Uniqueness of Oblivious Equilibrium in Dynamic Discrete Games*

João Macieira†

February 7, 2010

Abstract

The recently introduced concept of oblivious equilibrium aims at approximating Markov-Perfect Nash Equilibria (MPNE) when its calculation is computationally prohibitive in Ericson-Pakes-style games with many players. This paper examines and extends the oblivious equilibrium concept to dynamic discrete choice games. We consider a set of assumptions commonly posed in the applied dynamic games literature and establish oblivious equilibrium properties. In contrast to Ericson-Pakes models, we find that there is a unique oblivious equilibrium in the dynamic game for any number of players under the frequently posed assumption of independence of state transitions across players. We demonstrate that the distance between this equilibrium and any MPNE converges in probability to zero as the number of game players goes to infinity. Unlike the original work on oblivious equilibrium, our convergence result requires neither a "light tail" condition nor the absence of aggregate state variables.

Keywords: Dynamic discrete games, oblivious equilibrium, equilibrium uniqueness, convergence in probability, approximation.

JEL Codes: C61, C62, C73, L13.

* *Email address:* joaom07@vt.edu.

†Department of Economics, Virginia Tech.

1 Introduction

The introduction of the Markov-Perfect Nash Equilibrium (MPNE) concept by Maskin and Tirole (1988a, 1988b) has resulted in a growing interest on examining dynamic games over the last two decades. To facilitate the mapping between theoretical results and concrete applications, several building-block frameworks have been introduced in the literature. For example, Ericson and Pakes (1995) introduced what has since become the building-block of modeling dynamic oligopolies. Aguirregabiria and Mira (2007), Pesendorfer and Schmidt-Dengler (2008) and Berry, Ostrovsky and Pakes (2007) developed estimable dynamic discrete-choice games by building on earlier advances in single-agent dynamic models (e.g. Rust 1987, Hotz and Miller 1993, Hotz et al. 1994). However, two main difficulties persist on examining equilibria in dynamic games. First, the calculation of MPNE is computationally unaffordable unless the dynamic game has a small number of players and a low-dimensional state space. Concrete applications where this problem is illustrated and discussed include Pakes and McGuire (1994, 2001), Gowrisankaran (1999) and Gowrisankaran and Town (1997). Second, a dynamic game typically has multiple Markov-Perfect equilibria, even when the researcher assumes that players restrict attention to anonymous, symmetric pure strategies (see Doraszelski and Satterthwaitte 2009 for an extensive discussion). These difficulties can be handled when the researcher is primarily interested on estimating the parameters of the dynamic game (e.g. Aguirregabiria and Mira 2007, Pesendorfer and Schmidt-Dengler 2007, Pakes, Ostrovsky and Berry 2008, Bajari, Benkard and Levin 2007). However, unless the policy experiment of interest to the researcher is spanned by estimation output such as player policy functions (e.g., Macieira 2009a, Macieira 2009b, Sweeting 2009), the researcher must compute the equilibria of the game.

The literature on dynamic equilibrium computation has recently proposed approaches that deal with some of the difficulties on calculating MPNE. One of those difficulties is the *curse of dimensionality*, defined as the fact that the state space grows exponentially with the number of game players. Doraszelski and Judd (2009) replace the usual discrete-time setup of dynamic games with a continuous-time one and illustrate how this difference can mitigate the curse of dimensionality. However, their dynamic investment game with entry and exit does not rule out multiple equilibria. Doraszelski, Besanko, Satherthwaite and Kryukov (2009) propose a method of calculating the set of possible equilibria in the context of a model of firm competition with organizational learning and forgetting. Nonetheless, even if this model is appropriate for the application of interest to the researcher, the problem of which equilibrium to compute for policy experiments still persists. Aguirregabiria and Ho (2009b) and Aguirregabiria (2009) deal with this difficulty in dynamic discrete games by calculating a new equilibrium using a linear Taylor series approximation around a given equilibrium. Even though this equilibrium selection method leads to substantial computational savings, it requires an initial equilibrium to start with. In addition, this approach is valid only when the counterfactual experiment involves small changes in game parameters. Abbring and Campbell (2009) propose a dynamic version of the model of

Bresnahan and Reiss (1993) with homogeneous firms that result in a unique symmetric MPNE when players restrict attention to what they name last-in-first-out (LIFO) strategies. However, the uniqueness result does not hold for more general settings of interest to applied researchers (see Abbring, Campbell and Yang 2010). Finally, Weintraub, Benkard and Van Roy (2008a) introduce the concept of oblivious equilibrium (OE) that aims at approximating MPNE in Ericson-Pakes models with many players. However, their approximation results apply only to settings without aggregate shocks affecting all player's payoffs (e.g., market demand state). In addition, OE is not shown to be unique. Weintraub, Benkard and Van Roy (2009a) develop an *extended oblivious equilibrium* concept to accommodate aggregate shocks in their model. However, the approximation theorem presented in Weintraub, Benkard and Van Roy (2008a) is not extended to this setup despite their simulation results on OE approximations to MPNE.

This paper aims at filling these gaps by examining the properties of oblivious equilibrium in dynamic discrete choice games. We build on the setup of Aguirregabiria and Mira (2007) and extend the concept of oblivious equilibrium (OE) pioneered by Weintraub, Benkard and Van Roy (2008) to that framework. We find that oblivious equilibrium is unique for the class of dynamic discrete games where player-specific state transitions do not depend on the actions of its rivals. This is one of the most frequent setups considered in concrete applications of dynamic games (e.g. Collard-Wexler 2009, Sweeting 2009, Ryan 2009, Beresteanu and Ellickson 2009, Aguirregabiria and Ho 2009a 2009b, among others[1]). We also find that oblivious equilibrium approximates MPNE in for games under a less restrictive set of assumptions than in Weintraub, Benkard and Van Roy (2008). In particular, we find that the distance between player value functions under MPNE and OE converges in probability to zero as the number of players goes to infinity. Our results apply to settings where there exist aggregate state variables affecting all firm payoffs. In addition, we provide sufficient conditions under which there exists OE for infinitely many players. We illustrate our approach with an example of firm competition market entry and exit borrowed from Aguirregabiria and Mira (2007).

