic setting: consumer-level data allow us to see how individual consumers behave over time. However, this is exactly the reason why consumer level data sets are hard to collect, especially for products, such as some durables, that are purchased infrequently. For this reason, a number of applications have relied on aggregate data. We (informally) discuss identification and estimation with market level data.

2.2 Storable products

Many of the products purchased by consumers are storable so consumers can buy them for future consumption. A typical pricing pattern in these markets involves short lived price reductions with a return to the regular price. This pattern of prices generates an incentive for consumers to store the product when the price is low. Boizot, Robin, and Visser (2001) and Pesendorfer (2002) were among the first to study the effects of temporary price reductions and storability in Economics.⁶

2.2.1 Evidence

There is ample evidence that once faced with temporary price reductions consumers store for future consumption. For example, Pesendorfer (2002) using data for ketchup finds that holding the current price constant aggregate quantity sold depends on duration from previous sale. Hendel and Nevo (2006a, 2010), find similar evidence for other products.

Additional evidence for the existence of demand accumulation is provided by Hendel and Nevo (2006a) who use household level data to document patterns that are consistent with consumer

⁶An earlier Marketing literature examined some of the same issues, but the treatment there was generally not consistent with optimal dynamic behavior. See for example, Shoemaker (1979), Blattberg, Eppen and Lieberman (1981), Gupta, (1988), and Chiang (1991).

stockpiling behavior. For example, they show that the household's propensity to purchase on sales is correlated with proxies of storage costs, and that households in areas where houses are larger (with cheaper storage) buy more on sale. They also show that when purchasing on sale, duration to next purchase is longer. This is true both within households – for a given household when buying on sale the duration is longer – and across households – households who purchase more on sale also purchase less frequently. Finally, proxies for inventory are negatively correlated with quantity purchased and the probability of purchasing.

2.2.2 Implications

Given the evidence on demand accumulation, it is natural to ask what are the implications. The primary implication is for demand estimation, which in itself is an input for addressing important economic questions that we have discussed above.

Once we recognize that consumers can store the product, we need to separate between the short run response, to a temporary price change, and the long run response to either a temporary or permanent price change. For most economic applications we care about long run changes. If price changes in the data are permanent, then static estimation yields consistent estimates of the long run demand responses. Indeed, one way to estimate long run responses is to only use permanent price changes and ignore, to the extent possible, the temporary price changes. In many data sets the temporary price change constitute most or even all of the variation in prices. Therefore, dropping these price changes means a significant loss of efficiency, possibly even completely wiping out any price variation.

On the other hand, if price changes in the data are temporary, then static demand estimates over-estimate own price effects. The (large) demand response to a sale is attributed to an increase in consumption (which in a static model equals purchase), and not to an increase in storage. The decline in purchases after a sale coincides with an increase in price, and is mis-attributed as a decline in consumption. At the same time, static estimation under-estimates cross price effects. During a sale the quantity sold of competing products goes down, but static estimation misses an additional effect: the decrease in the quantity sold in the future. Intuitively, when a competing product was on sale in the past, consumers purchased to consume today and therefore, the relevant, or "effective," cross price is not the current cross price. The current price is (weakly) higher. Furthermore, when a (cross) product is on sale the current (cross) price is more likely to be the effective price. Both these effects bias the estimated cross price effect towards zero.

2.2.3 A Model of Consumer Stockpiling

Hendel and Nevo (2006b) propose the following model of consumer stockpiling, which we use to demonstrate some of the key issues faced by applied researchers.

The starting point is similar to the discrete choice model discussed above in Section 2.1.1. The consumer can purchase one of $J + 1$ brands which come in different sizes, which we index by $x \in \{1, 2, \dots, X\}$. Let d_{jxt} equal to 1 if the consumer purchases brand j of size x at time t , and 0 otherwise. Since the choice is discrete, stockpiling is achieved by buying larger sizes, and adding to existing inventory, rather than by buying multiple units on any given shopping trip. This assumption seems reasonable for the data used by Hendel and Nevo (2006b), where there were few purchases of multiple units. In other context this might not be reasonable and one would need to model the choice of multiple units.

The consumer also has to decide how much to consume each period.⁷ The per period utility consumer i obtains from consuming in t is

$$u_i(\mathbf{c}_t, \boldsymbol{\nu}_t) + \alpha_i m_t \quad (3)$$

where \mathbf{c}_t is a J dimensional vector of the quantities consumed of each brand, $\boldsymbol{\nu}_t$ is a J dimensional vector of shocks to utility that change the marginal utility from consumption and m_t is the utility from the outside good. In addition to utility from consumption, a consumer one-period utility has two other components. We assume that the consumer pays a cost $C_i(\mathbf{i}_t)$ for holding inventories \mathbf{i}_t , where \mathbf{i}_t is a vector of inventories by brand. There is also an instantaneous utility associated with preference for the purchased brand. At period $t = 1$, the purchase and consumption decisions, $\{\mathbf{c}, j, x\}$, are made to maximize

$$\begin{aligned} & \sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E}[u_i(\mathbf{c}_t, \boldsymbol{\nu}_t) - C_i(\mathbf{i}_t) + a_{jxt}\beta_i - \alpha_i p_{jxt} + \xi_{jxt} + \varepsilon_{ijxt} \mid \mathbf{s}_1] \\ \text{s.t. } & 0 \leq \mathbf{i}_t, \quad 0 \leq \mathbf{c}_t \quad \sum_{j,x} d_{jxt} = 1, \quad i_{j,t+1} = i_{j,t} + \sum_x d_{jxt}x_t - c_{j,t} \quad j = 1, \dots, J \end{aligned} \quad (4)$$

where \mathbf{s}_t is the information set at time t , δ is the discount factor, p_{jxt} is the price of purchasing quantity x of brand j , ξ_{jxt} is an unobserved (to the researcher) brand specific quality, a_{jxt} are observed product attributes and ε_{ijxt} is a random shock. We allow ξ_{jxt} to vary by brand in order

⁷An alternative of assuming that consumption is constant over time, but varying across households, seems attractive, especially for the type of products usually modeled. A slightly more general model than constant consumption allows for random shocks, \mathbf{v}_t , that determine consumption. Both these models are nested within our model and in principle can be tested. The results in Hendel and Nevo (2006b) suggest that consumption is mostly constant, but when inventory runs low consumers reduce consumption. This behavior is required to explain long periods of no purchase followed by periods of frequent purchases observed in the data. Indeed, it is this variation in inter-purchase time that identifies the utility from consumption.

to capture differentiation across products, and across sizes for reasons we discuss below. In principle, the brand preference can also vary across consumers.

The expectation $\mathbb{E}(\cdot)$ is taken with respect to the uncertainty regarding future shocks in the vectors $\boldsymbol{\nu}_t$ and $\boldsymbol{\varepsilon}_t$, and future prices (and other time varying attributes). We assume that ε_{jxt} is iid extreme value and that $\boldsymbol{\nu}_t$ is iid over time and across consumers with a known parametric distribution. Prices (and observed characteristics) evolve according to a first order Markov process.

Some aspects of the specification of this consumer decision problem deserve further explanation. First, we assume no physical depreciation of the product, though this assumption is easy to relax if needed. Second, we assume that a decision is made each period with perfect knowledge of current prices. Implicitly, we are assuming that the consumer visits the store every period. This assumption also helps us in the specification of consumer expectations regarding future prices. If consumers do not visit the store every period, we have to model the process by which they arrive at the store to determine the next set of prices they should expect.

At the moment, even with the simplifying assumptions already made, the vector of state variables is quite large and includes a J -dimensional vector of inventory holdings by brand, \mathbf{i}_t ; a $(K+1) * J * X$ -dimensional vector of prices and characteristics, \mathbf{p}_t ;⁸ a J -dimensional vector of consumption shocks, $\boldsymbol{\nu}_t$; and $J * X$ -dimensional vector of iid extreme value shocks, $\boldsymbol{\varepsilon}_t$. The vector of state variables at period t is $\mathbf{s}_t = (\mathbf{i}_t, \mathbf{p}_t, \boldsymbol{\nu}_t, \boldsymbol{\varepsilon}_t)$. Without a first order Markov assumption the state space would be even larger and would include additional lags of prices and characteristics.

Let $V_i(s_t)$ be the value function of consumer i . As usual in a dynamic programming problem, this value function can be obtained as the unique solution of a Bellman equation.

$$V_i(\mathbf{s}_t) = \max_{\{\mathbf{c}, j, x\}} \left\{ u_i(\mathbf{c}, \boldsymbol{\nu}_t) - C_i(\mathbf{i}_t) + a_{jxt}\beta_i - \alpha_i p_{jxt} + \xi_{jxt} + \varepsilon_{ijxt} + \delta \int V_i(\mathbf{s}_{t+1}) dF_s(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{c}, j, x) \right\} \quad (5)$$

where F_s represents the transition probability of the vector of state variables. Given that the state variables $(\boldsymbol{\nu}_t, \boldsymbol{\varepsilon}_t)$ are independently distributed over time, it is convenient to reduce the dimensionality of this dynamic programming problem by using a value function that is integrated over these iid random variables. The *integrated value function*, sometimes also called the ex-ante value function, is defined as $EV_i(\mathbf{i}_t, \mathbf{p}_t) \equiv \int V_i(\mathbf{s}_t) dF_\varepsilon(\boldsymbol{\varepsilon}_t) dF_\nu(\boldsymbol{\nu}_t)$, where F_ε and F_ν represent the CDFs of $\boldsymbol{\varepsilon}_t$ and $\boldsymbol{\nu}_t$, respectively. The value function EV_i is the unique solution of the integrated Bellman equation. Given the distributional assumptions on the shocks $\boldsymbol{\varepsilon}_t$ and $\boldsymbol{\nu}_t$, the integrated Bellman

⁸To keep notation simple we use \mathbf{p}_t to denote the observed variables at time t . These variables include prices and other observed variables.

equations is:

$$EV_i(\mathbf{i}_t, \mathbf{p}_t) = \max_{\mathbf{c}, x} \int \ln \left(\sum_j \exp \left\{ \begin{array}{l} u_i(\mathbf{c}, \boldsymbol{\nu}_t) - C_i(\mathbf{i}_t) + a_{jxt}\beta_i - \alpha_i p_{jxt} + \xi_{jxt} \\ + \delta \mathbb{E}[EV_i(\mathbf{i}_{t+1}, \mathbf{p}_{t+1}) \mid \mathbf{i}_t, \mathbf{p}_t, \mathbf{c}, j, x] \end{array} \right\} \right) dF(\boldsymbol{\nu}_t). \quad (6)$$

The main computational cost is to compute the functions EV_i . We now explore ways to reduce this cost.

2.2.4 Reducing the Dimension of the State Space

As it stands the state space is quite large and not workable for anything except very small number of products J . In order to reduce the state space, several additional assumptions are needed.

Inventories and Consumption We first explore ways to reduce the dimension of inventories needed to keep track of. One possible assumption is to assume products are perfect substitutes in consumption and storage

Assumption A1: $U_i(\mathbf{c}_t, \mathbf{v}_t) = U_i(c_t, \nu_t)$ and $C_i(\mathbf{i}_t) = C_i(i_t)$ where $c_t = \mathbf{1}'\mathbf{c}_t$, $\nu_t = \mathbf{1}'\boldsymbol{\nu}_t$, $i_t = \mathbf{1}'\mathbf{i}_t$, and $\mathbf{1}$ is a vector of 1's.

Under this assumption the inventory and the consumption shocks reduce to a scalar: we only need to keep track of a single inventory and a single consumption shock. Formally, now

$$EV_i(\mathbf{i}_t, \mathbf{p}_t) = EV_i(i_t, \mathbf{p}_t) \quad (7)$$

This assumption not only reduces the state space but, as we see below, it also allows us to modify the dynamic problem, which can significantly aid in the estimation of the model.

Taken literally, this assumption implies that there is no differentiation in consumption: the product is homogenous in use. Note, that through ξ_{jxt} and ε_{ijxt} we allow differentiation in purchase, as is standard in the IO literature. Indeed, it is well known that this differentiation is needed to explain purchasing behavior. This seemingly creates a tension in the model: products are differentiated at purchase but not in consumption. Before explaining how this tension is resolved we note that the tension is not only in the model but potentially in reality as well. Many products seem to be highly differentiated at the time of purchase but its hard to imagine that they are differentiated in consumption. For example, households tend to be extremely loyal to the laundry detergent brand they purchase – a typical household buys only 2-3 brands of detergent over a very long horizon – yet its hard to imagine that the usage and consumption are very different for different brands. One way to think of the model is to assume that there is a brand-specific utility

in consumption. As long as the utility in this component is linear and we can ignore, to a first order, discounting, then the brand specific utility in consumption is captured by ξ_{jxt} .⁹ This is the reason we want to let ξ_{jxt} vary by size. Indeed, the above suggests that $\xi_{jxt} = \xi_{jt} * x$.

Assumption A1 implies that the optimal consumption does not depend on which brand is purchased. Formally, let $c_k^*(\mathbf{s}_t; x, k)$ be the optimal consumption of brand k conditional on state \mathbf{s}_t and on purchase of size x of that brand. Lemma 1 in the appendix of Hendel and Nevo (2006b) shows that $c_k^*(\mathbf{s}_t; x, k) = c_j^*(\mathbf{s}_t; x, j) = c^*(\mathbf{s}_t; x)$. In words, the optimal consumption does not depend on the brand purchased, only on the size.

This result implies that the (integrated) Bellman equation in (6) can be written as:

$$EV_i(i_t, \mathbf{p}_t) = \max_{c,x} \int \ln \left(\sum_x \exp \{u_i(c_t, v_t) - C_i(i_t) + \omega_{ixt} + \delta \mathbb{E}[EV_i(i_{t+1}, \mathbf{p}_{t+1}) \mid i_t, \mathbf{p}_t, c, x]\} \right) dF_\nu(v_t) \quad (8)$$

where ω_{ixt} is the inclusive value from all brands of size x , as defined by equation (2), i.e., $\omega_{ixt} = \ln \left(\sum_j \exp(a_{jxt}\beta_i - \alpha_i p_{jxt} + \xi_{jxt}) \right)$. In words, the problem can now be seen as a choice between sizes, each with a utility given by the size-specific inclusive value (and extreme value shock). The dimension of the state space is still large and includes all prices, because we need all the prices to compute the evolution of the inclusive value. However, in combination with additional assumptions the modified problem is easier to estimate.

Finally, we note that if needed we could reduce the inventory to several types of products, rather than to a scalar. For example, suppose we are studying the breakfast cereal market, we could split the brands into kids cereal and adult cereals, such that within a group products are perfect substitutes. In this case we would need to keep two inventories – for adult and kids cereal – still significantly less than the number of brands.

Prices As we noted, even with Assumption A1 the state space is still large and includes all prices. Therefore, for a realistic number of products, the state space is still too large to be manageable. To further reduce it, we make an additional assumption (Assumption A4 in Hendel and Nevo, 2006b). Let $\boldsymbol{\omega}_{it}$ be a vector of inclusive values for the different sizes.

Assumption A2: $F(\boldsymbol{\omega}_{i,t+1} \mid \mathbf{s}_t) = F(\boldsymbol{\omega}_{i,t+1} \mid \boldsymbol{\omega}_{it}(\mathbf{p}_t))$

In words, the vector $\boldsymbol{\omega}_{it}$ contains all the relevant information in \mathbf{s}_t to obtain the probability distribution of $\boldsymbol{\omega}_{i,t+1}$ conditional on \mathbf{s}_t . Instead of all the prices (and attributes) we only need a

⁹See Hendel and Nevo (2006a) for the details of the argument. Erdem, Imai and Keane (2003) offer an alternative model that allows for two inventories. One can show that under the above assumptions their model is a private case of the one discussed here.

single index for each size. Two vectors of prices that yield the same (vector of) current inclusive values imply the same distribution of future inclusive values. This assumption is violated if individual prices have predictive power above and beyond the predictive power of ω_{it} . Therefore, if the inclusive values can be estimated outside the dynamic demand model, the assumption can be tested and somewhat relaxed by including additional statistics of prices in the state space. Note, that ω_{it} is consumer specific: different consumers value a given set of products differently and therefore this assumption does not further restrict the distribution of heterogeneity.

Given Assumptions A1 and A2 we can show (see Hendel and Nevo, 2006b) that

$$EV_i(i_t, \mathbf{p}_t) = EV_i(i_t, \omega_{it}(\mathbf{p}_t)) \quad (9)$$

In words, the expected future value only depends on a lower dimensional statistic of the full state vector.

2.3 Estimation

In this section we discuss the identification and estimation of the model. We assume that the researcher has access to consumer level data. Such data is widely available from several data collection companies and recently researchers in several countries have been able to gain access to such data for academic use.¹⁰ The data include the history of shopping behavior of a consumer over a period of one to three years. The researcher knows whether a store was visited, if a store was visited then which one, and what product (brand and size) was purchased and at what price. In many cases the hardest information to gather are the prices of products not purchased. From the view point of the model, the key information that is not observed is consumer inventory and consumption decisions.

The most straightforward way to estimate the model follows a similar algorithm to the one suggested by Rust (1987).¹¹ For a given set of parameters we solve the dynamic programming problem and obtain (deterministic) decision rules for purchases and consumption as a function of the state variables including the unobserved random shocks. Assuming a distribution for these shocks we derive a likelihood of observing each consumer's decision conditional on prices and inventory.

¹⁰See for example the ERIM data available at <http://research.chicagobooth.edu/marketing/databases/erim/index.aspx>, or the so called Stanford Basket described in Bell and Lattin (1998). For more recent datasets, see, for example, Griffith, Leicester, Leibtag and Nevo (2009) for a use of UK data; Einav, Leibtag and Nevo (2010) for US data; Bonnet and Dubois (2010) for French data; and Browning and Carro (2006) for Danish data.

¹¹For computational reasons methods based on conditional choice probabilities (Hotz and Miller, 1993, Hotz, Miller, Sanders and Smith, 1994, Aguirregabiria and Mira, 2002) have become quite popular. Since the model includes unobserved endogenous time varying state variables these methods cannot be directly applied here. However, the method of Arcidiacono and Miller (2009) could potentially be applied to the estimation of this model. See Section 3 for further discussion of these methods.

We nest this computation of the likelihood into the search for the values of the parameters that maximize the likelihood of the observed sample.

We face two hurdles in implementing the algorithm. First, consumption (a decision variable) and inventory (a state variable) are not observed. As we show below, this can be solved by using the model to derive the optimal consumption and the implied inventory. The second problem is the dimensionality of the state space. We discussed several assumptions that can be used to reduce the state space. Nevertheless, the computational problem is still quite difficult. We show how the computation can be significantly simplified by splitting the estimation into estimation of the brand choice conditional on size, which does not require solving the dynamic problem, and then estimating the choice of size, which requires solving a much simpler dynamic problem.

For the purpose of inference, since in some specification we want to allow for household fixed effects, we usually need to assume that the number of observations per household is very large.

As we noted above, its quite common in the IO literature to estimate static demand models using market level data. We are unaware of any paper that has tried to estimate the model we propose here using aggregate data. Hendel and Nevo (2010) estimate a simpler model using aggregate data.

2.3.1 Identification

Before discussing the details of estimation, we informally discuss identification. If inventory and consumption were observed, then identification using consumer level data follows standard arguments (see Rust, 1994 and Magnac and Thesmar, 2002, Aguirregabiria, 2010). However, we do not observe inventory or consumption so the question is which features of the data allow us to identify functions of these variables?