The concept of Oblivious Equilibrium (OE) is proposed by Weintraub, Benkard and Van Roy (2008a, 2009a) as an alternative to MPNE for solving dynamic games. It assumes that players restrict attention only to their individual state variables (and aggregate states, if any) when making a decision. Under this solution concept, players are *oblivious* by considering the steady-state value of their rivals' states in lieu of their currently observed states. This involves the calculation of the steady-state of a vector process containing the number of players in each possible state, denoted $\mathbf{s}_t$. However, the direct calculation of the steady-state of a Markov chain $\{\mathbf{s}_t, t \in \mathbb{Z}_0^+\}$ with state transition Υ as the solution to the fixed point problem of $\mathbf{s} = \Upsilon\mathbf{s}$ is computationally unnafordable when the vector $\mathbf{s}_t$ can assume many possible values. Even though Weintraub, Benkard and Van Roy (2008a, 2009a) avoid this problem by exploiting the properties of their Ericson-Pakes-

[1] For a comprehensive survey on dynamic discrete games, see Aguirregabiria and Mira (2009) and Aguirregabiria and Nevo (2010) for a more comprehensive survey contemplating dynamic oligopolies in general.

like game, this difficulty applies to dynamic discrete-choice games with many players. We deal with difficulty by building on the Multinomial Markov Distribution framework of Wang and Yang (1995). By deriving the probability-generating function of sums of individual stationary Markov Chains, the framework of Wang and Yang (1995) allows us to redefine the fixed-point problem $\mathbf{s} = \Upsilon\mathbf{s}$ by replacing its right-hand side with a vector of partial derivatives of the probability-generating function. This results in a powerful simplification of steady-state calculations and eases our proofs on OE properties. Even though this paper focuses on the relationship between OE and MPNE, our approach to steady-state representation in dynamic discrete-choice games is applicable to other equilibrium concepts, provided they can be represented by choice probability profiles.

The main contribution of Weintraub, Benkard and Van Roy (2008a) is a theorem that shows that a sequence of OE satisfies what they define as the Asymptotic Markovian Equilibrium (AME) property. That is, as the number of players goes to infinity, the expected net present value gain of deviating from an OE to a best-response prescribed by a Markovian strategy converges to zero. This result is used to claim that OE approximates MPNE well when the game has many players, since by definition an MPNE implies no net present values from equilibrium deviations for any number of players. Weintraub, Benkard and Van Roy (2009a) provide simulation results that agree with their claim, yet there are two important concerns that matter for practical applications. First, the expectation measuring value function differences is taken with respect to the long-run distribution under an oblivious strategy. This expected difference metric allows player payoffs to differ considerably across states under Markovian and oblivious strategies while "canceling out" when the expectation is taken, resulting in a value close to zero. Second, the theorem of Weintraub, Benkard and Van Roy (2008a) requires a "light tail" condition on the flow payoff of players. This condition aims at ruling out equilibria where the shares of game payoffs are concentrated in a few players. In this paper we deal with these issues in the context of dynamic discrete-choice games. We consider convergence in probability as our metric for approximation between net present values under different strategies. This type of convergence concept considers distance in absolute value and enjoys important properties under function continuity. This allows a direct mapping between net present values and choice probabilities, ensuring convergence at every possible state of the game. In addition, we do not require a "light tail" condition since OE is proven to be unique in our framework.

Most of dynamic models following the framework of Ericson and Pakes (1995) involve a state variable observable to all players that affects each agents payoff (e.g. number of potential consumers, state of demand). However, the theorem of Weintraub, Benkard and Van Roy (2008a) does not extend to this important setting. Although Weintraub, Benkard and Van Roy (2009a) provide simulation evidence indicating that the presence of aggregate states does not compromise the approximation power of OE, it is convenient to have a concrete proof agreeing with the computational evidence. Our paper also contributes in this dimension for the case of dynamic discrete-choice games. Our approximation result based on an asymptotic theorem also applies to models with aggregate states. In addition, OE is also unique in that situation. This latter result is grounded on the fact that OE in dynamic discrete games with independent transitions across players essentially replicate Markov decision models similar to the one Aguirregabiria and Mira (2002). Our uniqueness proof therefore builds on steps in the spirit of Aguirregabiria and Mira (2002) with proper modifications to the presence of a steady-state vector for $\mathbf{s}_t$.

This paper is organized as follows. Section 2 describes the dynamic discrete-choice model and the two equilibrium concepts considered for solving the game. Section 3 presents the properties of oblivious equilibrium and applicable asymptotic results describing approximation to MPNE for games with many players. Section 4 outlines our main conclusions and discusses possible extensions to our framework. All proofs not displayed in the main sections of the paper are presented in the Appendix.

2 The Dynamic Model

Our framework builds on the dynamic discrete choice games of Aguirregabiria and Mira (2007) but with some differences. We consider flow payoff functions which do not depend on the actions of rivals. That is, rival actions only interfere with future player payoffs through its stochastic impact on the state vector evolution. This is the case most frequently considered in the dynamic games literature and it is a basic feature of all work involving the oblivious equilibrium concept (e.g., Weintraub, Benkard and van Roy 2006, 2008a, 2008b, 2009a, 2009b, Xu 2008). In addition, we adapt the notation of Aguirregabiria and Mira (2007) to the one used by Weintraub, Benkard and Van Roy (2008a) to facilitate the definition of oblivious equilibrium in dynamic discrete games.

2.1 Model and Assumptions

Time is assumed discrete and indexed by $t = 0, 1, ..., \infty$. In each period, there are N players indexed by $i = 1, ..., N$ who simultaneously decide which action to take out of a set of discrete and finite options, denoted A_i. We represent player i's action by $a_{it} \in A_i$ and we let $\mathbf{a}_t = (a_{1t}, ..., a_{Nt})$ represent the vector of player choices at period t. In addition, we assume that all payoff-relevant features for each player can be encoded into a state vector. At the beginning of each period, a player is characterized by two state vectors: a vector of variables that are common-knowledge to all players, denoted $x_{it} \in X_i$, and a vector $\varepsilon_{it} \in \Sigma$ that is private information of player i. In what follows, $\boldsymbol{\varepsilon}_t = (\varepsilon_{i,t}, ..., \varepsilon_{N,t})$ represents a vector containing private information of all N players at period t.We assume that the set X_i has a discrete and finite support. In addition, we define the *industry state* $\mathbf{s}_t$ as the vector over player-specific observed states that specifies how many players out of N are in each state $x_{it} \in X_i$.