The same correlations and patterns we described in Section 2.2.1 to suggest that dynamics are relevant are the ones that identify the dynamic model. In particular, the individual level data provide the probability of purchase conditional on current prices, and past purchases of the consumer (amounts purchased and duration from previous purchases). Suppose that we see that this probability is not a function of past behavior, we would then conclude that dynamics are not relevant and that consumers are purchasing for immediate consumption and not for inventory. On the other hand, if we observe that the purchase probability is a function of past behavior, and we assume that preferences are stationary then we conclude that there is dynamic behavior.¹² Regarding the identification of storage costs, consider the following example. Suppose we observe

¹²Serial correlation in v_t might also generate a dependence of the purchase probability on past behavior. However, positive serial correlation in v_t generates positive dependence between past and current purchases, while the stockpiling model generates negative dependence between past and current purchases.

two consumers who face the same price process and purchase the same amount over a given period. However, one of them purchases more frequently than the other. This variation leads us to conclude that this consumer has higher storage costs. Therefore, the storage costs are identified from the average duration between purchases. The utility from consumption is identified from the variation in these duration times, holding the amount purchased constant. For example, a model of constant consumption cannot explain large variation in the duration times.

In some cases the researcher might not have consumer level, but only store or market level data. We are unsure if the model presented here is identified from aggregate data. Given the above discussion it might seem unlikely. However, a slightly simpler dynamic demand model for storable goods can be identified from aggregate store level data, as long as the aggregation corresponds to the timing of price changes (if we have weekly data we need the prices to be constant within the week). The variation in the data that identifies the model is dependence of total quantity sold on the duration from last sale. See Hendel and Nevo (2010) for details.

A key emphasis in static demand estimation is the potential endogeneity of prices. The concern is that prices, and sometimes other variables, are correlated with ξ_{jxt} . On the other hand, some researchers estimating dynamic demand have brushed this concern aside saying that papers that focus on endogeneity "missed the mark" (Erdem, Imai and Keane, 2003, pg 11). In our view, whether or not one should be concerned about endogeneity depends on the data structure, what is included in the model, and the institutional knowledge of the industry. Broad statements like endogeneity is not an issue in dynamic models or when using consumer level data, are generally not correct. In the estimation below we deal with endogeneity by (1) assuming that $\xi_{jxt} = \xi_{jx}$, i.e., does not vary over time, and control for it with fixed effects and (2) by using the simplified computational problem to control for time varying variables like advertising and promotions. In our discussion of durable goods we review a GMM method that closely follows the static demand estimation.

2.3.2 Estimation

The parameters of the model can be estimated via maximum likelihood following a similar algorithm to Rust (1987).¹³ Since inventory, one of the state variables, is not observed we need to impute it as part of the estimation. This can be done in the following way:

- (i) Guess an initial inventory distribution and draw from it for each consumer;

¹³This is subject to the caveat regarding the endogeneity of prices. See the discussion in the previous section.

- (ii) For a given value of the parameters solve the consumer problem and obtain the value and policy functions;
- (iii) Using the draws of inventory (from (i)), the computed consumption policy (from (ii)) and observed purchases obtain the sequence of inventory, and compute likelihood of the observed purchases;
- (iv) Repeat steps (ii) and (iii) to choose the parameters that maximize the likelihood of the observed data, possibly leaving out some of the initial observations to let the inventory process settle.
- (v) Update the initial guess of the distribution of inventory and repeat steps (i)-(iv).

The likelihood, in step (iii), of observing a sequence of purchasing decisions, $(\mathbf{d}_1, \dots, \mathbf{d}_T)$, as a function of the observed state variables, $(\mathbf{p}_1, \dots, \mathbf{p}_T)$ and observed demographic variables D_i is

$$\mathbb{P}(\mathbf{d}_1, \dots, \mathbf{d}_T | \mathbf{p}_1, \dots, \mathbf{p}_T, D_i) = \int \prod_{t=1}^T \mathbb{P}(\mathbf{d}_t | \mathbf{p}_t, i_t(\mathbf{d}_{t-1}, \dots, \mathbf{d}_1, v_{t-1}, \dots, v_1, i_1), v_t, D_i) dF(v_1, \dots, v_T) dF(i_1) \quad (10)$$

Inventory is a function of previous observed purchase (or no purchase) decisions, the previous consumption shocks and the initial inventory. The exact functional form of the dependence of inventory on past consumption shocks depends on the consumption policy. The probability inside the integral represents the integration over the set of epsilons that induce \mathbf{d}_t as the optimal choice. Using Assumptions A1 and A2 and the results from Section 2.2.4, this probability is given by

$$\mathbb{P}(j, x | \mathbf{p}_t, i_t, v_t, D_i) = \frac{\exp \{a_{jxt}\beta_i - \alpha_i p_{jxt} + \xi_{jx} + \max_c (u_i(c, v_t) - C_i(i_t) + \delta \mathbb{E} [EV_i(i_{t+1}, \boldsymbol{\omega}_{t+1}) | i_t, \boldsymbol{\omega}_t, c, j, x])\}}{\sum_{k,y} \exp \{a_{kyt}\beta_i - \alpha_i p_{kyt} + \xi_{ky} + \max_c (u_i(c, v_t) - C_i(i_t) + \delta \mathbb{E} [EV_i(i_{t+1}, \boldsymbol{\omega}_{t+1}) | i_t, \boldsymbol{\omega}_t, c, k, y])\}} \quad (11)$$

Hence, to compute the likelihood we only need to solve the dynamic problem in the reduced state space.

Splitting the Likelihood Its important to note that up to this point we used the stochastic structure of the problem, but we did not restrict the distribution of consumer heterogeneity. In particular, we can allow for the taste coefficients, α_i and β_i , to vary with both observed and unobserved factors, and estimate their distribution using the above joint likelihood of brand and size choice.

We now show that if we are willing to place some restrictions on the unobserved heterogeneity, we can significantly simplify the computational problem.

As we discussed above the optimal consumption is not brand specific so $M_i(\boldsymbol{\omega}_t, i_t, \nu_t, y) = \max_c (u_i(c, \nu_t) - C_i(i_t) + \delta \mathbb{E}[EV_i(i_{t+1}, \boldsymbol{\omega}_{t+1}) | i_t, \boldsymbol{\omega}_t, c, j, y])$ does not vary by brand j , conditional on a size y . Thus, we note that the above probability can be written as

$$\begin{aligned} \mathbb{P}(j, x | \mathbf{p}_t, i_t, \nu_t, D_i) &= \frac{\exp\{a_{jxt}\beta_i - \alpha_i p_{jxt} + \xi_{jx}\}}{\sum_k \exp\{a_{kxt}\beta_i - \alpha_i p_{kxt} + \xi_{kx}\}} \frac{\exp\{\omega_t^x + M(\boldsymbol{\omega}_t, i_t, \nu_t, x)\}}{\sum_y \exp\{\omega_t^y + M(\boldsymbol{\omega}_t, i_t, \nu_t, y)\}} \quad (12) \\ &= \mathbb{P}(j|x, \mathbf{p}_t, D_i) \mathbb{P}(x|\boldsymbol{\omega}_t, i_t, D_i) \end{aligned}$$

In order for this factorization to be useful in reducing the computational cost we need a conditional independence assumption:

Assumption A3 (conditional independence of heterogeneity): $F(\alpha_i, \beta_i | x_t, \mathbf{p}_t, D_i) = F(\alpha_i, \beta_i | \mathbf{p}_t, D_i)$ where x_t represents the chosen size.

This assumption is satisfied if heterogeneity is only a function of observed demographics, including possibly "fixed effects." If this assumption holds then:

$$\mathbb{P}(j|x_t, \mathbf{p}_t, D_i) = \int \mathbb{P}(j|x_t, \mathbf{p}_t, \alpha_i, \beta_i) dF(\alpha_i, \beta_i | x_t, \mathbf{p}_t, D_i) = \int \mathbb{P}(j|x_t, \mathbf{p}_t, \alpha_i, \beta_i) dF(\alpha_i, \beta_i | \mathbf{p}_t, D_i).$$

On the other hand, if the assumption does not hold we need to compute $F(\alpha_i, \beta_i | x_t, \mathbf{p}_t, D_i)$, which, in general, requires us to solve the dynamic programming problem.

To illustrate what Assumption A3 rules out consider the following example. Suppose there are two brands, A and B , offered in two sizes, L and S . There are two types of consumers each with equal mass. Type a prefer brand A , while type b prefer brand B . Suppose brand A goes on sale in size L , but not size S . Now consider the conditional choice probabilities

$$\mathbb{P}(A|L, \mathbf{p}_t) = \mathbb{P}(A|L, \mathbf{p}_t, a) \mathbb{P}(a|L, \mathbf{p}_t) + \mathbb{P}(A|L, \mathbf{p}_t, b) \mathbb{P}(b|L, \mathbf{p}_t)$$

Unconditionally, $\mathbb{P}(a) = \mathbb{P}(b) = 0.5$. But since brand A size L was on sale its likely that conditional on purchasing size L the mass of type a is higher than the mass of type b . To figure how much higher, we need to compute for each type the probability that they purchase size L . In general this requires solving the dynamic problem.

If assumption A3 holds we can compute the likelihood in the following three steps:

1. Estimate the parameters governing brand choice, α_i and β_i , by maximizing $\mathbb{P}(j|x_t, \mathbf{p}_t)$. This boils down to estimating a (static) conditional logit using only the options with size x_t .¹⁴ This estimation is static, can be done at low computational cost and can include many controls, which among other things help with concerns about the endogeneity of prices.
2. Use the estimated parameters to compute ω_{xit} and estimate the transition probability function $F(\omega_{i,t+1} | \omega_{it})$. Since this step is done once and outside the dynamic problem the transition probability can be estimated very flexibly (and Assumption A2 can be tested by testing if elements of \mathbf{p}_t have power in predicting $\omega_{i,t+1}$ above and beyond ω_t .)
3. Estimate the dynamic parameters – governing the utility from consumption, storage cost and the distribution of ν_t – using $\mathbb{P}(x|\omega_t, i_t)$, which require solving the modified dynamic program.

The split of the likelihood significantly reduces the computational cost and as a result a much richer model can be estimated, allowing for additional variables and rich patterns of observed heterogeneity. Among other things the control for additional variables somewhat reduces the concerns of price endogeneity. The results in Hendel and Nevo (2006b) suggest that this additional richness is important.

A final point worth emphasizing is that the split of the likelihood is separate from the simplification of the state space. The simplification of the state space relied on assumptions A1 and A2. The split in the utility also required Assumption A3.

2.4 Durable Products

Another area that has seen a lot of recent work on dynamics is the estimation of demand for durable products. There is a long tradition in IO of estimating static demand for durable products. Indeed some of the "classic" IO papers involve estimation of demand for durable goods (for example, see Bresnahan 1981, Berry Levinsohn and Pakes, 1995 among many others). In durable goods markets dynamics arise quite naturally since products are used in multiple periods. The durability of the product does not in itself imply that a static model cannot properly capture demand. For example, if consumers hold only a single variety, think a single car, and there are no transaction costs in resale (i.e., products can be sold and purchased costlessly) and no uncertainty about future resale prices, then a purchase of a durable can be seen a static period by period "rental." However, if these conditions do not hold then the current products owned impact purchases. Furthermore, if

¹⁴The idea is similar to the computation of fixed effects in a logit model estimated with panel data: the fixed effects can be partialled out with the right conditioning. Here the conditioning eliminates the dynamics.

consumers are forward looking expectations about future prices as well as quality of the available products impact current decisions.

There are several pricing patterns that can drive dynamics for durable goods. First, just like storable products, there could be temporary price changes that arise, for example in the case of cars, if gas prices temporarily increase or there are temporary discounts.¹⁵ However, a much more common pattern, observed across a wide range of industries are declining prices (and increased quality). This means that the trade off consumers face is between delaying purchase, and the utility obtained from it, with a lower price or higher quality in the future. This is the pattern we focus on.

2.4.1 Implications for Static Estimation

The implications for demand estimation of ignoring dynamics, if they are present in the data, depend on the exact details of the data generating process. For example, a temporary price cut, like the case of storable goods, causes static estimation to overestimate the own price elasticity (and under estimate the cross price elasticity). On the other hand, if gas prices temporarily spike we usually think that static estimates underestimate the impacts of a permanent price increase.

If the key dynamic are declining prices then, in general, it is harder to sign the direction of the bias in static estimation. It is useful to separate between two cases: with and without repeat purchase.

Without repeat purchase – once consumers purchase they leave the market forever – the main bias in static demand estimation is the failure to recognize that each period the potential market size is changing. The static demand model does not recognize that each period the demand curve is potentially changing because some (high willingness-to-pay) consumers have been skimmed off. Consider the following simple example to illustrate the point. Suppose consumers have a willingness to pay that is distributed uniformly on the unit interval, and a total mass of 100. Consumers are myopic and buy the product if the price is below their willingness to pay. Once consumers buy the product they are out of the market forever. This yields a well defined linear demand curve $Q = 100 - 100P$. Suppose we observe a sequence of prices equal to $(0.9, 0.8, 0.7, \dots, 0.1)$. Given the above demand structure the quantity sold over that same time horizon equals 10 units per period. A static demand model lead the researcher to conclude that consumers are not sensitive to price, since the same quantity is sold as prices decline, and estimate an own price elasticity of 0. So in

¹⁵Busse, Simester and Zettelmeyer (2010) study the 2005 Employee Discount Pricing, and show that its main effect was to induce consumers to purchase earlier.

this example the static model underestimates the price sensitivity. More generally, however, even in this example as we change the distribution of willingness to pay and the sequence of observed prices the conclusions might change. Of course, signing the effect is harder once we consider more general models with forward looking consumers.

There are two problems with the standard static random coefficients discrete choice model if there are no repeat purchases. First, the distribution of the random coefficients is likely to change over time as some consumer purchase and exit the market. For example, if prices fall over time it is likely that less price sensitive consumers purchase initially. Second, if consumers are forward looking then they realize there is an option value to not purchasing today. This option value is reflected in the value of the outside option, which in the static model is assumed constant.

With repeat purchases the issues are a bit different. First, the distribution of the consumers does not change, since consumers do not exit. However, consumers who previously purchased a product have a different value of no purchase since their alternative is to stay with their current product. Therefore, the problem with static estimation is that it does not account for the different value, across consumers and over time, of the outside option. Second, now when purchasing consumers do not forgo the option to purchase in the future. Indeed, consumers might find it optimal to buy an inferior option only to replace it shortly after.

2.4.2 A Model of Demand for Durable Goods

We now present a basic model of demand for durable differentiated products. Our presentation follows closely Gowrisankaran and Rysman (2009).¹⁶ The framework extends the static discrete choice model we presented in Section 2.1.1, and in ways is similar to the inventory model we presented in the previous section. Indeed, to some extent the role of inventory is equivalent to the role of the quality of the product already owned. So in the durable good model "stockpiling" means buying a higher quality product. The difference is in the trade-off faced by consumers. The typical price pattern for durable goods is a decreasing quality adjusted price. Faced with this price pattern for storable goods consumers would not stockpile instead they would buy a small amount for current consumption and buy in the future, when the price is lower, for future consumption. In durable goods markets consumers can buy a "small" amount only if they can rent, lease, or resell the used product with low transaction costs. If these options are not available, the consumer's trade-off is between waiting for a lower price, or higher quality product, and either forgoing consumption until then or purchasing a product now and retiring it earlier than needed.

¹⁶See also Melnikov (2001) and Conlon (2009).

The (conditional indirect) utility consumer i gets from product j at time t is given by:

$$u_{ijt} = \gamma_{ijt}^f - \alpha_i p_{jt} + \varepsilon_{ijt} \quad (13)$$

where $\gamma_{ijt}^f = a_{jt}\beta_i + \xi_{jt}$ defines the flow utility. The notation follows the definitions of the static model in Section 2.1.1. If the consumer does not purchase she gets the utility $u_{i0t} = \gamma_{i0t}^f + \varepsilon_{i0t}$ where

$$\gamma_{i0t}^f = \begin{cases} 0 & \text{if no previous purchase} \\ \gamma_{i\hat{j}t}^f & \text{if last purchase was product } \hat{j} \text{ at time } \hat{t} \end{cases} \quad (14)$$

This definition of the utility from the outside option is the main difference between the static model and the dynamic model. Once consumers purchase it changes their outside option. Thus, previous purchases impact current decisions a fact that forward looking consumers realize when they make current choices. Note, that implicitly in the definition of the no purchase option there is an assumption of repeated purchase: consumers are still on the market even after purchase, just with a different outside option.

Assuming that (i) the consumer holds at most a single product at any time and (ii) there is no resale market, then the Bellman equation of the consumer problem is given by

$$V_i(\varepsilon_{it}, \gamma_{i0t}^f, \mathbf{p}_t) = \max_{j=0, \dots, J} \left\{ u_{ijt} + \delta \mathbb{E}[EV_i(\gamma_{ijt}^f, \mathbf{p}_{t+1} | \mathbf{p}_t)] \right\} \quad (15)$$

where $EV_i(\gamma_{ijt}^f, \mathbf{p}_t) = \int V_i(\varepsilon_{it}, \gamma_{ijt}^f, \mathbf{p}_t) dF_\varepsilon(\varepsilon_{it})$, and \mathbf{p}_t represents the set of prices and other product characteristics at period t . The expectation is taken with respect to the uncertainty regarding future vector ε_t , future products, prices and attributes.

If there are no repeat purchases and no resale, then the consumer's problem is slightly different.¹⁷ Because there is no resale, without loss of generality, the utility, γ_{ijt}^f , can be seen as capturing the lifetime value from the product and there is no continuation value. Also, there is no need to keep track of the consumer's stock and $\gamma_{i0t}^f = 0$. The dynamics arise because of the option value of not purchasing. The value function in the no repeat purchase case is given by

$$V_i(\varepsilon_{it}, \mathbf{p}_t) = \max \left\{ \varepsilon_{i0t} + \delta \mathbb{E}[EV_i(\mathbf{p}_{t+1} | \mathbf{p}_t)], \max_{j=1, \dots, J} u_{ijt} \right\} \quad (16)$$

where now $EV_i(\mathbf{p}_t) = \int V_i(\varepsilon_{it}, \mathbf{p}_t) dF_\varepsilon(\varepsilon_{it})$. The first term within brackets represents the value of waiting to purchase in the future. The second term is the value of purchasing today. Because we do not have to keep track of the current holding the state space is reduced.

¹⁷See Melnikov (2001) and Conlon (2009) for applied examples and further discussion of the no-repeat purchase model.

2.4.3 Reducing the Dimension of the State Space

The main computational cost is computing the expected value function, which in the repeat purchase model equals $EV_i(\gamma_{ijt}^f, \mathbf{p}_t)$. The state space is very similar to what we saw in the storable goods problem and consists of the quality of the currently held product, which is equivalent to the inventory in the storable good problem, and the matrix of prices and current attributes required in order to form expectations regarding the future.

Holdings It is useful to briefly consider a somewhat more general model of the consumers holding. In this model the consumer can hold several varieties of the products and the utility from the different varieties interact with each other.¹⁸ There are several ways to model the flow utility in this case,¹⁹ but in all of them the state variable includes a vector describing the consumer's current holdings, and not a scalar. By assuming that the consumer only holds a single option at any point in time, we have reduced the state space to a scalar value of the holding, in the repeat purchase model, or avoided it all together in the no repeat purchase model. Thus, the assumption of holding only a single product serves the same purpose as Assumption A1 in the storable goods model, and reduces the dimension of the holdings, or inventory, variable.

To further understand the differences between the storable and durable goods model consider that a product can be characterized by two dimensions, which can be consumer specific: its quality (utility per use) and its quantity (or how many times it can be used). In the storable goods model we simplified the model by making assumptions on the quality (see the discussion following Assumption A1) and focused on the quantity. Here we leave the quality unrestricted but make assumptions on the quantity by assuming a single good and no depreciation. Allowing for depreciation that is a function of endogenously chosen usage makes the durable goods model closer to the storable goods model.

Note, that there is another similarity with the storable goods model. Here the utility carried forward is γ_{ijt}^f and not $\gamma_{ijt}^f + \varepsilon_{ijt}$. Thus, just like in the inventory model there is a separation between the utility at the time of purchase and the utility at the time of usage.

Finally, we note that an alternative way to reduce the state space is to allow for multiple purchases but assume no interaction in the utility and continue to assume no resale.