In what follows, it is convenient to define what Weintraub, Benkard and Van Roy (2008, 2009) call the *state of competitors* of player i. This state vector, denoted $\mathbf{s}_{-i,t}$, is defined by $\mathbf{s}_{-i,t}(x) = \mathbf{s}_t(x) - 1$ if $x_{it} = x$ and $\mathbf{s}_{-i,t}(x) = \mathbf{s}_t(x)$ otherwise. Finally, we assume that there is an aggregate shock state $\omega_t \in \Omega$ which is common-knowledge to all players. This aggregate shocks affects each players' payoff function, denoted $\tilde{\pi}_i(a_{it}, \omega_t, x_{it}, \mathbf{s}_{-i,t}, \varepsilon_{it})$. Here we differ from Weintraub, Benkard and Van Roy (2008, 2009) by considering a more general payoff function. Our payoff function depends on player action a_{it} and private information ε_{it}, while Weintraub, Benkard and Van Roy (2008, 2009) restrict attention to flow payoffs that depend only on player choice and observed states.

We assume that each firm decides its action to maximize expected discounted sum of payoffs

$$E\left[\sum_{\tau=t}^{\infty}\beta^{\tau-t}\tilde{\pi}_i(a_{it}, \omega_t, x_{it}, \mathbf{s}_{-i,t}, \varepsilon_{it})\middle|\, \omega_t, x_{it}, \mathbf{s}_{-i,t}, \varepsilon_{it}\right] \tag{2.1}$$

where $\beta \in [0, 1)$ is a discount factor common to all players. The conditional expectation is taken over current and future values of state variables and private shocks considering the state transition probability $p(\omega_{t+1}, x_{i,t+1}, \mathbf{s}_{-i,t+1}, \boldsymbol{\varepsilon}_{t+1}|\omega_t, x_{it}, , \mathbf{s}_{-i,t}, \boldsymbol{\varepsilon}_t, \mathbf{a}_t)$. As in Aguirregabiria and Mira (2007), we pose the following assumptions:

Assumption 1 *(Additive separability) The flow payoff function of each player is linear in private information for each $a_{it} \in A_i$. That is, $\tilde{\pi}_i(a_{it}, \omega_t, x_{it}, \mathbf{s}_{-i,t}, \varepsilon_{it}) = \pi_i(a_{it}, \omega_t, x_{it}, \mathbf{s}_{-i,t}) + \varepsilon_{it}(a_{it})$, where $\varepsilon_{it}(a_{it})$ denotes the private information value that player i obtains if he chooses action a_{it} at time t.*

Assumption 2 *(Conditional Independence) The state transition function satisfies the condition*

$$\begin{aligned} p(\omega_{t+1}, x_{i,t+1}, \mathbf{s}_{-i,t+1}, \boldsymbol{\varepsilon}_{t+1}|\omega_t, x_{it}, , \mathbf{s}_{-i,t}, \boldsymbol{\varepsilon}_t, \mathbf{a}_t) \quad = \quad & p_\varepsilon(\boldsymbol{\varepsilon}_{t+1}|\omega_{t+1}, x_{i,t+1}, \mathbf{s}_{-i,t+1}) \times \\ & p_x(x_{i,t+1}|\omega_t, x_{it}, , \mathbf{s}_{-i,t}, a_{it}) \times \\ & p_{\mathbf{s}_{-i}}(\mathbf{s}_{-i,t+1}|\omega_t, \mathbf{a}_t, x_{it}, , \mathbf{s}_{-i,t}) \times \\ & p_\omega(\omega_{t+1}|\omega_t, x_{it}, , \mathbf{s}_{-i,t}) \end{aligned} \tag{2.2}$$

Assumption 3 *(Independent Private Values) The private information vector $\boldsymbol{\varepsilon}_t$ is independently distributed across players given observed states $(\omega_{t+1}, x_{i,t+1}, \mathbf{s}_{-i,t+1})$. Its factors as*

$$p_\varepsilon(\boldsymbol{\varepsilon}_{t+1}|\omega_{t+1}, x_{i,t+1}, \mathbf{s}_{-i,t+1}) = \prod_{i=1}^{N} g_i(\varepsilon_{i,t+1}|\omega_{t+1}, x_{i,t+1}, \mathbf{s}_{-i,t+1}) \tag{2.3}$$

where, for each player $i = 1, ..., N$, $g_i(.|.)$ is an absolutely continuous density function measurable with respect to the Lebesgue measure.

The first assumption is posed for technical convenience and is shared by nearly all the dynamic discrete choice models. The second assumption imposes several restrictions on the probability transition function. First, the private information shocks are independent over time. Even though this is an important limitation, it is shared by most theoretical and empirical work on dynamic discrete games for yield model tractability. Second, the evolution of each player's individual future state value, $x_{i,t+1}$, is a stochastic function of observed states and own player's action, a_{it}. Even though other dynamic discrete game frameworks also consider the dependence on other player actions (e.g., Aguirregabiria and Mira 2007, Pesendorfer and Schmidt-Dengler 2008), most practical applications of dynamic games consider dependence only on own player's action (e.g. Collard-Wexler 2009, Ryan 2009, among others). The third assumption shares similarities with most dynamic discrete game frameworks and is also motivated by tractabilitiy. It is not as strong as in other mainstream works on dynamic discrete games, in which the conditional density $g_i(\varepsilon_{i,t+1}|\omega_{t+1}, x_{i,t+1}, \mathbf{s}_{-i,t+1})$ is replaced by its unconditional analog $g_i(\varepsilon_{i,t+1})$ (e.g. Aguirregabiria and Mira 2007, Pesendorfer and Schmidt-Dengler 2007).

In order to provide some intuition and fix an interpretation for our framework, we provide an example which consists of a modified version of Example 1 in Aguirregabiria and Mira (2007, page 6) that satisfies the above-listed assumptions:

Example 1 *(Market Entry and Exit) There are N players in the game consisting of firms deciding on operating or not in a given market. At each moment t, each player i can be in two situations represented by the sate variable $x_{i,t}$; either they are "potential entrants" (denoted $x_{i,t} = 0$) or they are "incumbents" (denoted $x_{i,t} = 1$). Thus the set of possible individual states is $X_i \equiv \{0, 1\}$. The industry state $\mathbf{s}_t$ is a vector which displays how many players are in each state in X_i. In our setup, if the market has n incumbents at time t, then $\mathbf{s}_t = (N - n, n)$. The competitor's state vector, $\mathbf{s}_{-i,t}$, is $(N - n - 1, n)$ if player i is potential entrant and it is $(N - n, n - 1)$ otherwise. We also assume that there is an aggregate shock variable ω_t representing market demand at time t. An example of the set of possible values for ω_t is $\Omega = \{a, b\}$ where $a, b \in \mathbb{R}^+$ and $a < b$. Intuitively, this would correspond to a model of "Low" and "High" demand states quantified by the values a and b, respectively.*

We assume that at period t each player must decide whether or not to become an incumbent for the next period. We denote player i's choice on operating in the market at period $t+1$ by $a_{it} = 1$ and the decision of not operating by $a_{it} = 0$. In our model, this implies an action set $A_i \equiv \{0, 1\}$ for every player i. The vector of private information shocks for chosing each alternative is $\varepsilon_{it} = (\varepsilon_{it}(0), \varepsilon_{it}(1))$. An example that satisfies our conditions on the distribution of private values is assumption that $\varepsilon_{it}(0)$ and $\varepsilon_{it}(1)$ are identically and independently distributed according to a Type I Extreme Value distribution. The flow payoff function is assumed to take the form

$$\tilde{\pi}_i(a_{it}, \omega_t, x_{it}, \mathbf{s}_{-i,t}, \varepsilon_{it}) = x_{it}\left(\theta_M \frac{\omega_t}{(2+\mathbf{s}_{-i,t}(2))^2} - \theta_{FC} + \theta_X(1-a_{it})\right) - \theta_E(1-x_{it})a_{it} + \varepsilon_{it}(a_{it})$$

where $\mathbf{s}_{-i,t}(2)$ denotes the second cell of $\mathbf{s}_{-i,t}$ (i.e., the number of player i's rival incumbents). In this formulation, a potential entrant bears no market profits at time t, and must pay an entry cost θ_E in case she decides to operate at period $t+1$. An incumbent receives market variable profits $\theta_M \omega_t / (2+\mathbf{s}_{-i,t}(2))^2$ net of fixed costs θ_{FC}. The expression for variable market profits could be obtained from two well-known models: a Cournot equilibrium with same marginal costs across firms, and the Salop (1979) circle city model of firm competition in prices. The incumbent also incurs into an extra payoff of θ_X in case she decides not to operate the market in period $t+1$. The parameter θ_X could be either positive or negative, depending on whether a market leave decision yields a scrap value or exit costs, respectively.

One example of observed state transitions satisfying our assumptions would be the case where $\{\omega_t, t \geq 0\}$ is an ergodic Markov chain and where $x_{i,t+1} = a_{it}$. That is, the decision of operating fully determines the player's state in the following period. The implied transition probabilities for every possible value of $x_{i,t+1}$ and $\mathbf{s}_{-i,t+1}$ given observed states and actions are defined by

$$p_x(x_{i,t+1}|\omega_t, x_{it}, , \mathbf{s}_{-i,t}, a_{it}) = \begin{cases} 1 \text{ if } x_{i,t+1} = a_{it} \\ 0 \text{ otherwise} \end{cases}$$

$$p_{\mathbf{s}_{-i}}(\mathbf{s}_{-i,t+1}|\omega_t, \mathbf{a}_t, x_{it}, , \mathbf{s}_{-i,t}) = \begin{cases} 1 \text{ if } \mathbf{s}_{-i,t+1} = \left(N - \mathbf{1}^T\mathbf{a}_t - (1-a_{it}), \mathbf{1}^T\mathbf{a} - a_{it}\right) \\ 0 \text{ otherwise} \end{cases}$$

where $\mathbf{1}^T$ is a $1 \times N$ matrix of ones.

2.2 Equilibria

In this section we provide a convenient representation of each agent's problem and the equilibrium concepts considered to solve the dynamic game. We assume that all players restrict attentione to stationary Markovian strategies. That is, player actions are deterministic functions of observed states and private information. We omit time subindexes and denote next period observed states, private information and competitor's state by ω', x_i', ε_i' and $\mathbf{s}_{-i}'$, respectively. We denote player i's choice function by $\sigma_i(\omega, x_i, \mathbf{s}_{-i}, \varepsilon_i)$, which are formally a mapping $\sigma_i : \Omega \times X_i \times \mathbf{S}_{-i} \times \Sigma_i \to A_i$. In what follows, it is convenient to define the conditional choice probabilities associated with the profile of strategy functions $\boldsymbol{\sigma} = \{\sigma_i(\omega, x_i, \mathbf{s}_{-i}, \varepsilon_i)\}_{i=1}^N$ as

$$P_i^{\sigma}(a_i|\omega, x_i, \mathbf{s}_{-i}) \equiv \Pr(\sigma_i(\omega, x_i, \mathbf{s}_{-i}, \varepsilon_i) = a_i|\omega, x_i, \mathbf{s}_{-i}) = \int \mathbf{1}(\sigma_i(\omega, x_i, \mathbf{s}_{-i}, \varepsilon_i) = a_i) g_i(d\varepsilon_i|\omega, x_i, \mathbf{s}_{-i}) \quad (2.4)$$