¹⁸We should note that even in static models the issue of multiple purchases and the interaction in utility, or through a budget constraint, is mostly an open question and usually ignored. The few exceptions are Hendel (1999), Dube (2004), Nevo, McCabe and Rubinfeld (2005) and Genztkow (2007).

¹⁹For example, the utility can be a function of the products held, or it can be a function of the characteristics of the products held.

Prices Even after reducing the dimension of the holding vector, for a realistic number of products the state space is still too large to be manageable. Like before we rely on the inclusive value to reduce the state space. Define the *dynamic inclusive value* as from the J inside alternatives as:

$$\Omega_{it}(\mathbf{p}_t) = \ln \left(\sum_{j=1}^J \exp(\gamma_{ijt}^f - \alpha_i p_{jt} + \delta \mathbb{E}[EV_i(\gamma_{ijt}^f, \mathbf{p}_{t+1}) | \mathbf{p}_t]) \right). \quad (17)$$

Note, that this definition is different in an important way from the definition given in Section 2.1.1. The above definition provides the expected *value*, including the future value, from the J options. The definition in Section 2.1.1 provides the expected *flow utility*, not accounting for the future value. The difference is not just semantic. The static definition basically provides a statistic that is a summary of (exogenous) prices and attributes of available products. The dynamic definition also includes (endogenous) future behavior of the agent. Once we impose a particular stochastic structure on the evolution of Ω_{it} a natural question is whether the imposed structure is consistent with the consumer optimization problem.

In order reduce the state space we make the modified version of Assumption A2 (Assumption A1, Inclusive Value Sufficiency, in Gowrisankaran and Rysman, 2009):

Assumption A2': $F(\Omega_{i,t+1} | \mathbf{p}_t) = F(\Omega_{i,t+1} | \Omega_{it}(\mathbf{p}_t))$

As before, the assumption assumes that the inclusive value is sufficient to compute the transition probabilities, but now it is the dynamic inclusive value, Ω_{it} . Furthermore, now there is a single inclusive value rather than a vector of size specific inclusive values

Using this assumption we can now write, in the repeat purchase model,

$$EV_i(\gamma_{i0t}^f, \mathbf{p}_t) = EV_i(\gamma_{i0t}^f, \Omega_{it}) = \ln \left(\exp(\Omega_{it}) + \exp \left(\gamma_{i0t}^f + \delta \mathbb{E}[EV_i(\gamma_{i0t}^f, \Omega_{it+1} | \Omega_{it})] \right) \right). \quad (18)$$

Like the storable goods problem Assumption A2' allows us to reduce the state space, but unlike the storable good problem we do not modify the dynamic problem. In the storable goods problem, Assumption A1 allowed us to modify the dynamic problem into a choice between sizes, rather than a choices among brand size combinations. The reason we could do that is that under Assumption A1 choices of the same size impacted the dynamics in the same way. Here we cannot modify the problem because we cannot generate such equivalence classes for the dynamics.

Things are a bit different in the no repeat purchase model. First, the state space can be reduced but the relevant definition of the inclusive value is the static one, given in Section 2.1.1, and not the one given in equation (17). Assuming a version of A2' for the inclusive values we can show that $EV_i(\mathbf{p}_t) = EV_i(\Omega_{it})$. Second, in no repeat purchase the dynamics involves a decision on when to

buy, but conditional on purchase the decision of which product to buy is static. Like the storable product model the dynamic problem can be simplified.

2.5 Estimation and Identification

In this section we discuss the identification and estimation of the model. There are several studies that have estimated demand for durable products using household level data.²⁰ However, recently many studies of demand for durable goods have relied on aggregate data. For this reason we focus our discussion on estimation with aggregate data.

2.5.1 Identification

If consumer level data are observed then, in principle, identification follows the standard arguments (Rust, 1994 and Magnac and Thesmar, 2002, Aguirregabiria, 2010).²¹ With aggregate data we do not observe the purchase history of each consumer, which makes identification significantly more difficult. To see this consider the example in Section 2.4.1. The sequence of quantities can be perfectly explained using a static model, with zero price sensitivity, or with the no repeat purchase dynamic model, which generated the example. Suppose that there are multiple products available in each period, then the model has to explain not just the time series variation in shares but also the cross sectional variation. The key to identifying the model and to separating the different alternative models is the ability of the models to explain both the cross sectional variation, across markets and products, and the time series variation.

We are unaware of a formal identification proof, and getting might be quite difficult. Standard identification proofs for static models require some form of substitution (e.g., what Berry and Haile, 2009, call connected substitutes) between products. In static models the substitution is between products in a given period, but here the requirement is for substitution over time and across products. This need not be satisfied. For example, if the price of a high quality product falls at time t it could actually increase the demand for a low quality product at $t - 1$, because some consumers might buy it for one period.

As we previously discussed, a key issue in static demand estimation has been the potential endogeneity of price. In dynamic demand models estimated using aggregate data the solution has followed closely the static literature, using GMM and moment conditions quite similar to the static

²⁰Many of these studies estimate static demand. For examples of dynamic demand see Erdem, Keane, Oncu and Strebel (2005), or Prince (2008).

²¹The standard arguments need to be adjusted for the existence of ξ_{jt} , but with enough observations these could be controlled for and then we are back in the standard case.

models.

2.5.2 Estimation

The estimation follows closely the method proposed by Berry, Levinsohn and Pakes (1995), but nests a solution of the dynamic programming problem inside the inner loop.²² The basic steps are as follows (see Gowrisankaran and Rysman (2009) for details.)

1. For a given value of the parameters and a vector of mean flow utility levels ($\gamma_{jt}^f = a_{jt}\beta + \xi_{jt}$) compute the predicted market shares by following
 - (a) For a number of simulated consumers, each with a α_i and β_i , calculate the dynamic inclusive value given by equation (17) (start the process with an initial guess on EV_i);
 - (b) Use these inclusive value to compute $F(\Omega_{it+1} | \Omega_{it}(\mathbf{p}_t))$;²³
 - (c) Use the estimated process to update EV_i ;
 - (d) Iterate the previous 3 steps until convergence;
 - (e) Use the estimated policy to simulate for each consumer the purchase path, assuming all consumers initially hold the outside good;
 - (f) Aggregate the consumers purchase decision to obtain market shares;
2. For a given value of the parameters, use the iteration proposed by Berry, Levinsohn and Pakes (1995) to compute the vector of γ_{jt}^f values. The iteration uses the markets shares computed in step 1.
3. As in Berry, Levinsohn and Pakes (1995) use the vector of γ_{jt}^f to compute moment conditions and search for the parameters than minimize a GMM objective function.

In the no repeat purchase model the computation can be simplified by if we add an assumption like Assumption A3, which limits the heterogeneity. See Melnikov (2001).

²²For an alternative computational method see Judd and Su (2009), Dube, Fox and Su (2010) and Conlon (2010).

²³In principle, the process here can be quite general. In reality, however, since the computation is nested with the computation the process has to be fairly simple. Gowrisankaran and Rysman (2009) assume that $\omega_{i,t+1} = \gamma_{1i} + \gamma_{2i}\omega_{i,t} + v_{it}$.

3 Dynamic Games of Oligopoly Competition

The study of firm behavior, especially in oligopoly, is at the heart of Industrial Organization. In many industries, a firm's current actions affect its future future profits, as well as the current and future profits of other firms in the industry. Supply-side dynamics can arise from different sources, including sunk costs of entry, partially irreversible investments, product repositioning costs, or learning-by-doing. Ignoring supply-side dynamics can potentially lead to biases in our estimates of structural parameters. More substantially, accounting for dynamics can change our view of the impact of competition in some industries, as well as our evaluation of public policies. The following examples illustrate these points.²⁴

Example 1. Product repositioning in differentiated product markets. Sweeting (2007) and Aguirregabiria and Ho (2009) are two examples of empirical applications that endogenize product attributes using a dynamic game of competition in a differentiated products industry. Sweeting estimates a dynamic game of oligopoly competition in the US commercial radio industry. The model endogenizes the choice of radio stations format (genre), and estimates product repositioning costs. Aguirregabiria and Ho (2009) study the contribution of different factors to explain airlines' adoption of *hub-and-spoke networks*. They propose and estimate a dynamic game of airline network competition where the number of direct connections that an airline has in an airport is an endogenous product characteristic. These studies highlight the two potential limitations of static models. First, a common assumption in many static (and dynamic) demand models is that product characteristics, other than prices, are exogenous. This assumption, if violated, can generate biases in the estimated parameters. The dynamic game acknowledges the endogeneity of some product characteristics and exploits the dynamic structure of the model to generate valid moment conditions for the consistent estimation of the structural parameters. A second important limitation of a static model of firm behavior is that it cannot recover the costs of repositioning product characteristics. As a result, the static model cannot address important empirical questions such as the effect of a merger on product repositioning.

Example 2. Evaluating the effects of regulation. Ryan (2006) provides another example of how ignoring the endogeneity of market structure, and its dynamics, can lead to misleading results. He studies the effects of the 1990 Amendments to the Clean Air Act on the US cement industry.

²⁴As before, we cannot provide a complete survey of the literature. Other examples of papers that study similar questions are Einav (2009), Collard-Wexler (2006), Macieira (2007), Krykov (2008), Hashmi and Van Biesebroeck (2010), Snider (2009), Suzuki (2010), Gowrisankaran et al.(2010), Walrath (2010), and Finger (2008), among others.

This environmental regulation added new categories of regulated emissions, and introduced the requirement of an environmental certification that cement plants have to pass before starting their operation. Ryan estimates a dynamic game of competition where the sources of dynamics are sunk entry costs and adjustment costs associated with changes in installed capacity. The estimated model shows that the new regulation had negligible effects on variable production costs but it increased significantly the sunk cost of opening a new cement plant. A static analysis, that ignores the effects of the policy on firms' entry-exit decisions, would conclude that the regulation had negligible effects on firms profits and consumer welfare. In contrast, the dynamic analysis shows that the increase in sunk-entry costs caused a reduction in the number of plants that in turn implied higher markups and a decline in consumer welfare.

Initial attempts to answer many of these questions were done using entry models in the spirit of Bresnahan and Reiss (1990, 1991) and Berry (1992). The simplest forms of these models use a reduced form profit function, in the sense that variable profits are not derived from explicit models of price or quantity competition, and static, in the sense that firms are not forward looking. These models have been used to explain cross-market variation in market structure that is assumed to be an equilibrium of an entry game. It is possible, to include predetermined variables in the payoff function (e.g., firm size, capacity, incumbent status), and to interpret the payoff function as an intertemporal value function (see Bresnahan and Reiss, 1993). Indeed, one could use panel data to estimate some of the parameters, or use price and quantity data to estimate the variable profits. These models typically are much easier to estimate than the dynamic games we discuss below, and therefore at times they might serve as a useful first cut of the data. The main limitation of this approach is that often the parameters do not have a clear economic interpretation in terms of costs or demand, and the model cannot be used for counterfactual policy experiments. Furthermore, empirical questions in IO that have to do with the effects of uncertainty on firm behavior and competition, or that try to distinguish between short-run and long-run effects of exogenous shocks, typically require the specification and estimation of dynamic structural models that explicitly take into account firms' forward-looking behavior. For these reasons, most of the recent work in IO dealing with industry dynamics has relied on more explicit modeling of dynamics, as in the model of Ericson and Pakes (1995). In Section 3.2 we briefly describe a simple version of this model that allows us to demonstrate our key points.

Sections 3.3 to 3.5 discuss some of the main econometric, computational, and modeling issues faced by applied researchers who want to estimate a dynamic game. The standard nested fixed

point algorithm, that has been used successfully in the estimation of single-agent models, is computationally unfeasible in actual applications of dynamic games. As a result, researchers have turned to alternative methods based on the ideas of Hotz and Miller (1993) and Aguirregabiria and Mira (2002), i.e., estimation methods based on conditional choice probabilities (CCP). We survey some of the methods that have been proposed to implement these ideas. We focus on several issues. First, we discuss the impact of multiple equilibria on identification, and present sufficient conditions for point identification of the structural parameters. We then turn to discuss the properties of an iterative procedure that has been proposed by Aguirregabiria and Mira (2007) to deal with one of the potential shortcomings of two-step CCP methods: finite sample bias. A recent paper by Pesendorfer and Schmidt-Dengler (2008) shows that indeed in some cases finite sample bias is reduced, but in other cases the iterative procedure actually increases the bias. We provide stability conditions on the equilibrium that guarantee the performance of the method and explain the results of Pesendorfer and Schmidt-Dengler.

Another main shortcoming of the CCP approach is the lack of unobserved firm or market level heterogeneity, beyond a firm level i.i.d. shock. In Section 3.4 we briefly discuss some new CCP methods that allow us to relax this assumption. In Section 3.5 we return to a theme that was a major part of our discussion of dynamic demand: methods to reduce the dimension of the state space. We show how the inclusive-value approach discussed above can be extended to dynamic games in order to reduce the computational burden in the solution and estimation of this class of models. We conclude in Section 3.6 with a description of an homotopy method that can be used to implement counterfactual experiments given the estimated model.

3.1 The Structure of Dynamic Games of Oligopoly Competition

We use a simple dynamic game of market entry-exit to illustrate the different issues and methods. Time is discrete and indexed by t . The game is played by N firms that we index by i . Let a_{it} be the decision variable of firm i at period t . In the entry-exit model we consider, the decision variable is a binary indicator of the event "firm i is active in the market at period t ". The action is taken to maximize the expected and discounted flow of profits in the market, $E_t(\sum_{r=0}^{\infty} \delta^r \Pi_{it+r})$ where $\delta \in (0, 1)$ is the discount factor, and Π_{it} is firm i 's profit at period t .

The profits of firm i at time t are given by $\Pi_{it} = VP_{it} - FC_{it} - EC_{it}$, where VP_{it} represents variable profits, FC_{it} is the fixed cost of operating, and EC_{it} is a one time entry cost. Following the standard structure in the Ericson-Pakes (1995) framework, incumbent firms in the market at period t compete in prices or quantities in a static Cournot or Bertand model. For example, the

variable profit function can take on the form²⁵

$$VP_{it}(a_{it}, a_{-it}) = a_{it} H_t \sum_{n=0}^{N-1} 1 \left\{ \sum_{j \neq i} a_{jt} = n \right\} \theta_{i,n}^{VP} \quad (19)$$

H_t is a measure of market size; $1\{\cdot\}$ is the indicator function; and $\sum_{j \neq i} a_{jt}$ is the number of active competitors of firm i at period t . The vector of parameters $\{\theta_{i,n}^{VP} : n = 0, 1, \dots, N-1\}$ represents firm i 's variable profit per-capita when there are other n competitors active in the market. We expect $\theta_{i,0}^{VP} \geq \theta_{i,1}^{VP} \geq \dots \geq \theta_{i,N-1}^{VP}$. The fixed cost is paid every period that the firm is active in the market, and it has the following structure, $FC_{it} = a_{it} (\theta_i^{FC} + \varepsilon_{it})$. θ_i^{FC} is a parameter that represents the mean value of the fixed operating cost of firm i . ε_{it} is a zero-mean shock that is private information of firm i . The entry cost is paid only if the firm was not active in the market at previous period: $EC_{it} = a_{it}(1 - s_{it}) \theta_i^{EC}$, where s_{it} is a binary indicator that is equal to 1 if firm i was active in the market in period $t-1$, i.e., $s_{it} \equiv a_{i,t-1}$, and θ_i^{EC} is a parameter that represents the entry cost of firm i . The specification of the primitives of the model is completed with the transition rules of the state variables. Market size follows an exogenous Markov process with transition probability function $F_H(H_{t+1}|H_t)$. The transition of the incumbent status is trivial, $s_{it+1} = a_{it}$. Finally, the private information shock ε_{it} is i.i.d. over time and independent across firms with CDF G_i .²⁶

Somewhat in contrast to static entry models, where both games of complete and incomplete information have been studied, the recent literature on empirical dynamic games has focused solely on games of incomplete information. The introduction of private information shocks ensures the existence of an equilibrium in pure strategies (Doraszelski and Satterthwaite, 2010). In addition, these random shocks are a convenient way to allow for econometric unobservables that can explain how agents with the same observable characteristics make different decisions.

Following Ericson and Pakes (1995), most of the recent literature in IO studying industry dynamics focuses on studying a Markov Perfect Equilibrium (MPE), as defined by Maskin and Tirole (1987, 1988a, 1988b). The key assumption in this solution concept is that players' strategies are functions of only payoff-relevant state variables. We use the vector \mathbf{x}_t to represent all the

²⁵This indirect variable profit function may come from the equilibrium of a static Bertrand game with differentiated product as in the framework presented in section 2.1. Suppose that all firms have the same marginal cost and product differentiation is symmetric. For instance, consumer utility of buying product i is $u_{it} = \nu - \alpha p_{it} + \varepsilon_{it}$, where ν and α are parameters, and ε_{it} is a consumer-specific i.i.d. random variable. Then, the equilibrium variable profit of an active firm depends only on the number of firms active in the market.

²⁶In this example, we consider that firms' entry-exit decisions are made at the beginning of period t and they are effective during the same period. An alternative timing that has been considered in some applications is that there is a one-period time-to-build, i.e., the decision is made at period t , and entry costs are paid at period t , but the firm is not active in the market until period $t+1$. The latter is in fact the timing of decisions in Ericson and Pakes (1995).

common knowledge state variables at period t , i.e., $\mathbf{x}_t \equiv (H_t, s_{1t}, s_{2t}, \dots, s_{Nt})$. In this model, the payoff-relevant state variables for firm i are $(\mathbf{x}_t, \varepsilon_{it})$.²⁷ Let $\alpha = \{\alpha_i(\mathbf{x}_t, \varepsilon_{it}) : i \in \{1, 2, \dots, N\}\}$ be a set of strategy functions, one for each firm. A MPE is a set of strategy functions α^* such that every firm is maximizing its value given the strategies of the other players. For given strategies of the other firms, the decision problem of a firm is a single-agent dynamic programming (DP) problem. Let $V_i^\alpha(\mathbf{x}_t, \varepsilon_{it})$ be the value function of this DP problem. This value function is the unique solution to the Bellman equation:

$$V_i^\alpha(\mathbf{x}_t, \varepsilon_{it}) = \max_{a_{it} \in \{0,1\}} \left\{ a_{it} (\Pi_i^\alpha(\mathbf{x}_t) - \varepsilon_{it}) + \delta \int V_i^\alpha(\mathbf{x}_{t+1}, \varepsilon_{it+1}) dG_i(\varepsilon_{it+1}) F_i^\alpha(\mathbf{x}_{t+1}|a_{it}, \mathbf{x}_t) \right\} \quad (20)$$

where $a_{it} (\Pi_i^\alpha(\mathbf{x}_t) - \varepsilon_{it})$ and $F_i^\alpha(\mathbf{x}_{t+1}|a_{it}, \mathbf{x}_t)$ are the expected one-period profit and the expected transition of the state variables, respectively, for firm i given the strategies of the other firms. By definition, the expected one-period profit $\Pi_i^\alpha(\mathbf{x}_t)$ is,

$$\Pi_i^\alpha(\mathbf{x}_t) = H_t \sum_{n=0}^{N-1} \Pr \left(\sum_{j \neq i} \alpha_j(\mathbf{x}_t, \varepsilon_{jt}) = n \mid \mathbf{x}_t \right) \theta_{i,n}^{VP} - \theta_i^{FC} - (1 - s_{it})\theta_i^{EC} \quad (21)$$

And the expected transition of the state variables is:

$$F_i^\alpha(\mathbf{x}_{t+1}|a_{it}, \mathbf{x}_t) = 1\{s_{it+1} = a_{it}\} \left[\prod_{j \neq i} \Pr(s_{j,t+1} = \alpha_j(\mathbf{x}_t, \varepsilon_{jt}) \mid \mathbf{x}_t) \right] F_H(H_{t+1} \mid H_t) \quad (22)$$

A player's best response function gives his optimal strategy if the other players behave, now and in the future, according to their respective strategies. In this model, the best response function of player i is $1\{\varepsilon_{it} \leq v_i^\alpha(\mathbf{x}_t)\}$, where $v_i^\alpha(\mathbf{x}_t)$ is the threshold value of ε_{it} that leaves firm i indifferent between alternative 0 and 1 given that the other players play strategies is α . According to our model,

$$v_i^\alpha(\mathbf{x}_t) \equiv \Pi_i^\alpha(\mathbf{x}_t) + \delta \int V_i^\alpha(\mathbf{x}_{t+1}, \varepsilon_{it+1}) dG_i(\varepsilon_{it+1}) [F_i^\alpha(\mathbf{x}_{t+1}|1, \mathbf{x}_t) - F_i^\alpha(\mathbf{x}_{t+1}|0, \mathbf{x}_t)] \quad (23)$$

A Markov perfect equilibrium (MPE) in this game is a set of strategy functions α^* such that for any player i and for any $(\mathbf{x}_t, \varepsilon_{it})$ we have that $\alpha_i^*(\mathbf{x}_t, \varepsilon_{it}) = 1\{\varepsilon_{it} \leq v_i^{\alpha^*}(\mathbf{x}_t)\}$.