As in Aguirregabiria and Mira (2007), the set of probabilities $\{P_i^{\sigma}(a_i|\omega, x_i, \mathbf{s}_{-i}) : a_i \in A_i\}_{i=1}^{N}$ represents expected behavior from player i from the perspective of the remaining players when player i abides by his strategy in the set $\boldsymbol{\sigma}$. We use this fact to provide a recursive representation of the player's problem 2.1. Under a strategy profile $\boldsymbol{\sigma}$, player i's problem can be represented by the Bellman equation

$$\begin{aligned} \hat{V}_i^{\sigma}(\omega, x_i, \mathbf{s}_{-i}, \varepsilon_i) \quad = \quad & \max_{a_i \in A_i} \{\pi_i(a_i, \omega, x_i, \mathbf{s}_{-i}) + \varepsilon_i(a_i) \\ & + \beta \sum_{\omega' \in \Omega_i} \sum_{x_i' \in X_i} \sum_{\mathbf{s}_{-i} \in S_{-i}} \left[\int \hat{V}_i^{\sigma}(\omega', x_i', \mathbf{s}_{-i}', \varepsilon_i') g_i(d\varepsilon_i'|\omega', x_i', \mathbf{s}_{-i}') \right] p_x(x_i'|\omega, x_i, , \mathbf{s}_{-i}, a_i) \times \\ & p_\omega(\omega'|\omega, x_i, \mathbf{s}_{-i}) \times p_{\mathbf{s}_{-i}}^{\sigma}(\mathbf{s}_{-i}'|\omega, x_i, \mathbf{s}_{-i}, a_i) \Big\} \end{aligned} \tag{2.5}$$

where $p_{\mathbf{s}_{-i}}^{\sigma}(\mathbf{s}_{-i}'|\omega, x_i, \mathbf{s}_{-i}, a_i)$ is the transition probability of player i's competitors state given that they follow the strategy profile $\boldsymbol{\sigma}$ and player i chooses a_i. Taking conditional expectations of 2.5 with respect to ε_i allows us to relate value functions with choice probabilities. The resulting expression is the integrated Bellman Equation

$$\begin{aligned} V_i^{\sigma}(\omega, x_i, \mathbf{s}_{-i}) \quad = \quad & \sum_{a_i \in A_i} P_i^{\sigma}(a_i|\omega, x_i, \mathbf{s}_{-i}) \{[\pi_i(a_i, \omega, x_i, \mathbf{s}_{-i}) + E(\varepsilon_i(a_i)|a_i, \omega, x_i, \mathbf{s}_{-i})] \\ & + \beta \sum_{\omega' \in \Omega_i} \sum_{x_i' \in X_i} \sum_{\mathbf{s}_{-i} \in S_{-i}} V_i^{\sigma}(\omega', x_i', \mathbf{s}_{-i}') p_x(x_i'|\omega, x_i, , \mathbf{s}_{-i}, a_i) \times \\ & p_\omega(\omega'|\omega, x_i, \mathbf{s}_{-i}) \times p_{\mathbf{s}_{-i}}^{\sigma}(\mathbf{s}_{-i}'|\omega, x_i, \mathbf{s}_{-i}, a_i) \Big\} \end{aligned} \tag{2.6}$$

2.2.1 Markov-Perfect Nash Equilibrium

The Markov-Perfect Nash Equilibrium (MPNE) pioneered by Maskin and Tirole (1988a, 1988b) has been the most frequently used solution concept for dynamic games. The following definition of MPNE in pure strategies builds on the work of Aguirregabiria and Mira (2007) with different notation adapted to our framework:

Definition 1 *A profile of Markovian strategies $\boldsymbol{\sigma}^*$ is a Markov-Perfect Nash Equilibrium if for every player $i = 1, .., N$ and any $(\omega, x_i, \mathbf{s}_{-i}, \varepsilon_i) \in \Omega \times X_i \times S_{-i} \times \Sigma_i$, $\sigma_i^*(\omega, x_i, \mathbf{s}_{-i}, \varepsilon_i)$ is a solution for the Bellman Equation 2.5 when all rivals of player i choose according to their strategy function described in $\boldsymbol{\sigma}^*$.*

As in Aguirregabiria and Mira (2007), the MPNE just defined can be represented in probability space by following the framework of Milgrom and Weber (1985). Working in probability space is particularly convenient for several reasons. First, it facilitates equilibrium existence proofs by transforming the strategy set into a bounded continuous set. Second, the right-hand side of the integrated Bellman Equation 2.6 consists of a contraction mapping operator that is continuous in choice probabilities. This property is os central importance for the proofs outlined below. Finally, integrating out private information from 2.5 allows applied researchers to map the structural model to concrete empirical settings (see for example, Aguirregabiria and Mira 2007, Pesendorfer and Schmidt-Dengler 2007, Beresteanu and Ellickson 2009).

Let $\mathbf{P} = (P_1, ..., P_N)$ denote a profile of choice probabilitiy functions associated with a Markovian strategy profile $\boldsymbol{\sigma}$ and let $(V_1^P, ..., V_N^P)$ denote the players' value fiunction vectors associated with $\mathbf{P}$. Then we can rewrite 2.6 with slight changes in notation as

$$\begin{aligned} V_i^P(\omega, x_i, \mathbf{s}_{-i}) \quad = \quad & \sum_{a_i \in A_i} P_i(a_i|\omega, x_i, \mathbf{s}_{-i}) \left\{ \left[\pi_i(a_i, \omega, x_i, \mathbf{s}_{-i}) + e_i^P\,(a_i, \omega, x_i, \mathbf{s}_{-i}) \right] \right. \\ & + \beta \sum_{\omega' \in \Omega_i} \sum_{x_i' \in X_i} \sum_{\mathbf{s}_{-i} \in S_{-i}} V_i^P(\omega', x_i', \mathbf{s}_{-i}') p_x(x_i'|\omega, x_i, , \mathbf{s}_{-i}, a_i) \times \\ & \left. p_\omega(\omega'|\omega, x_i, \mathbf{s}_{-i}) \times p_{\mathbf{s}_{-i}}^P(\mathbf{s}_{-i}'|\omega, x_i, \mathbf{s}_{-i}, a_i) \right\} \end{aligned} \tag{2.7}$$

where $e_i^P\,(a_i, \omega, x_i, \mathbf{s}_{-i})$ corresponds to the expected value of private information given observed states and the action chosen by player i equals $\sigma_i(\omega, x_i, \mathbf{s}_{-i}, \varepsilon_i)$. The fact that $e_i^P\,(a_i, \omega, x_i, \mathbf{s}_{-i})$ is a function only of player i's choice probability vector P_i and the conditional private information density $g_i(.|.)$ is established by following the same steps as in Aguirregabiria and Mira (2007, pp 9-10).