Given a strategy function $\alpha_i(\mathbf{x}_t, \varepsilon_{it})$, we define the corresponding *Conditional Choice Probability (CCP)* function as:

$$P_i(\mathbf{x}_t) \equiv \Pr(\alpha_i(\mathbf{x}_t, \varepsilon_{it}) = 1 \mid \mathbf{x}_t) = \int \alpha_i(\mathbf{x}_t, \varepsilon_{it}) dG_i(\varepsilon_{it}) \quad (24)$$

²⁷If private information shocks are serially correlated, the history of previous decisions contains useful information to predict the value of a player's private information, and it should be part of the set of payoff relevant state variables. Therefore, the assumption that private information is independently distributed over time has implications for the set of payoff-relevant state variables.

Since choice probabilities are integrated over the continuous variables in ε_{it} , they are lower dimensional objects than the strategies α . For instance, when both a_{it} and \mathbf{x}_t are discrete, CCPs can be described as vectors in a finite dimensional Euclidean space. In our entry-exit model, $P_i(\mathbf{x}_t)$ is the probability that firm i is active in the market given the state \mathbf{x}_t . By definition, given $\alpha_i(\mathbf{x}_t, \varepsilon_{it})$ the CCP $P_i(\mathbf{x}_t)$ is uniquely determined. If the private information shock ε_{it} : (a) is i.i.d. over time; (b) does not enter in the transition probability of \mathbf{x}_t (i.e., conditional independence assumption); and (c) enters additively in the expected one-period profit (i.e., additive separability), then given a CCP function $P_i(\mathbf{x}_t)$ there is a unique strategy function $\alpha_i(\mathbf{x}_t, \varepsilon_{it})$ compatible with it.²⁸ Therefore, there is a one-to-one relationship between strategy functions and CCPs. From now on, we use CCPs to represent players' strategies, and use the terms 'strategy' and 'CCP' as interchangeable. We also use $\Pi_i^{\mathbf{P}}$ and $F_i^{\mathbf{P}}$ instead of Π_i^α and F_i^α to represent the expected profit function and the transition probability function, respectively.

Based on the concept of CCP, we describe a representation of the equilibrium mapping and of a MPE that is particularly useful for the econometric analysis.²⁹ This representation has two main features: (1) a MPE is a vector of CCPs; and (2) a player's best response is an optimal response not only to the other players' strategies but also to his own strategy in the future. A MPE is a vector of CCPs, $\mathbf{P} \equiv \{P_i(\mathbf{x}_t) : i = 1, 2, \dots, N; \mathbf{x}_t \in \mathcal{X}\}$, such that for every firm and any state \mathbf{x}_t the following equilibrium condition is satisfied:

$$P_i(\mathbf{x}_t) = G_i \left(\Pi_i^{\mathbf{P}}(\mathbf{x}_t) + \delta \sum_{\mathbf{x}_{t+1}} [F_i^{\mathbf{P}}(\mathbf{x}_{t+1}|1, \mathbf{x}_t) - F_i^{\mathbf{P}}(\mathbf{x}_{t+1}|0, \mathbf{x}_t)] V_i^{\mathbf{P}}(\mathbf{x}_{t+1}) \right) \quad (25)$$

The right hand side of equation (25) is a best response probability function. $V_i^{\mathbf{P}}$ is the valuation operator of player i if every player behaves now and in the future according to their respective strategies in \mathbf{P} . We can obtain $V_i^{\mathbf{P}}$ as the unique solution of the recursive expression:

$$V_i^{\mathbf{P}}(\mathbf{x}_t) = P_i(\mathbf{x}_t) [\Pi_i^{\mathbf{P}}(\mathbf{x}_t) + e_i(P_i(\mathbf{x}_t))] + \delta \sum_{\mathbf{x}_{t+1}} V_i^{\mathbf{P}}(\mathbf{x}_{t+1}) F^{\mathbf{P}}(\mathbf{x}_{t+1}|\mathbf{x}_t) \quad (26)$$

where $e_i(P_i(\mathbf{x}_t))$ is the expectation $E_{\varepsilon_{it}}(\varepsilon_{it} | \varepsilon_{it} \leq G_i^{-1}(P_i(\mathbf{x}_t)))$, and the form of the function $e_i(\cdot)$ depends on the probability distribution of ε_{it} .³⁰ When the space \mathcal{X} is discrete and finite, we

²⁸Under conditions (a), (b), and (c), the best response function has the single-threshold form $1\{\varepsilon_{it} \leq v_i^\alpha(\mathbf{x}_t)\}$. Therefore, we can limit our analysis to the set of strategy functions with this threshold structure. This implies that CCP functions should have the form $P_i(\mathbf{x}_t) = G_i(v_i^\alpha(\mathbf{x}_t))$. Since the CDF function G_i is invertible everywhere, for given $P_i(\mathbf{x}_t)$ there is a unique threshold $v_i^\alpha(\mathbf{x}_t)$ that is compatible with this choice probability, i.e., $v_i^\alpha(\mathbf{x}_t) = G_i^{-1}(P_i(\mathbf{x}_t))$ where $G_i^{-1}(\cdot)$ is the inverse function of G_i . Thus, given a CCP function $P_i(\mathbf{x}_t)$, the unique strategy function compatible with it is $\alpha_i(\mathbf{x}_t, \varepsilon_{it}) = 1\{\varepsilon_{it} \leq G_i^{-1}(P_i(\mathbf{x}_t))\}$. This result can be extended to multinomial discrete choice model with a general number of J choice alternatives and to dynamic games with continuous decision variables.

²⁹For the general results, see the *Representation Lemma* in Aguirregabiria and Mira (2007).

³⁰If ε_{it} is normally distributed with zero mean and variance σ_i^2 , then $e_i(P_i(\mathbf{x}_t)) = -\sigma_i \phi(\Phi^{-1}(P_i(\mathbf{x}_t)))$, where ϕ is the PDF and Φ^{-1} is the inverse CDF of the standard normal. If ε_{it} is extreme value type I with dispersion parameter σ_i , we have that $e_i(P_i(\mathbf{x}_t)) = \gamma - \sigma_i \ln(P_i(\mathbf{x}_t))$, where γ is Euler's constant.

can obtain $V_i^{\mathbf{P}}$ as the solution of a system of linear equations of dimension $|\mathcal{X}|$. In vector form, $\mathbf{V}_i^{\mathbf{P}} = (\mathbf{I} - \delta \mathbf{F}^{\mathbf{P}})^{-1} P_i * [\mathbf{\Pi}_i^{\mathbf{P}} + e_i^{\mathbf{P}}]$, where \mathbf{I} is the identity matrix, $\mathbf{V}_i^{\mathbf{P}}$, $\mathbf{\Pi}_i^{\mathbf{P}}$, and $e_i^{\mathbf{P}}$ are $|\mathcal{X}| \times 1$ vectors, and $\mathbf{F}^{\mathbf{P}}$ is the $|\mathcal{X}| \times |\mathcal{X}|$ transition matrix with elements $F^{\mathbf{P}}(\mathbf{x}_{t+1}|\mathbf{x}_t)$. We represent the equilibrium mapping in matrix form as $\Psi(\mathbf{P}, \boldsymbol{\theta})$, such that a MPE associated with a vector of structural parameters $\boldsymbol{\theta}$ is a fixed point $\mathbf{P} = \Psi(\mathbf{P}, \boldsymbol{\theta})$.

The valuation and the best response operators can be further simplified for the class of models where the expected profit function is multiplicatively separable in the structural parameters. In our entry-exit model, $\Pi_i^{\mathbf{P}}(\mathbf{x}_t) = \mathbf{z}_i^{\mathbf{P}}(\mathbf{x}_t)\boldsymbol{\theta}_i$, where $\boldsymbol{\theta}_i$ is the vector of structural parameters $(\theta_{i,0}^{VP}, \theta_{i,1}^{VP}, \dots, \theta_{i,N-1}^{VP}, \theta_i^{FC}, \theta_i^{EC})'$, and $\mathbf{z}_i^{\mathbf{P}}(\mathbf{x}_t)$ is a vector of known functions $\{H_t \Pr(\sum_{j \neq i} a_{jt} = 0 | \mathbf{x}_t, \mathbf{P}), \dots, H_t \Pr(\sum_{j \neq i} a_{jt} = N-1 | \mathbf{x}_t, \mathbf{P}), -1, -(1-s_{it})\}$. If ε_{it} is normally distributed with zero mean and variance σ_i^2 , we have that $e_i(P_i(\mathbf{x}_t)) = -\sigma_i \phi(\Phi^{-1}(P_i(\mathbf{x}_t)))$. Therefore, the valuation operator is multiplicatively separable in the structural parameters: $\mathbf{V}_i^{\mathbf{P}} = \mathbf{W}_{z,i}^{\mathbf{P}} \boldsymbol{\theta}_i + \mathbf{W}_{e,i}^{\mathbf{P}} \sigma_i$, where $\mathbf{W}_i^{\mathbf{P}}$ is the matrix $(\mathbf{I} - \delta \mathbf{F}^{\mathbf{P}})^{-1} P_i * \mathbf{Z}_i^{\mathbf{P}}$, and $\mathbf{W}_{e,i}^{\mathbf{P}}$ is the vector $(\mathbf{I} - \delta \mathbf{F}^{\mathbf{P}})^{-1} P_i * e_i^{\mathbf{P}}$, with $\mathbf{Z}_i^{\mathbf{P}}$ being a matrix with rows the vectors $\mathbf{z}_i^{\mathbf{P}}(\mathbf{x}_t)$, P_i is the vector of CCPs for player i , $e_i^{\mathbf{P}}$ is the vector with elements $\phi(\Phi^{-1}(P_i(\mathbf{x}_t)))$, and $*$ is the element-by-element product. Then, the best response probability function is:

$$P_i(\mathbf{x}_t) = \Phi \left(\tilde{\mathbf{z}}_i^{\mathbf{P}}(\mathbf{x}_t) \frac{\boldsymbol{\theta}_i}{\sigma_i} + \tilde{e}_i^{\mathbf{P}}(\mathbf{x}_t) \right) \quad (27)$$

where $\tilde{\mathbf{z}}_i^{\mathbf{P}}(\mathbf{x}_t)$ is equal to $\mathbf{z}_i^{\mathbf{P}}(\mathbf{x}_t) + \delta \sum_{\mathbf{x}_{t+1}} [F_i^{\mathbf{P}}(\mathbf{x}_{t+1}|1, \mathbf{x}_t) - F_i^{\mathbf{P}}(\mathbf{x}_{t+1}|0, \mathbf{x}_t)] W_{z,i}^{\mathbf{P}}(\mathbf{x}_{t+1})$, and $\tilde{e}_i^{\mathbf{P}}(\mathbf{x}_t)$ is $\delta \sum_{\mathbf{x}_{t+1}} [F_i^{\mathbf{P}}(\mathbf{x}_{t+1}|1, \mathbf{x}_t) - F_i^{\mathbf{P}}(\mathbf{x}_{t+1}|0, \mathbf{x}_t)] W_{e,i}^{\mathbf{P}}(\mathbf{x}_{t+1})$.

3.2 Data, Identification, and Estimation

3.2.1 Data

In most applications of dynamic games in empirical IO the researcher observes a random sample of M markets, indexed by m , over T periods of time, where the observed variables consists of players' actions and state variables. In the standard application in IO, the values of N and T are small, but M is large. Two aspects of the data deserve some comments. For the moment, we consider that the industry and the data are such that: (a) each firm is observed making decisions in every of the M markets; and (b) the researcher knows all the payoff relevant market characteristics that are common knowledge to the firms. We describe condition (a) as a data set with *global players*. For instance, this is the case in a retail industry characterized by competition between large retail chains which are potential entrants in any of the local markets that constitute the industry. With this type of data we can allow for rich firm heterogeneity that is fixed across markets and time

by estimating firm-specific structural parameters, θ_i . This 'fixed-effect' approach to deal with firm heterogeneity is not feasible in data sets where most of the competitors can be characterized as *local players*, i.e., firms specialized in operating in a few markets. Condition (b) rules out the existence of unobserved market heterogeneity. Though it is a convenient assumption, it is also unrealistic for most applications in empirical IO. In section 3.4 we present estimation methods that relax conditions (a) and (b) and deal with unobserved market and firm heterogeneity.

3.2.2 Identification with multiple equilibria

Multiple equilibria are the rule, rather than the exception, in most dynamic games. We now discuss the implications of multiple equilibria for identification. Equilibrium uniqueness is neither a necessary nor a sufficient condition for the identification of a model (Jovanovic, 1989). To see this, consider a model with vector of structural parameters $\theta \in \Theta$, and define the mapping $C(\theta)$ from the set of parameters Θ to the set of measurable predictions of the model. Multiple equilibria implies that the mapping $C(\cdot)$ is a correspondence. A model is not point-identified if at the observed data the inverse mapping C^{-1} is a correspondence. In general, C being a function (i.e., equilibrium uniqueness) is neither a necessary nor a sufficient condition for C^{-1} being a function (i.e., for point identification).

To illustrate the identification of a game with multiple equilibria, we start with a simple binary choice game with identical players and where the equilibrium probability P is implicitly defined as the solution of the condition $P = \Phi(-1.8 + \theta P)$, where θ is a structural parameter, and $\Phi(\cdot)$ is the CDF of the standard normal. Suppose that the true value θ_0 is 3.5. It is possible to verify that the set of equilibria associated with θ_0 is $C(\theta_0) = \{P^{(A)}(\theta_0) = 0.054, P^{(B)}(\theta_0) = 0.551, \text{ and } P^{(C)}(\theta_0) = 0.924\}$. The game has been played M times and we observe players' actions for each realization of the game $\{a_{im} : i, m\}$. Let P_0 be the population probability $\Pr(a_{im} = 1)$. Without further assumptions the probability P_0 can be estimated consistently from the data. For instance, a simple frequency estimator $\hat{P}_0 = (NM)^{-1} \sum_{i,m} a_{im}$ is a consistent estimator of P_0 . Without further assumption, we do not know the relationship between population probability P_0 and the equilibrium probabilities in $C(\theta_0)$. If all the sample observations come from the same equilibrium, then P_0 should be one of the points in $C(\theta_0)$. However, if the observations come from different equilibria in $C(\theta_0)$, then P_0 is a mixture of the elements in $C(\theta_0)$. To obtain identification, we can assume that every observation in the sample comes from the same equilibrium. Under this condition, since P_0 is an equilibrium associated with θ_0 , we know that $P_0 = \Phi(-1.8 + \theta_0 P_0)$. Given that $\Phi(\cdot)$ is an invertible function, we have that $\theta_0 = (\Phi^{-1}(P_0) + 1.8)/P_0$. Provided that P_0

is not zero, it is clear that θ_0 is point identified regardless the existence of multiple equilibria in the model.³¹

The basic idea in this example can be extended to get identification in our class of dynamic games. We make the following assumptions.

Assumption: Single equilibrium in the data. Every observation in the sample comes from the same equilibrium, i.e., for any observation (m, t) , $\mathbf{P}_{mt}^0 = \mathbf{P}^0$.

Assumption: No unobserved common-knowledge variables. The only unobservables for the econometrician are the private information shocks ε_{imt} and the structural parameters θ .

The distribution of ε_{imt} is known up to a scale parameter. For the sake of concreteness, consider that ε_{imt} is a normal random variable with zero mean and variance σ_i . Under these assumptions the vector of population CCPs \mathbf{P}^0 is an equilibrium of the model associated with θ^0 (i.e., it is not a mixture of equilibria) and it is identified from the data. Since \mathbf{P}^0 is an equilibrium, the condition $P_i^0(\mathbf{x}_{mt}) = \Phi(\bar{\mathbf{z}}_i^{\mathbf{P}^0}(\mathbf{x}_{mt}) \theta_i^0/\sigma_i + \tilde{e}_i^{\mathbf{P}^0}(\mathbf{x}_{mt}))$ is satisfied for any firm i and any state \mathbf{x}_{mt} . We can re-write this equilibrium condition as a linear-in-parameters model $Y_{imt} = X_{imt} \theta_i^0/\sigma_i$, where $Y_{imt} \equiv \Phi^{-1}(P_i^0(\mathbf{x}_{mt})) - \tilde{e}_i^{\mathbf{P}^0}(\mathbf{x}_{mt})$, and $X_{imt} \equiv \bar{\mathbf{z}}_i^{\mathbf{P}^0}(\mathbf{x}_{mt})$. A necessary and sufficient condition for the identification of θ_i^0/σ_i is that the variance-covariance matrix $E(X_{imt}'X_{imt})$ is non-singular, or equivalently, that the variables in X_{imt} are not perfectly collinear. The variables in X_{imt} are expected present values of the variables in the one-period expected profit, $\mathbf{z}_i^{\mathbf{P}^0}(\mathbf{x}_{mt})$. In general, these variables and their expected present values are not collinear, and therefore θ_i^0/σ_i is identified. Under the single-equilibrium-in-the-data assumption, the multiplicity of equilibria in the model does not play any role in the identification of the structural parameters.³²

3.2.3 Estimation: Maximum Likelihood and Two-Step Methods

In principle, estimation of dynamic games could follow the same methods as the estimation of single-agent dynamic structural models. For example, one could imagine using a nested fixed point algorithm that maximizes a sample criterion function over the space of structural parameters, and solves for the equilibrium of the model, assuming it is unique, for each trial value of the parameters.

³¹The *single-equilibrium-in-the-data assumption* plays a key role in this identification result. Suppose that a fraction λ of the observations comes from the stable equilibrium $P^{(A)}(\theta_0) = 0.054$ and a fraction $1 - \lambda$ from the other stable equilibrium $P^{(C)}(\theta_0) = 0.924$, but the researcher does not know which observation comes from one equilibrium or the other. Therefore, $P_0 = \lambda P^{(A)}(\theta_0) + (1 - \lambda)P^{(C)}(\theta_0)$. Note that P_0 is not an equilibrium of the model but a mixture of two equilibria. If the researcher ignores this mixture and imposes the assumption of a *single-equilibrium-in-the-data assumption*, the estimator of θ_0 will be inconsistent.

³²The single-equilibrium-in-the-data assumption is a sufficient for identification but it is not necessary. Sweeting (2009), Aguirregabiria and Mira (2009), and Paula and Tang (2010) present conditions for the point-identification of static games of incomplete information when there are multiple equilibria in the data.

While this approach has been quite successful in single agent problems, it is problematic in games. The existence of multiple equilibria significantly increases the computational burden, especially if we use standard estimation methods such as maximum likelihood or GMM. In this section we discuss how the literature has dealt with this issue.