We can explicitly write the fixed point problem that defines equilibrium probabilities $\mathbf{P}^*$ by taking the following two steps. First, we stack the state-specific elements of 2.7 for each player i and define a continuous matrix operator on $\mathbf{P}$, denoted $\Gamma_i(\mathbf{P})$, as is the unique solution for V_i^P given a profile $\mathbf{P}$. Second, we combine the Definition 1 with 2.4 to represent MPNE in probability space as the solution to the fixed point problem involving $\Gamma(\mathbf{P})$, where $\Gamma(\mathbf{P}) = (\Gamma_1(\mathbf{P}), ..., \Gamma_N(\mathbf{P}))$. The system representation of 2.7 following the stacking procedure described for the first step is

$$V_i^P = \sum_{a_i \in A_i} P_i(a_i) * \left[\pi_i(a_i) + e_i^P\,(a_i) + \beta F_i^P(a_i) V_i^P \right] \tag{2.8}$$

where $P_i(a_i)$, $\pi_i(a_i)$ and $e_i^P\,(a_i)$ are column vectors of dimension $R = |\Omega \times X_i \times \mathbf{S}_{-i}|$ stacking state and action-specific elements of choice probability, payoff and conditional expected private information, respectively. The $R \times R$ matrix $F_i^P(a_i)$ contains the transition probabilities conditional on player i chosing action a_i, while $*$ represents the Hadamard element-by-element matrix product. As in Aguirregabiria and Mira

(2007), the unique solution for the player-specific value function vector V_i^P defined by the linear system in 2.8 given P is

$$\varphi_i(\mathbf{P}) = \left(I - \beta F_i^P\right)^{-1} \left\{ \sum_{a_i \in A_i} P_i(a_i) * \left[\pi_i(a_i) + e_i^P\,(a_i)\right] \right\} \tag{2.9}$$

where $F_i^P = \sum_{a_i \in A_i} P_i(a_i) * F_i^P(a_i)$.

Using the definition of choice probability in 2.4 to the particular case of a strategy that solves player i's problem 2.5, an operator mapping the choice probabilities profile $\mathbf{P}$ onto the the set of player-specific choice probabilities $\Psi_i(\mathbf{P}) \equiv \left(\Psi_i(a_1|\omega, x_i, \mathbf{s}_{-i}; \mathbf{P}), ..., \Psi_i(a_{|A_i|}|\omega, x_i, \mathbf{s}_{-i}; \mathbf{P})\right)$ is defined as

$$\begin{aligned}
\Psi_i(a_i|\omega, x_i, \mathbf{s}_{-i}; \mathbf{P}) \quad &= \quad \int \mathbf{1}(a_i = \sigma_i(\omega, x_i, \mathbf{s}_{-i}, \varepsilon_i)) g_i(d\varepsilon_i|\omega, x_i, \mathbf{s}_{-i}) && (2.10)\\
&\int \mathbf{1} \left\{ a_i = \arg\max_{a_i \in A_i} \{\pi_i(a_i, \omega, x_i, \mathbf{s}_{-i}) + \varepsilon_i(a_i) \right. && (2.11)\\
&+\beta \sum_{\omega' \in \Omega_i} \sum_{x_i' \in X_i} \sum_{\mathbf{s}_{-i} \in S_{-i}} \varphi_i(\omega', x_i', \mathbf{s}'_{-i}; \mathbf{P}) p_x(x_i'|\omega, x_i, , \mathbf{s}_{-i}, a_i) \times \\
&\left. p_\omega(\omega'|\omega, x_i, , \mathbf{s}_{-i}) \times p^P_{\mathbf{s}_{-i}}(\mathbf{s}'_{-i}|\omega, x_i, \mathbf{s}_{-i}, a_i) \Big) \right\} g_i(d\varepsilon_i|\omega, x_i, \mathbf{s}_{-i})
\end{aligned}$$

where $\varphi_i(\omega', x_i', \mathbf{s}'_{-i}; \mathbf{P})$ regards the value function cell of $\varphi_i(\mathbf{P})$ associated with the combination of states $(\omega', x_i', \mathbf{s}'_{-i})$ computed using choice probability profile $\mathbf{P}$.

An MPNE in choice probabilities space consists of any profile $\mathbf{P}^*$ that is a solution the fixed point problem $\mathbf{P} = \Psi(\mathbf{P})$, where $\Psi(\mathbf{P}) = (\Psi_1(\mathbf{P}), ..., \Psi_N(\mathbf{P}))$. Existence of $\mathbf{P}^*$ such that $\mathbf{P}^* = \Psi(\mathbf{P}^*)$ follows directly from Brower's fixed point theorem. However, the set of solutions for this fixed-point problem is rarely a singleton; Doraszelski and Satterthwaite (2009) discuss and illustrate multiplicity of MPNE for closely-related, general dynamic games. In an attempt to deal with multiplicity of equilibria while allowing good approximations to MPNE, we define the concept of oblivious equilibrium for the model outlined above in the next section.

2.2.2 Oblivious Equilibrium

So far we have outlined a model similar to Aguirregabiria and Mira (2007) with some modifications in notation and assumptions. We now define the concept of oblivious equilibrium pioneered by Weintraub, Benkard and Van Roy (2008a) in the context of the model outlined above. As their original work was aimed at examining models similar to the one of Ericson and Pakes (1995), we modify some aspects of equilibrium definition while keeping its intuition.

The concept of oblivious equilibrium was motivated by the fact that computation of MPNE is not computationally affordable when a dynamic game has more than a few players. This problem is due to two facts. First, increasing the number of players increases the number of dynamic programming problems required to obtain at least one (out of possibly many) equilibria. Second, the dimensionality of the state space grows exponentially with the number of players. This *curse of dimensionality* imposes substantial computational and memory storage problems, which cannot be handled with computational resources commonly available to researchers. To deal with this problem, Weintraub, Benkard and Van Roy (2008) proposed the concept of *oblivious equilibrium.* In this equilibrium concept, each agent solves its dynamic problem considering its own state, x_i, while treating the competitors' state, $\mathbf{s}_{-i}$, as a constant. This constant is the steady-state value of the industry state vector computed under the assumption that all players choose optimal strategies in the fashion just described. This approach transforms the multi-agent problem into a single-agent problem by treating other players' states as constant, leading to substantial computational gains. For an Ericson-Pakes game with no aggregate shocks, Weintraub, Benkard and Van Roy (2008) demonstrate that the expected difference between value functions computed under MPNE and oblivious strategies converges to zero as the number of game players goes to infinity[2]. Their proof does not extend to a game with aggregate shocks affecting all player payoffs, although simulation results indicate that oblivious equilibrium may still provide good approximations to MPNE in that situation (for a discussion Weintraub, Benkard and Van Roy 2009).