The use of a 'extended' or 'pseudo' likelihood (or alternatively GMM criterion) function plays an important role in the different estimation methods. For arbitrary values of the vector of structural parameters θ and firms' strategies \mathbf{P} , we define the following likelihood function of observed players' actions $\{a_{imt}\}$ conditional on observed state variables $\{\mathbf{x}_{mt}\}$:

$$Q(\theta, \mathbf{P}) = \sum_{i,m,t} a_{imt} \ln \Phi(\tilde{\mathbf{z}}_{imt}^{\mathbf{P}} \theta_i + \tilde{e}_{imt}^{\mathbf{P}}) + (1 - a_{imt}) \ln \Phi(-\tilde{\mathbf{z}}_{imt}^{\mathbf{P}} \theta_i - \tilde{e}_{imt}^{\mathbf{P}}) \quad (28)$$

where, for the sake of concreteness, we consider that private information shocks are normally distributed, and for notational simplicity we use θ_i to represent θ_i/σ_i , $\tilde{\mathbf{z}}_{imt}^{\mathbf{P}} \equiv \tilde{\mathbf{z}}_i^{\mathbf{P}}(\mathbf{x}_{mt})$, and $\tilde{e}_{imt}^{\mathbf{P}} \equiv \tilde{e}_i^{\mathbf{P}}(\mathbf{x}_{mt})$. We call $Q(\theta, \mathbf{P})$ a Pseudo Likelihood function because players' CCPs in \mathbf{P} are arbitrary and do not represent the equilibrium probabilities associated with θ implied by the model. An implication of using arbitrary CCPs, instead of equilibrium CCPs, is that likelihood Q is a function and not a correspondence³³ To compute this pseudo likelihood, a useful construct is the representation of equilibrium in terms of CCPs, which we presented above.

(a) Full Maximum Likelihood. The dynamic game imposes the restriction that the strategies in \mathbf{P} should be in equilibrium. The ML estimator is defined as the pair $(\hat{\theta}_{MLE}, \hat{\mathbf{P}}_{MLE})$ that maximizes the pseudo likelihood subject to the constraint that the strategies in $\hat{\mathbf{P}}_{MLE}$ are equilibrium strategies associated with $\hat{\theta}_{MLE}$. That is,

$$(\hat{\theta}_{MLE}, \hat{\mathbf{P}}_{MLE}) = \arg \max_{(\theta, \mathbf{P})} Q(\theta, \mathbf{P}) \quad (29)$$

$$\text{s.t. } P_i(\mathbf{x}_{mt}) = \Phi(\tilde{\mathbf{z}}_i^{\mathbf{P}}(\mathbf{x}_{mt})\theta_i + \tilde{e}_i^{\mathbf{P}}(\mathbf{x}_{mt})) \text{ for any } (i, \mathbf{x}_{mt}) \in I \times \mathcal{X}$$

This is a constrained ML estimator that satisfies the standard regularity conditions for consistency, asymptotic normality and efficiency of ML estimation.³⁴ The numerical solution of the constrained optimization problem that defines these estimators requires one to search over an extremely large dimensional space. In the empirical applications of dynamic oligopoly games, the vector of probabilities \mathbf{P} includes thousands or millions of elements. Searching for an optimum in that kind of

³³In fact, Q is the likelihood function of a standard probit model for a_{imt} with explanatory variables $\tilde{\mathbf{z}}_{imt}^{\mathbf{P}}$ and $\tilde{e}_{imt}^{\mathbf{P}}$.

³⁴Similarly, we could define a Pseudo GMM criterion function. The GMM estimator is defined as the pair $(\hat{\theta}_{GMM}, \hat{\mathbf{P}}_{GMM})$ that minimizes the pseudo GMM criterion subject to the constraint that the strategies in $\hat{\mathbf{P}}_{GMM}$ are equilibrium strategies associated with $\hat{\theta}_{GMM}$.

space is computationally demanding. Su and Judd (2008) have proposed to use a MPEC algorithm, which is a general purpose algorithm for the numerical solution of constrained optimization problems. However, even using the most sophisticated algorithm such as MPEC, the optimization with respect to $(\mathbf{P}, \boldsymbol{\theta})$ can be extremely demanding when \mathbf{P} has a high dimension.

(b) Two-step methods. In order to avoid this large computational cost, alternative two step methods have been explored. In this class of models, for given \mathbf{P} , the best response probability function $G_i(\mathbf{z}_i^{\mathbf{P}}(\mathbf{x}_t) \boldsymbol{\theta}_i + \hat{c}_{imt}^{\mathbf{P}})$ has the structure in a standard binary choice model with an index that is linear in parameters. A pseudo likelihood function based on these best response probabilities is globally concave in the structural parameters, and therefore optimization of $Q(\boldsymbol{\theta}, \mathbf{P})$ with respect to $\boldsymbol{\theta}$ for given \mathbf{P} is a simple task. Furthermore, the multiplicative separability of the valuation operator $\mathbf{V}_i^{\mathbf{P}}$ in the structural parameters³⁵ implies that while it is costly to compute $\mathbf{V}_i^{\mathbf{P}}$ for multiple values of \mathbf{P} , it is much cheaper to compute it for multiple values of $\boldsymbol{\theta}_i$. Two-step estimation methods exploit this particular structure of the model.

Let \mathbf{P}^0 be the vector with the population values of the probabilities $P_i^0(\mathbf{x}) \equiv \Pr(a_{imt} = 1 | \mathbf{x}_{mt} = \mathbf{x})$ for every firm i and any value of \mathbf{x} . Under the assumptions of "no unobserved common knowledge variables" and "single equilibrium in the data", the CCPs in \mathbf{P}^0 represent also firms' strategies in the only equilibrium that is played in the data. These probabilities can be estimated consistently using standard nonparametric methods. Let $\hat{\mathbf{P}}^0$ be a consistent nonparametric estimator of \mathbf{P}^0 . The two-step estimator of $\boldsymbol{\theta}^0$ is defined as $\hat{\boldsymbol{\theta}}_{2S} = \arg \max_{\boldsymbol{\theta}} Q(\boldsymbol{\theta}, \hat{\mathbf{P}}^0)$. Under standard regularity conditions, this two-step estimator is root-M consistent and asymptotically normal. This idea was originally exploited, for estimation of single agent problems, by Hotz and Miller (1993) and Hotz, Miller, Sanders and Smith (1994). It was expanded to the estimation of dynamic games by Aguirregabiria and Mira (2007), Bajari, Benkard and Levin (2007), Pakes, Ostrovsky and Berry (2007), and Pesendorfer and Schmidt-Dengler (2008). Pesendorfer and Schmidt-Dengler show that different estimators can be described as a general class of two-step estimators of dynamic games. An estimator within this class can be described using the following Minimum Distance (or Asymptotic Least Squares) approach:

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \left[\hat{\mathbf{P}}^0 - \Psi(\hat{\mathbf{P}}^0, \boldsymbol{\theta}) \right]' \mathbf{A}_M \left[\hat{\mathbf{P}}^0 - \Psi(\hat{\mathbf{P}}^0, \boldsymbol{\theta}) \right] \quad (30)$$

where \mathbf{A}_M is a weighting matrix. Each estimator within this general class is associated with a particular choice of the weighting matrix. The asymptotically optimal estimator within this class has a

³⁵The vector of values $\mathbf{V}_i^{\mathbf{P}}$ is equal to $\mathbf{W}_{z,i}^{\mathbf{P}} \boldsymbol{\theta}_i + \mathbf{W}_{e,i}^{\mathbf{P}} \sigma_i$, where $\mathbf{W}_i^{\mathbf{P}}$ is the matrix $(\mathbf{I} - \delta \mathbf{F}^{\mathbf{P}})^{-1} P_i * \mathbf{Z}_i^{\mathbf{P}}$, and $\mathbf{W}_{e,i}^{\mathbf{P}}$ is the vector $(\mathbf{I} - \delta \mathbf{F}^{\mathbf{P}})^{-1} P_i * e_i^{\mathbf{P}}$. Calculating $\mathbf{W}_{z,i}^{\mathbf{P}}$ and $\mathbf{W}_{e,i}^{\mathbf{P}}$ is significantly simpler than solving for the DP problem of one player, and much simpler than solving for an equilibrium of the dynamic game.

weighting matrix equal to the the inverse of $[\mathbf{I} - \partial\Psi(\mathbf{P}^0, \boldsymbol{\theta}^0)/\partial\mathbf{P}]' \Sigma_{\hat{\mathbf{P}}} [\mathbf{I} - \partial\Psi(\mathbf{P}^0, \boldsymbol{\theta}^0)/\partial\mathbf{P}]$, where $\Sigma_{\hat{\mathbf{P}}}$ is the variance matrix of the initial nonparametric estimator $\hat{\mathbf{P}}^0$. Pesendorfer and Schmidt-Dengler show that this estimator is asymptotically equivalent to the maximum ML estimator. Therefore, there is no loss of asymptotic efficiency by using a two-step estimator of the structural parameters instead of the MLE.

The main advantage of these two-step estimators is their computational simplicity. The first step is a simple nonparametric regression, and the second step is the estimation of a standard discrete choice model with a criterion function that in most applications is globally concave (e.g., such as the likelihood of a standard probit model in our entry-exit example). The main computational burden comes from the calculation of the present values $W_{z,i}^{\hat{\mathbf{P}}}(\mathbf{x})$ and $W_{e,i}^{\hat{\mathbf{P}}}(\mathbf{x})$. Though the computation of these present values may be subject to a curse of dimensionality (see Section 3.5), the cost of obtaining a two-step estimator is several orders of magnitude smaller than solving (just once) for an equilibrium of the dynamic game. In most applications, this makes the difference between being able to estimate the model or not.

These two-step estimators have some important limitations. A first limitation is the restrictions imposed by the assumption of no unobserved common knowledge variables. Ignoring persistent unobservables, if present, can generate important biases in the estimation of structural parameters. We deal with this issue in Section 3.4. A second problem is finite sample bias. The initial nonparametric estimator can be very imprecise in the samples available in actual applications, and this can generate serious finite sample biases in the two-step estimator of structural parameters. In dynamic games with heterogeneous players, the number of observable state variables is proportional to the number of players and therefore the so called *curse of dimensionality in nonparametric estimation* (and the associated bias of the two-step estimator) can be particularly serious. The source of this bias is well understood in two-step methods: $\hat{\mathbf{P}}$ enters nonlinearly in the sample moment conditions that define the estimator, and the expected value of a nonlinear function of $\hat{\mathbf{P}}$ is not equal to that function evaluated at the expected value of $\hat{\mathbf{P}}$. The larger the variance or the bias of $\hat{\mathbf{P}}$, the larger the bias of the two-step estimator of $\boldsymbol{\theta}_0$.³⁶

³⁶The asymptotically efficient two-step estimator proposed by Pesendorfer and Schmidt-Dengler does not deal with the finite sample bias problem. In fact, as it is well-known in the literature of covariance-structure models, the finite sample bias of this asymptotically optimal estimator can be significantly more severe than the one of the standard two-step estimator because the estimation of the optimal weighting matrix is also contaminated by the imprecise nonparametric estimator and it contributes to increase the finite sample bias (see Altonji, and Segal, 1996, and Horowitz 1998). In the context of two-step estimation of dynamic games, Pakes, Ostrovsky, and Berry (2007) present Monte Carlo experiments supporting this concern.

3.2.4 Recursive K-step estimators

To deal with finite sample bias, Aguirregabiria and Mira (2002, 2007) consider a recursive K-step extension. Given the two-step estimator $\hat{\theta}_{2S}$ and the initial nonparametric estimator of CCPs, $\hat{\mathbf{P}}^0$, we can construct a new estimator of CCPs, $\hat{\mathbf{P}}^1$, such that $\hat{P}_i^1(\mathbf{x}) = \Phi\left(\bar{\mathbf{z}}_i^{\hat{\mathbf{P}}^0}(\mathbf{x}) \hat{\theta}_{i,2S} + \tilde{e}_i^{\hat{\mathbf{P}}^0}(\mathbf{x})\right)$. This estimator exploits the parametric structure of the model, and the structure of best response functions. It seems intuitive that this new estimator of CCPs has better statistical properties than the initial nonparametric estimator, i.e., smaller asymptotic variance, and smaller finite sample bias and variance. As we explain below, this intuition is correct as long as the equilibrium that generated the data is (Lyapunov) stable. Under this condition, it seems natural to obtain a new two-step estimator by replacing $\hat{\mathbf{P}}^0$ with $\hat{\mathbf{P}}^1$ as the estimator of CCPs. The same argument can be applied recursively to generate a sequence of K -step estimators. Given an initial consistent nonparametric estimator $\hat{\mathbf{P}}^0$, the sequence of estimators $\{\hat{\theta}^K, \hat{\mathbf{P}}^K : K \geq 1\}$ is defined as $\hat{\theta}^K = \arg \max_{\theta} Q(\theta, \hat{\mathbf{P}}^{K-1})$, where $\hat{\mathbf{P}}^K = \Psi(\hat{\mathbf{P}}^{K-1}, \hat{\theta}^K)$. To study the properties of these K-step estimators, it is convenient to represent the sequence $\{\hat{\mathbf{P}}^K : K \geq 1\}$ as the result of iterating in a fixed point mapping. For arbitrary \mathbf{P} , define the mapping $\varphi(\mathbf{P}) \equiv \Psi(\mathbf{P}, \hat{\theta}(\mathbf{P}))$ where $\hat{\theta}(\mathbf{P}) \equiv \arg \max_{\theta} Q(\theta, \mathbf{P})$. The mapping $\varphi(\mathbf{P})$ is called the Nested Pseudo Likelihood (NPL) mapping. The sequence of estimators $\{\hat{\mathbf{P}}^K : K \geq 1\}$ can be obtained by successive iterations in the mapping φ starting with the nonparametric estimator $\hat{\mathbf{P}}^0$, i.e., for $K \geq 1$, $\hat{\mathbf{P}}^K = \varphi(\hat{\mathbf{P}}^{K-1})$.

Monte Carlo experiments in Aguirregabiria and Mira (2002, 2007), and Kasahara and Shimotsu (2008a, 2009) show that iterating in the NPL mapping can reduce significantly the finite sample bias of the two-step estimator. The Monte Carlo experiments in Pesendorfer and Schmidt-Dengler (2008) present a different, more mixed, picture. While for some of their experiments NPL iteration reduces the bias, in other experiments the bias remains constant or even increases. A closer look at the Monte Carlo experiments in Pesendorfer and Schmidt-Dengler shows that the NPL iterations provide poor results in those cases where the equilibrium that generates the data is not (Lyapunov) stable. As we explain below, this is not a coincidence. It turns out that the computational and statistical properties of the sequence of K-step estimators depend critically on the stability of the NPL mapping around the equilibrium in the data. Lyapunov stability of the NPL mapping is also important for the properties of the methods that have been proposed so far to deal with unobserved heterogeneity in the estimation of dynamic games and that we present in section 3.5. Therefore, it is important to analyze this issue in more detail here.

(a) Lyapunov stability. Let \mathbf{P}^* be a fixed point of the NPL mapping such that $\mathbf{P}^* = \varphi(\mathbf{P}^*)$.

We say that the mapping φ is Lyapunov-stable around the fixed point \mathbf{P}^* if there is a neighborhood of \mathbf{P}^* , \mathcal{N} , such that successive iterations in the mapping φ starting at $\mathbf{P} \in \mathcal{N}$ converge to \mathbf{P}^* . A necessary and sufficient condition for Lyapunov stability is that the *spectral radius* of the Jacobian matrix $\partial\varphi(\mathbf{P}^*)/\partial\mathbf{P}'$ is smaller than one. The neighboring set \mathcal{N} is denoted the dominion of attraction of the fixed point \mathbf{P}^* .³⁷ Similarly, if \mathbf{P}^* is an equilibrium of the mapping $\Psi(\cdot, \boldsymbol{\theta})$, we say that this mapping is Lyapunov stable around \mathbf{P}^* if and only if the *spectral radius* of the Jacobian matrix $\partial\Psi(\mathbf{P}^*, \boldsymbol{\theta})/\partial\mathbf{P}'$ is smaller than one.

There is a relationship between the stability of the NPL mapping and of the equilibrium mapping $\Psi(\cdot, \boldsymbol{\theta}^0)$ around \mathbf{P}^0 (i.e., the equilibrium that generates the data). The Jacobian matrices of the NPL and equilibrium mapping are related by the following expression (see Kasahara and Shimotsu, 2009): $\partial\varphi(\mathbf{P}^0)/\partial\mathbf{P}' = M(\mathbf{P}^0) \partial\Psi(\mathbf{P}^0, \boldsymbol{\theta}^0)/\partial\mathbf{P}'$, where $M(\mathbf{P}^0)$ is an idempotent projection matrix.³⁸ In single-agent dynamic programming models, the Jacobian matrix $\partial\Psi(\mathbf{P}^0, \boldsymbol{\theta}^0)/\partial\mathbf{P}'$ is zero (i.e., zero Jacobian matrix property, Aguirregabiria and Mira, 2002). Therefore, for that class of models $\partial\varphi(\mathbf{P}^0)/\partial\mathbf{P}' = 0$ and the NPL mapping is Lyapunov stable around \mathbf{P}^0 . In dynamic games, $\partial\Psi(\mathbf{P}^0, \boldsymbol{\theta}^0)/\partial\mathbf{P}'$ is not zero. However, given that $M(\mathbf{P}^0)$ is an idempotent matrix, it is possible to show that the spectral radius of $\partial\varphi(\mathbf{P}^0)/\partial\mathbf{P}'$ is not larger than the spectral radius of $\partial\Psi(\mathbf{P}^0, \boldsymbol{\theta}^0)/\partial\mathbf{P}'$. Therefore, Lyapunov stability of \mathbf{P}^0 in the equilibrium mapping implies stability of the NPL mapping.

(b) Convergence of NPL iterations. Suppose that the true equilibrium in the population, \mathbf{P}^0 , is Lyapunov stable with respect to the NPL mapping. This implies that with probability approaching one, as M goes to infinity, the (sample) NPL mapping is stable around a consistent nonparametric estimator of \mathbf{P}^0 . Therefore, the sequence of K -step estimators converges to a limit $\hat{\mathbf{P}}_{\text{lim}}^0$ that is a fixed point of the NPL mapping, i.e., $\hat{\mathbf{P}}_{\text{lim}}^0 = \varphi(\hat{\mathbf{P}}_{\text{lim}}^0)$. It is possible to show that this limit $\hat{\mathbf{P}}_{\text{lim}}^0$ is a consistent estimator of \mathbf{P}^0 (see Kasahara and Shimotsu, 2009). Therefore, under Lyapunov stability of the NPL mapping, if we start with a consistent estimator of \mathbf{P}^0 and iterate in the NPL mapping, we converge to a consistent estimator that is an equilibrium of the model. It is possible to show that this estimator is asymptotically more efficient than the two-step estimator (Aguirregabiria and Mira, 2007).

Pesendorfer and Schmidt-Dengler (2010) present an example where the sequence of K -step estimators converges to a limit estimator that is not consistent. As implied by the results presented

³⁷The spectral radius of a matrix is the maximum absolute eigenvalue. If the mapping φ is twice continuously differentiable, then the spectral radius is a continuous function of \mathbf{P} . Therefore, if φ is Lyapunov stable at \mathbf{P}^* , for any \mathbf{P} in the dominion of attraction of \mathbf{P}^* we have that the spectral radius of $\partial\varphi(\mathbf{P})/\partial\mathbf{P}'$ is also smaller than one.

³⁸The idempotent matrix $M(\mathbf{P}^0)$ is $I - \Psi_{\theta}(\Psi'_{\theta} \text{diag}\{\mathbf{P}^0\}^{-1} \Psi_{\theta})^{-1} \Psi'_{\theta} \text{diag}\{\mathbf{P}^0\}^{-1}$, where $\Psi_{\theta} \equiv \partial\Psi(\mathbf{P}^0, \boldsymbol{\theta}^0)/\partial\boldsymbol{\theta}'$.

above, the equilibrium that generates the data in their example is not Lyapunov stable.

The concept of Lyapunov stability of the best response mapping at an equilibrium means that if we marginally perturb players' strategies, and then allow players to best respond to the new strategies, then we will converge to the original equilibrium. To us this seems like a plausible equilibrium selection criterion. Ultimately, whether an unstable equilibrium is interesting depends on the application and the researchers taste. Nevertheless, at the end of this section we present simple modified versions of the NPL method that can deal with data generated from an equilibrium that is not stable.

(c) Reduction of finite sample bias. Kasahara and Shimotsu (2008a, 2009) derive a second order approximation to the bias of the K-step estimators. They show that the key component in this bias is the distance between the first step and the second step estimators of \mathbf{P}^0 , i.e., $\left\| \varphi(\hat{\mathbf{P}}^0) - \hat{\mathbf{P}}^0 \right\|$. An estimator that reduces this distance is an estimator with lower finite sample bias. Therefore, based on our discussion in point (b) above, the sequence of K-step estimators are decreasing in their finite sample bias if and only if the NPL mapping is Lyapunov stable around \mathbf{P}^0 . The Monte Carlo experiments in Pesendorfer and Schmidt-Dengler (2008) illustrate this point. They implement experiments using different DGPs: in some of them the data is generated from a stable equilibrium, and in others the data come from a non-stable equilibrium. It is simple to verify (see Aguirregabiria and Mira, 2010) that the experiments where NPL iterations do not reduce the finite sample bias are those where the equilibrium that generates the data is not (Lyapunov) stable.

(d) Modified NPL algorithms. Note that Lyapunov stability can be tested after obtaining the first NPL iteration. Once we have obtained the two-step estimator, we can calculate the Jacobian matrix $\partial\varphi(\hat{\mathbf{P}}^0)/\partial\mathbf{P}'$ and its eigenvalues, and then check whether Lyapunov stability holds at $\hat{\mathbf{P}}^0$. If the applied researcher considers that his data may have been generated by an equilibrium that is not stable, then it will be worthwhile to compute this Jacobian matrix and its eigenvalues. If Lyapunov stability holds at $\hat{\mathbf{P}}^0$, then we know that NPL iterations reduce the bias of the estimator and converge to a consistent estimator. When the condition does not hold, then the solution to this problem is not simple. Though the researcher might choose to use the two-step estimator, the non-stability of the equilibrium has also important negative implications on the properties of this simple estimator.³⁹ In this context, Kasahara and Shimotsu (2009) propose alternative

³⁹Non-stability of the NPL mapping at \mathbf{P}^0 implies that the asymptotic variance of the two-step estimator of \mathbf{P}^0 is larger than asymptotic variance of the nonparametric reduced form estimator. To see this, note that the two-step estimator of CCPs is $\hat{\mathbf{P}}^1 = \varphi(\hat{\mathbf{P}}^0)$, and applying the delta method we have that $Var(\hat{\mathbf{P}}^1) = [\partial\varphi(\mathbf{P}^0)/\partial\mathbf{P}'] Var(\hat{\mathbf{P}}^0) [\partial\varphi(\mathbf{P}^0)/\partial\mathbf{P}']'$. If the spectral radius of $\partial\varphi(\mathbf{P}^0)/\partial\mathbf{P}'$ is greater than 1, then $Var(\hat{\mathbf{P}}^1) > Var(\hat{\mathbf{P}}^0)$. This is a puzzling

recursive estimators based on fixed-point mappings other than the NPL that, by construction, are stable. Iterating in these alternative mappings is significantly more costly than iterating in the NPL mapping, but these iterations guarantee reduction of the finite sample bias and convergence to a consistent estimator.

Aguirregabiria and Mira (2010) propose two modified versions of the NPL algorithm that are simple to implement and that always converge to a consistent estimator with better properties than two-step estimators. A *first modified-NPL-algorithm* applies to dynamic games. The first NPL iteration is standard but in every successive iteration best response mappings are used to update guesses of each player’s own future behavior without updating beliefs about the strategies of the other players. This algorithm always converges to a consistent estimator even if the equilibrium generating the data is not stable and it reduces monotonically the asymptotic variance and the finite sample bias of the two-step estimator. The *second modified-NPL-algorithm* applies to static games and it consists in the application of the standard NPL algorithm both to the best response mapping and to the inverse of this mapping. If the equilibrium that generates the data is unstable in the best response mapping, it should be stable in the inverse mapping. Therefore, the NPL applied to the inverse mapping should converge to the consistent estimator and should have a largest value of the pseudo likelihood that the estimator that we converge to when applying the NPL algorithm to the best response mapping. Aguirregabiria and Mira illustrate the performance of these estimators using the examples in Pesendorfer and Schmidt-Dengler (2008 and 2010).

3.3 Dealing with Unobserved Heterogeneity

So far, we have maintained the assumption that the only unobservables for the researcher are the private information shocks that are i.i.d. over firms, markets, and time. In most applications in IO, this assumption is not realistic and it can be easily rejected by the data. Markets and firms are heterogenous in terms of characteristics that are payoff-relevant for firms but unobserved to the researcher. Not accounting for this heterogeneity may generate significant biases in parameter estimates and in our understanding of competition in the industry. For instance, in the empirical applications in Aguirregabiria and Mira (2007) and Collard-Wexler (2006), the estimation of a model without unobserved market heterogeneity implies estimates of strategic interaction between firms (i.e., competition effects) that are close to zero or even have the opposite sign to the one expected under competition. In both applications, including unobserved heterogeneity in the models

result because the estimator $\hat{\mathbf{P}}^0$ is nonparametric while the estimator $\hat{\mathbf{P}}^1$ exploits most of the structure of the model. Therefore, the non-stability of the equilibrium that generates the data is an issue for this general class of two-step or sequential estimators.

results in estimates that show significant and strong competition effects.

Aguirregabiria and Mira (2007), Aguirregabiria, Mira, and Roman (2007), and Arcidiacono and Miller (2008) have proposed methods for the estimation of dynamic games that allow for persistent unobserved heterogeneity in players or markets. Here we concentrate on the case of permanent unobserved market heterogeneity in the profit function. Arcidiacono and Miller (2008) propose a method that combines the NPL method, that we present here, with an EM algorithm, and they consider a more general framework that includes unobserved heterogeneity that can vary over time according to Markov chain process and that can enter both in the payoff function and in the transition of the state variables.⁴⁰

Consider our entry-exit model, but now the profit of firm i if active in market m includes a term ξ_m that is unobserved to the researcher:

$$\Pi_{imt} = H_{mt} \sum_{n=0}^{N-1} 1 \left\{ \sum_{j \neq i} a_{jmt} = n \right\} \theta_{i,n}^{VP} - \theta_i^{FC} - (1 - s_{imt}) \theta_i^{EC} - \sigma_{\xi_i} \xi_m - \varepsilon_{imt} \quad (31)$$

σ_{ξ_i} is a parameter, and ξ_m is a time-invariant 'random effect' that is common knowledge to the players but unobserved to the researcher.⁴¹ The distribution of this random effect has the following properties: (A.1) it has a discrete and finite support $\{\xi^1, \xi^2, \dots, \xi^L\}$, each value in the support of ξ represents a 'market type', and we index market types by $\ell \in \{1, 2, \dots, L\}$; (A.2) it is i.i.d. over markets with probability mass function $\lambda_\ell \equiv \Pr(\xi_m = \xi^\ell)$; and (A.3) it does not enter into the transition probability of the observed state variables, i.e., $\Pr(\mathbf{x}_{mt+1} \mid \mathbf{x}_{mt}, \mathbf{a}_{mt}, \xi_m) = F_x(\mathbf{x}_{mt+1} \mid \mathbf{x}_{mt}, \mathbf{a}_{mt})$. Without loss of generality, ξ_m has mean zero and unit variance because the mean and the variance of ξ_m are incorporated in the parameters θ_i^{FC} and σ_{ξ_i} , respectively. Also, without loss of generality, the researcher knows the points of support $\{\xi^\ell : \ell = 1, 2, \dots, L\}$ though the probability mass function $\{\lambda_\ell\}$ is unknown.

Assumption (A.1) is common when dealing with permanent unobserved heterogeneity in dynamic structural models. The discrete support of the unobservable implies that the contribution of a market to the likelihood (or pseudo likelihood) function is a finite mixture of likelihoods under the different possible best responses that we would have for each possible market type. With continuous support we would have an infinite mixture of best responses and this could complicate significantly the computation of the likelihood. Nevertheless, as we illustrate before, using a pseudo likelihood approach and a convenient parametric specification of the distribution of ξ_m simplifies

⁴⁰In fact, the framework that we present here can be generalized to include unobserved market heterogeneity that varies over time according to a Markov chain with finite support.

⁴¹In this example, we include unobserved heterogeneity only in the fixed cost. However, the estimation methods that we present here can deal with richer forms of unobserved heterogeneity, e.g., in fixed costs, variable profits, and entry costs.

this computation such that we can consider many values in the support of this unobserved variable at a low computational cost. Assumption (A.2) is also standard when dealing with unobserved heterogeneity. Unobserved spatial correlation across markets does not generate inconsistency of the estimators that we present here because the likelihood equations that define the estimators are still valid moment conditions under spatial correlation. Incorporating spatial correlation in the model, if present in the data, would improve the efficiency of the estimator but at a significant computational cost. Assumption (A.3) can be relaxed, and in fact the method by Arcidiacono-Miller deals with unobserved heterogeneity both in payoffs and transition probabilities.

Each market type ℓ has its own equilibrium mapping (with a different level of profits given ξ^ℓ) and its own equilibrium. Let \mathbf{P}_ℓ be a vector of strategies (CCPs) in market-type ℓ : $\mathbf{P}_\ell \equiv \{P_{i\ell}(\mathbf{x}_t) : i = 1, 2, \dots, N; \mathbf{x}_t \in \mathcal{X}\}$.⁴² It is straightforward to extend the description of an equilibrium mapping in CCPs to this model. A vector of CCPs \mathbf{P}_ℓ is a MPE for market type ℓ if and only if for every firm i and every state \mathbf{x}_t we have that: $P_{i\ell}(\mathbf{x}_t) = \Phi\left(\tilde{\mathbf{z}}_i^{\mathbf{P}_\ell}(\mathbf{x}_t, \xi^\ell) \boldsymbol{\theta}_i + \tilde{e}_i^{\mathbf{P}_\ell}(\mathbf{x}_t, \xi^\ell)\right)$, where now the vector of structural parameters $\boldsymbol{\theta}_i$ is $\{\theta_{i,0}^{VP}, \dots, \theta_{i,N-1}^{VP}, \theta_i^{FC}, \theta_i^{EC}, \sigma_{\xi_i}\}$ that includes σ_{ξ_i} , and the vector $\tilde{\mathbf{z}}_i^{\mathbf{P}_\ell}(\mathbf{x}_t, \xi^\ell)$ has a similar definition as before with the only difference that it has one more component associated with $-\xi^\ell$. Since the points of support $\{\xi^\ell : \ell = 1, 2, \dots, L\}$ are known to the researcher, he can construct the equilibrium mapping for each market type.

Let λ be the vector of parameters in the probability mass function of ξ , i.e., $\lambda \equiv \{\lambda_\ell : \ell = 1, 2, \dots, L\}$, and let \mathbf{P} be the set of CCPs for every market type, $\{\mathbf{P}_\ell : \ell = 1, 2, \dots, L\}$. The (conditional) pseudo log likelihood function of this model is $Q(\boldsymbol{\theta}, \lambda, \mathbf{P}) = \sum_{m=1}^M \log \Pr(\mathbf{a}_{m1}, \mathbf{a}_{m2}, \dots, \mathbf{a}_{mT} \mid \mathbf{x}_{m1}, \mathbf{x}_{m2}, \dots, \mathbf{x}_{mT}; \boldsymbol{\theta}, \lambda, \mathbf{P})$. We can write this function as $\sum_{m=1}^M \log q_m(\boldsymbol{\theta}, \lambda, \mathbf{P})$, where $q_m(\boldsymbol{\theta}, \lambda, \mathbf{P})$ is the contribution of market m to the pseudo likelihood:

$$q_m(\boldsymbol{\theta}, \lambda, \mathbf{P}) = \sum_{\ell=1}^L \lambda_{\ell|\mathbf{x}_{m1}} \left[\prod_{i,t} \Phi\left(\tilde{\mathbf{z}}_{imlt}^{\mathbf{P}_\ell} \boldsymbol{\theta}_i + \tilde{e}_{imlt}^{\mathbf{P}_\ell}\right)^{a_{imt}} \Phi\left(-\tilde{\mathbf{z}}_{imlt}^{\mathbf{P}_\ell} \boldsymbol{\theta}_i - \tilde{e}_{imlt}^{\mathbf{P}_\ell}\right)^{1-a_{imt}} \right] \quad (32)$$

where $\tilde{\mathbf{z}}_{imlt}^{\mathbf{P}_\ell} \equiv \tilde{\mathbf{z}}_i^{\mathbf{P}_\ell}(\mathbf{x}_{mt}, \xi^\ell)$, $\tilde{e}_{imlt}^{\mathbf{P}_\ell} \equiv \tilde{e}_i^{\mathbf{P}_\ell}(\mathbf{x}_{mt}, \xi^\ell)$, and $\lambda_{\ell|\mathbf{x}}$ is the conditional probability $\Pr(\xi_m = \xi^\ell \mid \mathbf{x}_{m1} = \mathbf{x})$. The conditional probability distribution $\lambda_{\ell|\mathbf{x}}$ is different from the unconditional distribution λ_ℓ . In particular, ξ_m is not independent of the predetermined endogenous state variables that represent market structure. For instance, we expect a negative correlation between the indicators of incumbent status, s_{imt} , and the unobserved component of the fixed cost ξ_m , i.e., markets where it is more costly to operate tend to have a smaller number of incumbent firms. This is the so called *initial conditions problem* (Heckman, 1981). In short panels (for T relatively small), not

⁴²The introduction of unobserved market heterogeneity also implies that we can relax the assumption of only ‘a single equilibrium in the data’ to allow for different market types to have different equilibria.

taking into account this dependence between ξ_m and \mathbf{x}_{m1} can generate significant biases, similar to the biases associated to ignoring the existence of unobserved market heterogeneity. There are different ways to deal with the initial conditions problem in dynamic models (Heckman, 1981). One possible approach is to derive the joint distribution of \mathbf{x}_{m1} and ξ_m implied by the equilibrium of the model. That is the approach proposed and applied in Aguirregabiria and Mira (2007) and Collard-Wexler (2006). Let $\mathbf{p}^{\mathbf{P}^\ell} \equiv \{p^{\mathbf{P}^\ell}(\mathbf{x}_t) : \mathbf{x}_t \in \mathcal{X}\}$ be the ergodic or steady-state distribution of \mathbf{x}_t induced by the equilibrium \mathbf{P}^ℓ and the transition F_x . This stationary distribution can be simply obtained as the solution to the following system of linear equations: for every value $\mathbf{x}_t \in \mathcal{X}$, $p^{\mathbf{P}^\ell}(\mathbf{x}_t) = \sum_{\mathbf{x}_{t-1} \in \mathcal{X}} p^{\mathbf{P}^\ell}(\mathbf{x}_{t-1}) F_x^{\mathbf{P}^\ell}(\mathbf{x}_t | \mathbf{x}_{t-1})$, or in vector form, $\mathbf{p}^{\mathbf{P}^\ell} = \mathbf{F}_x^{\mathbf{P}^\ell} \mathbf{p}^{\mathbf{P}^\ell}$ subject to $\mathbf{p}^{\mathbf{P}^\ell} \mathbf{1} = 1$. Given the ergodic distributions for the L market types, we can apply Bayes' rule to obtain:

$$\lambda_{\ell|\mathbf{x}_{m1}} = \frac{\lambda_\ell p^{\mathbf{P}^\ell}(\mathbf{x}_{m1})}{\sum_{\ell'=1}^L \lambda_{\ell'} p^{\mathbf{P}^{\ell'}}(\mathbf{x}_{m1})} \quad (33)$$

Note that given the CCPs $\{\mathbf{P}^\ell\}$, this conditional distribution does not depend on parameters in the vector $\boldsymbol{\theta}$, only on the distribution λ . Given this expression for the probabilities $\{\lambda_{\ell|\mathbf{x}_{m1}}\}$, we have that the pseudo likelihood in (32) only depends on the structural parameters $\boldsymbol{\theta}$ and $\boldsymbol{\lambda}$ and the incidental parameters \mathbf{P} .

For the estimators that we discuss here, we maximize $Q(\boldsymbol{\theta}, \lambda, \mathbf{P})$ with respect to $(\boldsymbol{\theta}, \lambda)$ for given \mathbf{P} . Therefore, the ergodic distributions $p^{\mathbf{P}^\ell}$ are fixed during this optimization. This implies a significant reduction in the computational cost associated with the initial conditions problem. Nevertheless, in the literature of finite mixture models, it is well known that optimization of the likelihood function with respect to the mixture probabilities λ is a complicated task because the problem is plagued with many local maxima and minima. To deal with this problem, Aguirregabiria and Mira (2007) introduce an additional parametric assumption on the distribution of ξ_m that simplifies significantly the maximization of $Q(\boldsymbol{\theta}, \lambda, \mathbf{P})$ for fixed \mathbf{P} . They assume that the probability distribution of unobserved market heterogeneity is such that the only unknown parameters for the researcher are the mean and the variance which are included in θ_i^{FC} and $\sigma_{\xi i}$, respectively. Therefore, they assume that the distribution of ξ_m (i.e., the points of support and the probabilities λ_ℓ) are known to the researcher. For instance, we may assume that ξ_m has a discretized standard normal distribution with an arbitrary number of points of support L . Under this assumption, the pseudo likelihood function is maximized only with respect to $\boldsymbol{\theta}$ for given \mathbf{P} . Avoiding optimization with respect to λ simplifies importantly the computation of the different estimators that we describe below.

NPL estimator. As defined above, the NPL mapping $\varphi(\cdot)$ is the composition of the equilibrium mapping and the mapping that provides the maximand in $\boldsymbol{\theta}$ to $Q(\boldsymbol{\theta}, \mathbf{P})$ for given \mathbf{P} . That is, $\varphi(\mathbf{P}) \equiv \Psi(\hat{\boldsymbol{\theta}}(\mathbf{P}), \mathbf{P})$ where $\hat{\boldsymbol{\theta}}(\mathbf{P}) \equiv \arg \max_{\boldsymbol{\theta}} Q(\boldsymbol{\theta}, \mathbf{P})$. By definition, an NPL fixed point is a pair $(\hat{\boldsymbol{\theta}}, \hat{\mathbf{P}})$ that satisfies two conditions: (a) $\hat{\boldsymbol{\theta}}$ maximizes $Q(\boldsymbol{\theta}, \hat{\mathbf{P}})$; and (b) $\hat{\mathbf{P}}$ is an equilibrium associated to $\hat{\boldsymbol{\theta}}$. The NPL estimator is defined as the NPL fixed point with the maximum value of the likelihood function. The NPL estimator is consistent under standard regularity conditions (Aguirregabiria and Mira, 2007, Proposition 2).

When the equilibrium that generates the data is Lyapunov stable, we can compute the NPL estimator using a procedure that iterates in the NPL mapping, as described in section 3.2 to obtain the sequence of K-step estimators (i.e., NPL algorithm). The main difference is that now we have to calculate the steady-state distributions $\mathbf{p}(\mathbf{P}_\ell)$ to deal with the initial conditions problem. However, the pseudo likelihood approach also reduces significantly the cost of dealing with the initial conditions problem. This NPL algorithm proceeds as follows. We start with L arbitrary vectors of players' choice probabilities, one for each market type: $\{\hat{\mathbf{P}}_\ell^0 : \ell = 1, 2, \dots, L\}$. Then, we perform the following steps. Step 1: For every market type we obtain the steady-state distributions and the probabilities $\{\lambda_{\ell|x_{m1}}\}$. Step 2: We obtain a pseudo maximum likelihood estimator of $\boldsymbol{\theta}$ as $\hat{\boldsymbol{\theta}}^1 = \arg \max_{\boldsymbol{\theta}} Q(\boldsymbol{\theta}, \hat{\mathbf{P}}^0)$. Step 3: Update the vector of players' choice probabilities using the best response probability mapping. That is, for market type ℓ , firm i and state \mathbf{x} , $\hat{P}_{i\ell}^1(\mathbf{x}) = \Phi(\tilde{\mathbf{z}}_i^{\hat{\mathbf{P}}_\ell^0}(\mathbf{x}, \xi^\ell) \hat{\boldsymbol{\theta}}_i^1 + \tilde{e}_i^{\hat{\mathbf{P}}_\ell^0}(\mathbf{x}, \xi^\ell))$. If, for every type ℓ , $\|\hat{\mathbf{P}}_\ell^1 - \hat{\mathbf{P}}_\ell^0\|$ is smaller than a predetermined small constant, then stop the iterative procedure and keep $\hat{\boldsymbol{\theta}}^1$ as a candidate estimator. Otherwise, repeat steps 1 to 4 using $\hat{\mathbf{P}}^1$ instead of $\hat{\mathbf{P}}^0$.