The computational challenges on calculating MPNEs in models with many players also apply to dynamic discrete games similar to the ones of Aguirregabiria and Mira (2007) and Pesendorfer and Schmidt-Dengler (2007). This fact motivates a definition of oblivious equilibrium in the spirit of Weintraub, Benkard and Van Roy (2009). Let $\bar{\mathbf{s}}^P(\omega)$ denote the steady-state industry vector given a profile of player choice probabilities $\mathbf{P}$ and the current value of the aggregate shock ω. In what follows, it is convenient to follow the probability-generating function framework of Wang and Yang (1995) to concisely define the steady-state industry vector. The probability-generating function of a discrete random variable Y is defined as $G(z) = E(z^Y)$. It is by definition a power series whose coefficients are positive since $E(z^Y) = \sum_{y=0}^{\infty} \Pr(Y = y)z^y$. The computation of probabilities and expectations using $G(z)$ is straightforward since $\Pr(Y = k) = G^{(k)}(0)/k!$ and $E(Y) = \lim_{z\to 1} G^{(1)}(z)$, where $G^{(l)}(z)$ denotes the l^{th} derivative of G with respect to z.

By definition, we have $\bar{\mathbf{s}}^P(\omega) \equiv N\boldsymbol{\eta}^P(\omega)$, where $\boldsymbol{\eta}^P(\omega)$ is a column vector containing the fraction of players that, in the steady-state, will be in each state considered in X_i. Under our assumption that the value of player i's individual state at period $t+1$ is not stochastically dependent on its rivals actions', the transition matrix for $x_{i,t+1}$ given the player's choice probabilities P_i and observed states is

[2] To be exact, the limit is taking by letting one of the game parameters - the market size - go to infinity. The result is established using the property that the equilibrium number of players is increasing in market size.

$$\begin{array}{ccc} & & \begin{array}{ccc} & x_{i,t+1} & \\ 1 & \cdots & J \end{array} \\ & \begin{array}{c} 1 \\ \vdots \\ J \end{array} & \begin{bmatrix} p_{1,1}(\omega, \mathbf{s}_{-i}) & \cdots & p_{1,J}(\omega, \mathbf{s}_{-i}) \\ \vdots & & \vdots \\ p_{J,1}(\omega, \mathbf{s}_{-i}) & \cdots & p_{J,J}(\omega, \mathbf{s}_{-i}) \end{bmatrix} \end{array} \quad (2.12)$$

where $p_{l,j}(\omega, \mathbf{s}_{-i}) = \sum_{a_i \in A_i} p_x(x_i' = j|\omega, x_i = l, \mathbf{s}_{-i}, a_i)P_i(a_i|\omega, x_i = l, \mathbf{s}_{-i})$. Using Theorem 1 in Wang and Yang (1995), the probability-generating function for the industry state vector $\mathbf{s}$ given that the game has N players is

$$\hat{G}_N(\mathbf{z}; \boldsymbol{\eta}^P(\omega), \omega, \mathbf{s}) = \left[\boldsymbol{\eta}^P(\omega)^T * (1, z_2, .., z_J)\right] Q^{N-1}(\mathbf{z}; \omega, \mathbf{s})\mathbf{1} \quad (2.13)$$

where the cells of $Q(\mathbf{z}; \omega, \mathbf{s})$ are defined by $[p_{l,j}(\omega, \mathbf{s}_{-i})z_j]_{l,j \in \{1,...,J\}}$, , $Q^k(.)$ represents the k^{th} power of matrix $Q(\mathbf{z}; \omega, \mathbf{s})$, and $\mathbf{1}$ is a $J \times 1$ vector of ones. Without loss of generality, we consider the probability generating function $G_N(\mathbf{z}; \bar{\mathbf{s}}^P(\omega), \omega, \mathbf{s})$ by merely multiplying and dividing $\boldsymbol{\eta}^P(\omega)$ by N. Using the expectation formula for probability-generating functions, the steady-state industry vector $\bar{\mathbf{s}}^P(\omega)$ is formally defined as the solution to the fixed-point problem on the vector $\mathbf{s}$

$$\mathbf{s} = T(\bar{\mathbf{s}}^P(\omega), \omega, \mathbf{s}) \quad (2.14)$$

where

$$T(\bar{\mathbf{s}}^P(\omega), \omega, \mathbf{s}) = \begin{bmatrix} \lim_{z_1 \to 1} \dfrac{\partial}{\partial z_1} G_N((z_1, 0, ..., 0); \bar{\mathbf{s}}^P(\omega), \omega, \mathbf{s}) \\ \vdots \\ \lim_{z_J \to 1} \dfrac{\partial}{\partial z_J} G_N((0, ..., 0, z_J); \bar{\mathbf{s}}^P(\omega), \omega, \mathbf{s}) \end{bmatrix} \quad (2.15)$$

The column vector $T(\bar{\mathbf{s}}^P(\omega), \omega, \mathbf{s})$ implicitly assumes that there exists at least one steady-state vector $\bar{\mathbf{s}}^P(\omega)$ by including it as one of its arguments. As our current focus is on the definition of oblivious equilibrium for our dynamic discrete game, we postpone the discussion about existence and uniqueness of a solution for 2.14 to the next section. Our exposition proceeds by assuming the existence of a solution $\bar{\mathbf{s}}^P(\omega)$ for industry steady-state given current aggregate shock value ω and a choice probability profile $\mathbf{P}$.