The NPL algorithm, upon convergence, finds an NPL fixed point. To guarantee consistency, the researcher needs to start the NPL algorithm from different CCP's in case there are multiple NPL fixed points. This situation is similar to using a gradient algorithm, designed to find a local root, in order to obtain an estimator which is defined as a global root. Of course, this global search aspect of the method makes it significantly more costly than the application of the NPL algorithm in models without unobserved heterogeneity. This is the additional computational cost that we have to pay for dealing with unobserved heterogeneity. Note, however, that this global search can be parallelized in a computer with multiple processors.

Arcidiacono and Miller (2008) extend this approach in several interesting and useful ways. First, they consider a more general form of unobserved heterogeneity that may enter both in the payoff function and in the transition of the state variables. Second, to deal with the complexity in

the optimization of the likelihood function with respect to the distribution of the finite mixture, they combine the NPL method with an EM algorithm. Third, they show that for a class of dynamic decision models, that includes but it is not limited to optimal stopping problems, the computation of the inclusive values $\tilde{z}_{imlt}^{\mathbf{P}_\ell}$ and $\tilde{e}_{imlt}^{\mathbf{P}_\ell}$ is simple and it is not subject to a 'curse of dimensionality', i.e., the cost of computing these value for given \mathbf{P}_ℓ does not increase exponentially with the dimension of the state space. Together, these results provide a relatively simple approach to estimate dynamic games with unobserved heterogeneity of finite mixture type. Note that Lyapunov stability of each equilibrium type that generates the data is a necessary condition for the NPL and the Arcidiacono-Miller algorithms to converge to a consistent estimator.

Kasahara and Shimotsu (2008). The estimators of finite mixture models in Aguirregabiria and Mira (2007) and Arcidiacono and Miller (2008) consider that the researcher cannot obtain consistent nonparametric estimates of market-type CCPs $\{\mathbf{P}_\ell^0\}$. Kasahara and Shimotsu (2008b), based on previous work by Hall and Zhou (2003), have derived sufficient conditions for the nonparametric identification of market-type CCPs $\{\mathbf{P}_\ell^0\}$ and the probability distribution of market types, $\{\lambda_\ell^0\}$. Given the nonparametric identification of market-type CCPs, it is possible to estimate structural parameters using a two-step approach similar to the one described above. However, this two-step estimator has three limitations that do not appear in two-step estimators without unobserved market heterogeneity. First, the conditions for nonparametric identification of \mathbf{P}^0 may not hold. Second, the nonparametric estimator in the first step is a complex estimator from a computational point of view. In particular, it requires the minimization of a sample criterion function with respect to the large dimensional object \mathbf{P} .⁴³ This is in fact the type of computational problem that we wanted to avoid by using two-step methods instead of standard ML or GMM. Finally, the finite sample bias of the two-step estimator can be significantly more severe when \mathbf{P}^0 incorporates unobserved heterogeneity and we estimate it nonparametrically.

3.4 Reducing the State Space

Although two-step and sequential methods are computationally much cheaper than full solution-estimation methods, they are still impractical for applications where the dimension of the state space is large. The cost of computing exactly the matrix of present values $\mathbf{W}_{z,i}^{\mathbf{P}}$, increases cubically with the dimension of the state space. In the context of dynamic games, the dimension of the state space increases exponentially with the number of heterogeneous players. Therefore, the cost of

⁴³Furthermore, the criterion function is not globally concave/convex and its optimization requires a global search.

computing the matrix of present values may become intractable even for a relatively small number of players.

A simple approach to deal with this curse of dimensionality is to assume that players are homogeneous and the equilibrium is symmetric. For instance, in our dynamic game of market entry-exit, when firms are heterogeneous, the dimension of the state space is $|H| * 2^N$, where $|H|$ is the number of values in the support of market size H_t . To reduce the dimensionality of the state space, we need to assume that: (a) only the number of competitors (and not their identities) affects the profit of a firm; (b) firms are homogeneous in their profit function; and (c) the selected equilibrium is symmetric. Under these conditions, the payoff relevant state variables for a firm i are $\{H_t, s_{it}, n_{t-1}\}$ where s_{it} is its own incumbent status, and n_{t-1} is the total number of active firms at period $t - 1$. The dimension of the state space is $|H| * 2 * (N + 1)$ that increases only linearly with the number of players.⁴⁴ It is clear that the assumption of homogeneous firms and symmetric equilibrium can reduce substantially the dimension of the state space, and it can be useful in some empirical applications. Nevertheless, there are many applications where this assumption is too strong. For instance, in applications where firms produce differentiated products.

To deal with this issue, Hotz, Miller, Sanders and Smith (1994) proposed an estimator that uses Monte Carlo simulation techniques to approximate the values $\mathbf{W}_{z,i}^P$. Bajari, Benkard, and Levin (2007) have extended this method to dynamic games and to models with continuous decision variables. This approach has proved useful in some applications. Nevertheless, it is important to be aware that in those applications with large state spaces, simulation error can be sizeable and it can induce biases in the estimation of the structural parameters. In those cases, it is worthwhile to reduce the dimension of the state space by making additional structural assumptions. That is the general idea in the inclusive-value approach that we have discussed in section 2 and that can be extended to the estimation of dynamic games. Different versions of this idea have been proposed and applied by Nevo and Rossi (2008), Maceria (2007), Rossi (2009), and Aguirregabiria and Ho (2009).

To present the main ideas, we consider here a dynamic game of quality competition in the spirit of Pakes and McGuire (1994), the differentiated product version of Ericson-Pakes model. There are N firms in the market, that we index by i , and B brands or differentiated products, that we index by b . The set of brands sold by firm i is $\mathcal{B}_i \subset \{1, 2, \dots, B\}$. Demand is given by a model similar to that of Section 2.1: consumers choose one of the B products offered in the market, or the outside good. The utility that consumer h obtains from purchasing product b at time t is

⁴⁴This is a particular example of the 'exchangeability assumption' proposed by Pakes and McGuire (2001).

$U_{hbt} = x_{bt} - \alpha p_{bt} + u_{hbt}$, where x_{bt} is the quality of the product, p_{bt} is the price, α is a parameter, and u_{hbt} represents consumer specific taste for product b . These idiosyncratic errors are identically and independently distributed over (h, b, t) with type I extreme value distribution. If the consumer decides not to purchase any of the goods, she chooses the outside option that has a mean utility normalized to zero. Therefore, the aggregate demand for product b is $q_{bt} = H_t \exp\{x_{bt} - \alpha p_{bt}\} [1 + \sum_{b'=1}^B \exp\{x_{b't} - \alpha p_{b't}\}]^{-1}$, where H_t represents market size at period t . The market structure of the industry at time t is characterized by the vector $\mathbf{x}_t = (H_t, x_{1t}, x_{2t}, \dots, x_{Bt})$. Every period, firms take as given current market structure and decide simultaneously their current prices and their investment in quality improvement. The one-period profit of firm i can be written as

$$\Pi_{it} = \sum_{b \in \mathcal{B}_i} (p_{bt} - mc_b) q_{bt} - FC_b - (c_b + \varepsilon_{bt}) a_{bt} \quad (34)$$

where $a_{bt} \in \{0, 1\}$ is the binary variable that represents the decision to invest in quality improvement of product b ; mc_b , FC_b , and c_b are structural parameters that represent marginal cost, fixed operating cost, and quality investment cost for product b , respectively; and ε_{bt} is an iid private information shock in the investment cost. Product quality evolves according to a transition probability $f_x(x_{bt+1}|a_{bt}, x_{bt})$. For instance, in Pakes-McGuire model, $x_{bt+1} = x_{bt} - \zeta_t + a_{bt} v_{bt}$ where ζ_t and v_{bt} are two independent and non-negative random variables that are independently and identically distributed over (b, t) .

In this model, price competition is static. The Nash-Bertrand equilibrium determines prices and quantities as functions of market structure \mathbf{x}_t , i.e., $p_b^*(\mathbf{x}_t)$ and $q_b^*(\mathbf{x}_t)$. Firms' quality choices are the result of a dynamic game. The one-period profit function of firm i in this dynamic game is $\Pi_i(\mathbf{a}_{it}, \mathbf{x}_t) = \sum_{b \in \mathcal{B}_i} (p_b^*(\mathbf{x}_t) - mc_b) q_b^*(\mathbf{x}_t) - FC_b - (c_b + \varepsilon_{bt}) a_{bt}$, where $\mathbf{a}_{it} \equiv \{a_{bt} : b \in \mathcal{B}_i\}$. This dynamic game of quality competition has the same structure as the game that we have described in Section 3.1 and it can be solved and estimated using the same methods. However, the dimension of the state space increases exponentially with the number of products and the solution and estimation of the model becomes impractical even when B is not too large.

Define the *cost adjusted inclusive value* of firm i at period t as $\omega_{it} \equiv \log[\sum_{b \in \mathcal{B}_i} \exp\{x_{bt} - \alpha mc_b\}]$. This value is closely related to the inclusive value that we have discussed in Section 2. It can be interpreted as the net quality level, or a value added of sort, that the firm is able to produce in the market. Under the assumptions of the model, the variable profit of firm i in the Nash-Bertrand equilibrium can be written as a function of the vector of inclusive values $\boldsymbol{\omega}_t \equiv (\omega_{1t}, \omega_{2t}, \dots, \omega_{Nt}) \in \Omega$, i.e., $\sum_{b \in \mathcal{B}_i} (p_b^*(\mathbf{x}_t) - mc_b) q_b^*(\mathbf{x}_t) = vp_i(\boldsymbol{\omega}_t)$. Therefore, the one-period profit Π_{it} is a function $\tilde{\Pi}_i(\mathbf{a}_{it}, \boldsymbol{\omega}_t)$. The following assumption is similar to Assumption A2 made in Section 2 and

it establishes that given vector ω_t the rest of the information contained in the in \mathbf{x}_t is redundant for the prediction of future values of ω .

Assumption: The transition probability of the vector of inclusive values ω_t from the point of view a firm (i.e., conditional on a firm's choice) is such that $\Pr(\omega_{t+1} \mid \mathbf{a}_{it}, \mathbf{x}_t) = \Pr(\omega_{t+1} \mid \mathbf{a}_{it}, \omega_t)$.

Under these assumptions, ω_t is the vector of payoff relevant state variables in the dynamic game. The dimension of the space Ω increases exponentially with the number of firms but not with the number of brands. Therefore, the dimension of Ω can be much smaller than the dimension of the original state space of \mathbf{x}_t in applications where the number of brands is large relative to the number of firms.

Of course, the assumption of sufficiency of ω_t in the prediction of next period ω_{t+1} is not trivial. In order to justify it we can put quite strong restrictions on the stochastic process of quality levels. Alternatively, it can be interpreted in terms of limited information, and/or bounded rationality. For instance, a possible way to justify this assumption is that firms face the same type of computational burdens that we do. Limiting the information that they use in their strategies reduces a firm's computational cost of calculating a best response.

Note that the dimension of the space of ω_t still increases exponentially with the number of firms. To deal with this curse of dimensionality, Aguirregabiria and Ho (2009) consider a stronger inclusive value / sufficiency assumption. Let vp_{it} the variable profit of firm i at period t . Assumption: $\Pr(\omega_{it+1}, vp_{it+1} \mid \mathbf{a}_{it}, \mathbf{x}_t) = \Pr(\omega_{it+1}, vp_{it+1} \mid \mathbf{a}_{it}, \omega_{it}, vp_{it})$. Under this assumption, the vector of payoff relevant state variables in the decision problem of firm i is (ω_{it}, vp_{it}) and the dimension of the space of (ω_{it}, vp_{it}) does not increase with the number of firms.

3.5 Counterfactual experiments with multiple equilibria

One of the attractive features of structural models is that they can be used to predict the effects of new counterfactual policies. This is a challenging exercise in a model with multiple equilibria. Under the assumption that our data has been generated by a single equilibrium, we can use the data to identify which of the multiple equilibria is the one that we observe. However, even under that assumption, we still do not know which equilibrium will be selected when the values of the structural parameters are different to the ones that we have estimated from the data. For some models, a possible approach to deal with this issue is to calculate all of the equilibria in the counterfactual scenario and then draw conclusions that are robust to whatever equilibrium is selected. However, this approach is of limited applicability in dynamic games of oligopoly competition because the

different equilibria typically provide contradictory predictions for the effects we want to measure.

Here we describe a simple homotopy method that has been proposed in Aguirregabiria (2009) and applied in Aguirregabiria and Ho (2009). Under the assumption that the equilibrium selection mechanism, which is unknown to the researcher, is a smooth function of the structural parameters, we show how to obtain a Taylor approximation to the counterfactual equilibrium. Despite the equilibrium selection function is unknown, a Taylor approximation of that function, evaluated at the estimated equilibrium, depends on objects that the researcher knows.

Let $\Psi(\boldsymbol{\theta}, \mathbf{P})$ be the equilibrium mapping such that an equilibrium associated with $\boldsymbol{\theta}$ can be represented as a fixed point $\mathbf{P} = \Psi(\boldsymbol{\theta}, \mathbf{P})$. Suppose that there is an equilibrium selection mechanism in the population under study, but we do not know that mechanism. Let $\boldsymbol{\pi}(\boldsymbol{\theta})$ be the selected equilibrium given $\boldsymbol{\theta}$. The approach here is quite agnostic with respect to this equilibrium selection mechanism: it only assumes that there is such a mechanism, and that it is a smooth function of $\boldsymbol{\theta}$. Since we do not know the mechanism, we do not know the form of the mapping $\boldsymbol{\pi}(\boldsymbol{\theta})$ for every possible $\boldsymbol{\theta}$. However, we know that the equilibrium in the population, \mathbf{P}^0 , and the vector of the structural parameters in the population, $\boldsymbol{\theta}^0$, belong to the graph of that mapping, i.e., $\mathbf{P}^0 = \boldsymbol{\pi}(\boldsymbol{\theta}^0)$.

Let $\boldsymbol{\theta}^*$ be the vector of parameters under the counterfactual experiment that we want to analyze. We want to know the counterfactual equilibrium $\boldsymbol{\pi}(\boldsymbol{\theta}^*)$ and compare it to the factual equilibrium $\boldsymbol{\pi}(\boldsymbol{\theta}^0)$. Suppose that Ψ is twice continuously differentiable in θ and P . The following is the key assumption to implement the homotopy method that we describe here.

Assumption: The equilibrium selection mechanism is such that $\boldsymbol{\pi}(\cdot)$ is a continuous differentiable function within a convex subset of Θ that includes $\boldsymbol{\theta}^0$ and $\boldsymbol{\theta}^*$.

That is, the equilibrium selection mechanism does not "jump" between the possible equilibria when we move over the parameter space from $\boldsymbol{\theta}^0$ to $\boldsymbol{\theta}^*$. This seems a reasonable condition when the researcher is interested in evaluating the effects of a change in the structural parameters but "keeping constant" the same equilibrium type as the one that generates the data.

Under these conditions, we can make a Taylor approximation to $\boldsymbol{\pi}(\boldsymbol{\theta}^*)$ around $\boldsymbol{\theta}^0$ to obtain:

$$\boldsymbol{\pi}(\boldsymbol{\theta}^*) = \boldsymbol{\pi}(\boldsymbol{\theta}^0) + \frac{\partial \boldsymbol{\pi}(\boldsymbol{\theta}^0)}{\partial \boldsymbol{\theta}'} (\boldsymbol{\theta}^* - \boldsymbol{\theta}^0) + O(\|\boldsymbol{\theta}^* - \boldsymbol{\theta}^0\|^2) \quad (35)$$

We know that $\boldsymbol{\pi}(\boldsymbol{\theta}^0) = \mathbf{P}^0$. Furthermore, by the implicit function theorem, $\partial \boldsymbol{\pi}(\boldsymbol{\theta}^0) / \partial \boldsymbol{\theta}' = \partial \Psi(\boldsymbol{\theta}^0, \mathbf{P}^0) / \partial \boldsymbol{\theta}' + \partial \Psi(\boldsymbol{\theta}^0, \mathbf{P}^0) / \partial \mathbf{P}' \partial \boldsymbol{\pi}(\boldsymbol{\theta}^0) / \partial \boldsymbol{\theta}'$. If \mathbf{P}^0 is not a singular equilibrium then $I - \partial \Psi(\boldsymbol{\theta}^0, \mathbf{P}^0) / \partial \mathbf{P}'$ is not a singular matrix and $\partial \boldsymbol{\pi}(\boldsymbol{\theta}^0) / \partial \boldsymbol{\theta}' = (I - \partial \Psi(\boldsymbol{\theta}^0, \mathbf{P}^0) / \partial \mathbf{P}')^{-1} \partial \Psi(\boldsymbol{\theta}^0, \mathbf{P}^0) / \partial \boldsymbol{\theta}'$. Solving this expression into the Taylor approximation, we have the following approximation to the

counterfactual equilibrium:

$$\hat{\mathbf{P}}^* = \hat{\mathbf{P}}^0 + \left(I - \frac{\partial \Psi(\hat{\boldsymbol{\theta}}^0, \hat{\mathbf{P}}^0)}{\partial \mathbf{P}'} \right)^{-1} \frac{\partial \Psi(\hat{\boldsymbol{\theta}}^0, \hat{\mathbf{P}}^0)}{\partial \boldsymbol{\theta}'} (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}^0) \quad (36)$$

where $(\hat{\boldsymbol{\theta}}^0, \hat{\mathbf{P}}^0)$ represents our consistent estimator of $(\boldsymbol{\theta}^0, \mathbf{P}^0)$. It is clear that $\hat{\mathbf{P}}^*$ can be computed given the data and $\boldsymbol{\theta}^*$. Under our assumptions, $\hat{\mathbf{P}}^*$ is a consistent estimator of the linear approximation to $\boldsymbol{\pi}(\boldsymbol{\theta}^*)$.

As in any Taylor approximation, the order of magnitude of the error depends on the distance between the value of the structural parameters in the factual and counterfactual scenarios. Therefore, this approach can be inaccurate when the counterfactual experiment implies a large change in some of the parameters. For these cases, we can combine the Taylor approximation with iterations in the equilibrium mapping. Suppose that \mathbf{P}^* is a (Lyapunov) stable equilibrium. And suppose that the Taylor approximation $\hat{\mathbf{P}}^*$ belongs to the dominion of attraction of \mathbf{P}^* . Then, by iterating in the equilibrium mapping $\Psi(\boldsymbol{\theta}^*, \cdot)$ starting at $\hat{\mathbf{P}}^*$ we will obtain the counterfactual equilibrium \mathbf{P}^* . Note that this approach is substantially different to iterating in the equilibrium mapping $\Psi(\boldsymbol{\theta}^*, \cdot)$ starting with the equilibrium in the data $\hat{\mathbf{P}}^0$. This approach will return the counterfactual equilibrium \mathbf{P}^* *if and only if* $\hat{\mathbf{P}}^0$ belongs to the dominion of attraction of \mathbf{P}^* . This condition is stronger than the one establishing that the Taylor approximation $\hat{\mathbf{P}}^*$ belongs to the dominion of attraction of \mathbf{P}^* .⁴⁵

4 Concluding Comments

In this paper we have surveyed several challenges that we consider particularly important for applied work in estimation of dynamic demand and dynamic games. Our discussion of the two areas was mostly separate, reflecting to a large extent these two literatures have developed almost separately. In our view, an interesting area for future work is better integration and cross fertilization. We see several directions in which future work might proceed.

(a) *Further model simplification.* While we discussed several ways to simplify the computation and estimation of the models, the methods and computation are still quite complex and have limited applications. For demand for storable goods, Hendel and Nevo (2010) offer an alternative simple model that can be (easily) estimated using aggregate data. They make several not trivial

⁴⁵ Aguirregabiria (2009) provides examples where iterating in $\Psi(\boldsymbol{\theta}^*, \cdot)$ starting from $\hat{\mathbf{P}}^0$ returns an equilibrium that is not $\boldsymbol{\pi}(\boldsymbol{\theta}^*)$ (i.e., it is not of the same 'type' as the equilibrium \mathbf{P}^0) while the iterations starting at $\hat{\mathbf{P}}^*$ converge to the desired counterfactual equilibrium.

assumptions, the most important for simplifying the computation is that consumers can store at most for a known and pre-determined number of periods. With these assumption they show that the storable goods model is identified from aggregate data and does not require solving the dynamic programming problem. Thus the computational cost is of the same order as the computational cost of a static demand model. We think this sort of careful economic modeling as potentially useful both in the modeling of dynamic demand and in dynamic games.

(b) *Integration of dynamic demand and supply models.* Another promising avenue for future research is the combination of dynamic demand and dynamic supply. Most of the literature on estimation of dynamic games has concentrated on dynamics in supply but has ignored dynamics in demand and most of the literature on dynamic demand has not allowed for dynamic supply.⁴⁶ This is obviously an important limitation in the current state of the literature. As we move towards combining the two areas we believe that the modeling simplifications discussed in (a) would a particularly useful way to proceed. For example, with a simpler demand model Hendel and Nevo (2010) are able to add a supply side to the dynamic demand.

(c) *Identification.* Our discussion of identification of dynamic demand was informal, which reflects the state of the literature. A productive future avenue for research is to formally derive identification conditions, especially for estimation using aggregate data.

(d) *Estimation methods.* We see several directions for future work in estimation methods. First, estimation based on conditional choice probability has been successfully applied elsewhere but has rarely been used for estimating dynamic demand. In large part, because the first generation of these estimators could not allow for persistent unobserved heterogeneity. With the emergence of new estimators, that we presented in Section 3.3, we suspect that we will see more use of these methods in estimation of dynamic demand. Second, Bayesian estimation methods are particularly efficient, from a computational point of view, when multiple integration is cheaper than optimization. As we have seen for some of the estimators that we have presented above, optimization is particularly costly because it requires global search over a large dimensional space. This seems a good scenario to apply Bayesian estimation methods. Though Bayesian methods have been recently proposed for the estimation of single-agent dynamic structural models (Norets, 2009, Imai, Nair, and Ching, 2009), this type of methods have not been extended yet to deal with games with multiple equilibria.

(e) *Multiple equilibria.* We see a couple of directions for future work here. First, in the macroeconomic literature of Dynamic Stochastic General Equilibrium (DSGE) models, the standard

⁴⁶Goettler and Gordon (2009) is one of the few exceptions.

approach to deal with multiple equilibria is to linearize the equilibrium mapping. This seems a reasonable approach when we consider that multiple equilibria is not really an important feature of the model that is needed to explain the data, but more of a nuisance associated to the nonlinearity of the model. Though the idea of linearize the equilibrium mapping is related to two-step methods presented above, it is a different approach and it will be interesting to explore it. Alternatively, instead of treating multiple equilibrium as a nuisance we might consider whether it might actually aid in identification. Sweeting (2009) exploits multiple equilibria in a static entry game to gain identification ⁴⁷ This idea has not been explored for dynamic games.

(f) *Strategic uncertainty and beliefs out of equilibrium.* If the researcher believes that multiplicity of equilibria is a real issue in competition in actual markets, then firms may face significant strategic uncertainty in the sense that they may not know the strategies that other firms are playing. This strategic uncertainty can be particularly important in the context of oligopoly competition. Firms tend to be secretive about their own strategies, and it can be in their own interest to hide or even to misrepresent their own strategies. The identification and estimation of dynamic oligopoly games when firms' beliefs are out of equilibrium is an interesting area of further research.

(g) *Applications.* A final area where we hope to see more progress is in applications. Our discussion above focused on particular areas and applications, but we hope that the methods we discussed can and will be applied more widely.

⁴⁷See also de Paula and Tang (2010).

References

- [1] Aguirregabiria, V. (2009): "A Method for Implementing Counterfactual Experiments in Models with Multiple Equilibria," Working Papers tecipa-381, University of Toronto.
- [2] Aguirregabiria, V. (2010): "Another Look at the Identification of Dynamic Discrete Decision Processes: An Application to Retirement Behavior," *Journal of Business and Economic Statistics*, 28(2), 201-218.
- [3] Aguirregabiria, V. and C-Y. Ho (2009): "A dynamic oligopoly game of the US airline industry: Estimation and policy experiments," Manuscript. University of Toronto.
- [4] Aguirregabiria, V., and P. Mira (2002): "Swapping the nested fixed point algorithm: A class of estimators for Markov decision models," *Econometrica*, 70(4), 1519-43.
- [5] Aguirregabiria, V. and P. Mira (2007): "Sequential estimation of dynamic discrete games," *Econometrica* 75, 1-53.
- [6] Aguirregabiria, V. and P. Mira (2010): "Stability of Equilibria and Properties of Sequential Estimators: Problems and Solutions," Manuscript. University of Toronto.
- [7] Aguirregabiria, V., P. Mira, and H. Roman (2007): "An estimable dynamic model of entry, exit and growth in oligopoly retail markets," *American Economic Review* 97, 449-454.
- [8] Akerlof, G. (1970): "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84(3), 488-500.
- [9] Altonji, J. and L. Segal (1996): "Small-Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business & Economic Statistics*, 14(3), 353-366.
- [10] Arcidiacono, P. and R. Miller (2009): "CCP Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity," manuscript. Tepper School of Business. Carnegie Mellon University.
- [11] Bajari, P. & Benkard, C. L. (2005): "Demand estimation with heterogeneous consumers and unobserved product characteristics: A hedonic approach," *Journal of Political Economy*, 113, 1239-1276.
- [12] Bajari, P., L. Benkard and J. Levin (2007): "Estimating dynamic models of imperfect competition," *Econometrica* 75, 1331-1370.
- [13] Beresteanu, A. and P. Ellickson (2005): "The Dynamics of Retail Oligopoly," Manuscript. Department of Economics. Duke University.
- [14] Berry, S. (1992): "Estimation of a model of entry in the airline industry", *Econometrica*, 60(4), 889-917.
- [15] Berry, S. and P. Haile (2009): "Identification in Differentiated Products Markets Using Market Level Data," manuscript. Department of Economics. Yale University.
- [16] Berry, S., Levinsohn, J., & Pakes, A. (1995): "Automobile prices in market equilibrium," *Econometrica*, 63, 841-890.
- [17] Boizot, C., J-M Robin, and M. Visser (2001): "The Demand for Food Products: An Analysis of Interpurchase Times and Purchased Quantities," *Economic Journal*, 111, 391-419.
- [18] Bonnet, C. and P. Dubois (2010): "Inference on vertical contracts between manufacturers and retailers allowing for nonlinear pricing and resale price maintenance," *RAND Journal of Economics*, 41(1), 139-164.

- [19] Bresnahan, T. (1981): "Departures from marginal-cost pricing in the American automobile industry: Estimates for 1977–1978," *Journal of Econometrics*, 17(2), 201-227.
- [20] Bresnahan, T. (1987): "Competition and Collusion in the American Automobile Industry: The 1955 Price War," *Journal of Industrial Economics*, 35(4), 457-482.
- [21] Bresnahan, T., and P. Reiss (1990): "Entry in monopoly markets", *Review of Economic Studies*, 57, 531-553.
- [22] Bresnahan, T., and P. Reiss (1991a): "Empirical models of discrete games", *Journal of Econometrics* 48, 57-81.
- [23] Bresnahan, T., and P. Reiss (1991b): "Entry and competition in concentrated markets", *Journal of Political Economy*, 99, 977-1009.
- [24] Bresnahan, T., and P. Reiss (1993): "Measuring the Importance of Sunk Costs," *Annales d'Economie et de Statistique*, 34.
- [25] Carranza, J. (2006): "Demand for durable goods and the dynamics of quality," Unpublished manuscript, University of Wisconsin.
- [26] Collard-Wexler, A. (2006): "Demand Fluctuations in the Ready-Mix Concrete Industry," Manuscript. New York University.
- [27] De Paula, A. and X. Tang (2010): "Inference of Signs of Interaction Effects in Simultaneous Games with Incomplete Information," PIER Working Paper 10-021. UPenn.
- [28] Doraszelski, U., and Satterthwaite, M. (2010): "Computable Markov-Perfect Industry Dynamics," *RAND Journal of Economics*, 41(2), 215–243.
- [29] Dunne, T., S. Klimek, M. Roberts, and Y. Xu (2009): "Entry, Exit and the Determinants of Market Structure," manuscript. Department of Economics. Pennsylvania State University.
- [30] Eichenbaum, M., N. Jaimovich, and S. Rebelo (2008): "Reference Prices and Nominal Rigidities," NBER Working Paper 13829.
- [31] Einav, L. (2009): "Not All Rivals Look Alike: An Empirical Model for Discrete Games with Asymmetric Rivals," *Economic Inquiry*, 48(2), 369-390.
- [32] Einav, L., E. Leibtag, and A. Nevo (2009): "Recording Discrepancies in Nielsen Homescan Data: Are they Present and Do they Matter?," *Quantitative Marketing and Economics*, 8(2).
- [33] Erdem, T., S. Imai and M. P. Keane (2003): "Brand and Quantity Choice Dynamics under Price Uncertainty," *Quantitative Marketing and Economics* 1, 5-64.
- [34] Ericson, R., and A. Pakes (1995): "Markov-Perfect industry dynamics: A framework for empirical work", *Review of Economic Studies*, 62, 53-82.
- [35] Esteban, S. & Shum, M. (2007). "Durable goods oligopoly with secondary markets: The case of automobiles," *RAND Journal of Economics*, 38, 332-354.
- [36] Finger, S. (2008): "An Empirical Analysis of R&D Competition in the Chemicals Industry," manuscript. University of South Carolina.
- [37] General Motors Corporation (1939): "The dynamics of automobile demand," based upon papers presented at a joint meeting of the American Statistical Association and the Econometric Society in Detroit, Michigan on December 27, 1938.
- [38] Goettler, R. and B. Gordon (2009): "Does AMD spur Intel to innovate more?" manuscript. University of Chicago Booth School of Business.

- [39] Goldberg, P. (1995): "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," *Econometrica*, 63, 891-951.
- [40] Gowrisankaran, G. and M. Rysman (2009): "Dynamics of Consumer Demand for New Durable Goods," manuscript. University of Arizona.
- [41] Gowrisankaran, G., C. Lucarelli, P. Schmidt-Dengler, and R. Town (2010): "Government policy and the dynamics of market structure: Evidence from Critical Access Hospitals," manuscript. University of Arizona.
- [42] Griffith, R., A. Leicester, E. Leibtag, and A. Nevo (2009) "Consumer Shopping Behavior: How Much Do Consumers Save?," *Journal of Economics Perspectives*, 23(2), 99-120.
- [43] Hall, P. and X-H Zhou (2003): "Nonparametric Estimation of Component Distributions in a Multivariate Mixture," *Annals of Statistics*, 31(1), 201-224.
- [44] Hartmann, W. (2006): "Intertemporal effects of consumption and their implications for demand elasticity estimates," *Quantitative Marketing and Economics*, 4(4), 325-349.
- [45] Hashmi, A. and J. Van Biesebroeck (2010): "Market Structure and Innovation: A Dynamic Analysis of the Global Automobile Industry," NBER Working Paper 15959.
- [46] Hausman, J., G. Leonard, and J. Zona (1994): "Competitive Analysis with Differentiated Products," *Annales D'Economie et de Statistique*, 34.
- [47] Heckman, J. (1981): "The incidental parameters problem and the problem of initial conditions in estimating a discrete time - discrete data stochastic process," in C. Manski and D. McFadden (eds.), *Structural Analysis of Discrete Data with Econometric Applications*. MIT Press.
- [48] Hendel, I. and A. Nevo (2004): "Intertemporal Substitution and Storable Products," *Journal of the European Economic Association*, 2(2-3), pages 536-547.
- [49] Hendel, I., and A. Nevo (2006a): "Measuring the Implications of Sales and Consumer Inventory Behavior. *Econometrica* 74, 1637-1674.
- [50] Hendel, I. and A. Nevo (2006b): "Sales and Consumer Inventory," *RAND Journal of Economics*, 37(3), 543-561.
- [51] Hendel, I. and A. Nevo (2010): "A Simple Model of Demand Anticipation," manuscript. Department of Economics. Northwestern University.
- [52] Horowitz, J. (1998): "Bootstrap Methods for Covariance Structures," *Journal of Human Resources*, 33, 39-61.
- [53] Hotz, J., and R.A. Miller (1993): "Conditional choice probabilities and the estimation of dynamic models," *Review of Economic Studies*, 60, 497-529.
- [54] Hotz, J., R.A. Miller, S. Sanders, and J. Smith (1994): "A simulation estimator for dynamic models of discrete choice," *Review of Economic Studies*, 61, 265-89.
- [55] Imai, S., N. Jain, and A. Ching (2009): "Bayesian Estimation of Dynamic Discrete Choice Models," *Econometrica*, 77(6), 1865-1899.
- [56] Jovanovic, B. (1989): "Observable implications of models with multiple equilibria", *Econometrica*, 57, 1431-1437.
- [57] Kasahara, H. and K. Shumitsu (2008a): "Pseudo-likelihood Estimation and Bootstrap Inference for Structural Discrete Markov Decision Models," *Journal of Econometrics*, 146(1), 2008, 92-106.

- [58] Kasahara, H. and K. Shumitsu (2008b): "Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices," *Econometrica*, 77(1), 2009, 135-175.
- [59] Kasahara, H. and K. Shumitsu (2009): "Sequential Estimation of Structural Models with a Fixed Point Constraint," manuscript. University of Western Ontario.
- [60] Kehoe, P. and V. Midrigan (2008): "Temporary Price Changes and the Real Effects of Monetary Policy," NBER Working Papers 14392.
- [61] Kryukov, Y. (2010): "Dynamic R&D and the Effectiveness of Policy Intervention in the Pharmaceutical Industry," manuscript. Tepper School of Business. Carnegie Mellon University.
- [62] Lee, R. (2009): "Dynamic Demand Estimation in Platform and Two-Sided Markets," manuscript. Stern School of Business. New York University.
- [63] Lou, W., D. Prentice, and X. Yin (2008): "The Effects of Product Ageing on Demand: The Case of Digital Cameras," MPRA Paper 13407.
- [64] Macieira, J. (2007): "Extending the frontier: A structural model of investment and technological competition in the supercomputer industry," manuscript. Virginia Tech.
- [65] Magnac, T. and D. Thesmar (2002): "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70(2), 801-816.
- [66] Maskin, E. and Tirole, J. (1987): "A Theory of Dynamic Oligopoly, III: Cournot Competition," *European Economic Review*, 31, 947-968.
- [67] Maskin, E. and Tirole, J. (1988a): "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs," *Econometrica*, 56, 549-569.
- [68] Maskin, E. and Tirole, J. (1988b): "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles," *Econometrica*, 56, 571-599.
- [69] McFadden, D. (1974): "Conditional Logit Analysis of Qualitative Choice Behavior," in P. Zarembka (ed.), *Frontiers in Econometrics*, 105-142, Academic Press. New York.
- [70] Melnikov, O., 2000, Demand for Differentiated Durable Products: The Case of the U.S. Computer Printer Market. Manuscript. Department of Economics, Yale University.
- [71] Nakamura, E. and J. Steinsson (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, 123(4), 1415-1464.
- [72] Nevo, A. (2001): "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69(2), 307-342.
- [73] Nevo, A. and Rossi, F. (2008): "An approach for extending dynamic models to settings with multi-product firms," *Economics Letters*, 100(1), 49-52.
- [74] Norets, A. (2009): "Inference in Dynamic Discrete Choice Models with Serially Correlated Unobserved State Variables," *Econometrica*, 77, 1665-1682.
- [75] Pakes, A. (1994): "Dynamic structural models, problems and prospects," in C. Sims (ed.) *Advances in Econometrics*. Sixth World Congress, Cambridge University Press.
- [76] Pakes, A. and P. McGuire (1994): "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model," *RAND Journal of Economics*, 25(4), 555-589.
- [77] Pakes, A. and P. McGuire (2001) "Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'Curse of Dimensionality'," *Econometrica*, 69, 1261-1281.

- [78] Pakes, A., M. Ostrovsky, and S. Berry (2007): "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)," *Rand Journal of Economics* 38, 373-399.
- [79] Perrone, H. (2009): "Inventories, Unobservable Heterogeneity and Long Run Price Elasticities," manuscript. Departement of Economics and Business at Universitat Pompeu Fabra
- [80] Pesendorfer, M. (2002): "Retail Sales: A Study of Pricing Behavior in Supermarkets," *Journal of Business*, 75(1), pages 33-66.
- [81] Pesendorfer, M. and P. Schmidt-Dengler (2003): "Identification and Estimation of Dynamic Games," NBER Working Papers 9726, National Bureau of Economic Research.
- [82] Pesendorfer, M. and P. Schmidt-Dengler (2008): "Asymptotic Least Squares Estimators for Dynamic Games," *Review of Economic Studies*, 75(3), pages 901-928.
- [83] Pesendorfer, M. and P. Schmidt-Dengler (2010): "Sequential Estimation of Dynamic Discrete Games: A Comment," *Econometrica*, 78(2), pages 833-842.
- [84] Porter, R. (1983): "A Study of Cartel Stability: The Joint Executive Committee, 1880-1886," *Bell Journal of Economics*, 14(2), 301-314.
- [85] Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55(5), 999-1033.
- [86] Rust, J. (1994): "Estimation of dynamic structural models, problems and prospects: discrete decision processes", in C. Sims (ed.) *Advances in Econometrics. Sixth World Congress*, Cambridge University Press.
- [87] Ryan, S. (2006): "The Costs of Environmental Regulation in a Concentrated Industry," manuscript. Department of Economics. MIT.
- [88] Samuelson, P. (1947): "Foundations of Economic Analysis,"
- [89] Schiraldi, P. (2010): "Automobile Replacement: a Dynamic Structural Approach," manuscript. London School of Economics.
- [90] Shcherbakov, O. (2008): "Measuring switching costs in the television industry," manuscript. University of Arizona.
- [91] Smith, J. M. (1982), "Evolution and the Theory of Games," Cambridge University Press.
- [92] Snider, C. (2009): "Predatory incentives and predation policy: the American Airlines case," manuscript. Department of Economics. UCLA.
- [93] Su C., and K. Judd (2008): "Constrained Optimization Approaches to Estimation of Structural Models," Discussion Papers 1460, Northwestern University, Center for Mathematical Studies in Economics and Management Science.
- [94] Suzuki, J. (2010): "Land Use Regulation as a Barrier to Entry: Evidence from the Texas Lodging Industry," manuscript. Department of Economics. University of Toronto.
- [95] Sweeting, A. (2007): "Dynamic Product Repositioning in Differentiated Product Markets: The Case of Format Switching in the Commercial Radio Industry," NBER Working Paper #13522.
- [96] Sweeting, A. (2009): "The Strategic Timing of Radio Commercials: An Empirical Analysis Using Multiple Equilibria," *RAND Journal of Economics*, 40(4), 710-742.
- [97] Walrath, M. (2010): "Religion as an Industry: Estimating a Strategic Entry Model for Churches," manuscript. University of St. Thomas.