As in Weintraub, Benkard and Van Roy (2009), we assume that agents take $\bar{\mathbf{s}}^P(\omega)$ as exogenously given in their dynamic problem by neglecting the impact of their individual choices on the computation of industry steady-state. That is, instead of considering the current value of competitor state vector $\mathbf{s}_{-i}$, agents take the industry steady-state function $\bar{\mathbf{s}}^P(\omega)$ as given and replace $\mathbf{s}_{-i}$ with the implied steady-state vector for competitor state, denoted $\bar{\mathbf{s}}^P_{-i}(\omega)$. This is the reason for naming this equilibrium as *oblivious*, in the sense that each agent disregards the true information on its rivals and replaces it with a steady-state constant vector. The latter is a particularly convenient value to consider, since it allows the agent to simplify considerably its optimization problem by restricting attention to its own state x_i and the aggregate shock ω as the relevant state variables to solve for her optimal policy. Consequently, the sequential representation of each player's objective function is now

$$E\left[\sum_{\tau=t}^{\infty}\beta^{\tau-t}\tilde{\pi}_i(a_{it},\omega_t,x_{it},\bar{\mathbf{s}}^P_{-i,t}(\omega_t),\varepsilon_{it})\middle|\,\omega_t,x_{it},\varepsilon_{it}\right] \tag{2.16}$$

where $\mathbf{s}_{-i,t}$ has been removed from the information set since ω_t suffices to compute the vector $\bar{\mathbf{s}}^P_{-i,t}(\omega_t)$. To solve this dynamic problemy, each player no longer considers the set $\Theta \equiv \Omega \times X_i \times \mathbf{S}_{-i} \times \Sigma_i$ to form a policy function; instead, she restricts attention to the set $\tilde{\Theta} \equiv \Omega \times X_i \times \Sigma_i$ and considers the mapping of oblivious policies $\tilde{\sigma} : \tilde{\Theta} \to A_i$. We define the *oblivious choice probabilities* is a fashion similar to 2.4 with the necessary modifications:

$$\tilde{P}_i^{\tilde{\sigma}}(a_i|\omega,x_i) \equiv \Pr(\tilde{\sigma}_i(\omega,x_i,\varepsilon_i)=a_i|\omega,x_i) = \int \mathbf{1}(\tilde{\sigma}_i(\omega,x_i,\varepsilon_i)=a_i)g_i(d\varepsilon_i|\omega,x_i,\bar{\mathbf{s}}^{\tilde{P}}_{-i}(\omega)) \tag{2.17}$$

where $\tilde{\boldsymbol{\sigma}} = \{\tilde{\sigma}_i(\omega,x_i,\varepsilon_i)\}_{i=1}^N$ represents a profile of oblivious strategy functions and $\tilde{P}$ is the choice probability profile associated with $\tilde{\boldsymbol{\sigma}}$. We assume that each player acknowledges that its rivals restrict attention to oblivious strategies. Consistent with this assumption, the set of oblivious choice probabilities $\{\tilde{P}_i^{\tilde{\sigma}}(a_i|\omega,x_i) : a_i \in A_i\}_{i=1}^N$ represents expected behavior from player i from the perspective of the remaining players when player i abides by the oblivious strategy profile $\tilde{\boldsymbol{\sigma}}$.

The recursive representation of the player i's problem 2.16 given oblivious strategy profile $\tilde{\boldsymbol{\sigma}}$ is represented by the Oblivious Bellman equation

$$\begin{aligned}\hat{V}_i^{\tilde{\sigma}}(\omega,x_i,\bar{\mathbf{s}}^{\tilde{P}}_{-i}(\omega),\varepsilon_i) \quad = \quad & \max_{a_i\in A_i}\Big\{\pi_i(a_i,\omega,x_i,\bar{\mathbf{s}}^{\tilde{P}}_{-i}(\omega))+\varepsilon_i(a_i) \\ & +\beta\sum_{\omega'\in\Omega_i}\sum_{x_i'\in X_i}\left[\int \hat{V}_i^{\tilde{\sigma}}(\omega',x_i',\bar{\mathbf{s}}^{\tilde{P}}_{-i}(\omega'),\varepsilon_i')g_i(d\varepsilon_i'|\omega',x_i',,\bar{\mathbf{s}}^{\tilde{P}}_{-i}(\omega'))\right]\times \\ & p_x(x_i'|\omega,x_i,\bar{\mathbf{s}}^{\tilde{P}}_{-i}(\omega),a_i)\times p_\omega(\omega'|\omega,x_i,,\bar{\mathbf{s}}^{\tilde{P}}_{-i}(\omega))\Big\}\end{aligned} \tag{2.18}$$

We now provide a definition of oblivious equilibrium for the dynamic discrete games:

Definition 2 *A profile of oblivious strategies* $\tilde{\boldsymbol{\sigma}}^*$ *is an Oblivious Equilibrium if for every player* $i = 1, .., N$ *and any* $(\omega, x_i, \varepsilon_i) \in \tilde{\Theta}$, *the following conditions hold:*

(i) $\sigma_i^*(\omega, x_i, \varepsilon_i)$ *is a solution for the Bellman Equation 2.18 when all rivals of player* i *choose according to their strategy function described in* $\tilde{\boldsymbol{\sigma}}^*$;

(ii) $\overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega)$ *is a solution to the steady-state problem 2.14 for the choice probability profile* $\tilde{\mathbf{P}}^*$ *implied by the oblivious strategy profile* $\tilde{\boldsymbol{\sigma}}^*$.

There are two key differences between 2.5 and 2.18. The first is that the competitor state vector is merely a function of ω in 2.18, not a time-varying vector governed by a transition as in 2.5. Secondly, the discounted future payoff term in 2.18 does not contain a summation over all possible values of $\mathbf{s}_{-i} \in S_{-i}$. This follows from the assumption that, from all players' perspective, the temporal evolution of $\mathbf{s}_{-i}$ is now governed by changes in the aggregate shock ω through the vector function $\overline{\mathbf{s}}_{-i}^{P}(\omega)$. These two properties result in two important simplifications of player i's problem. First, 2.18 is now a single-agent problem, since each player i takes $\overline{\mathbf{s}}_{-i}^{P}(\omega)$ as exogenously given and so her relevant state space consists only of its individual states - x_i and ε_i and an exogenously-evolving aggregate shock ω. Second, there is a considerable reduction on state space dimension as a result of restricting attention to the competitors' steady-state vector rather considering the transition matrix for $\mathbf{s}_{-i}$.

In a slight abuse of notation, we define the oblivious value function by removing $\overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega)$ from its argument set. Taking conditional expectations of 2.18 with respect to ε_i, the integrated Oblivious Bellman Equation becomes

$$\begin{aligned} \tilde{V}_i^{\tilde{\sigma}}(\omega, x_i) &= \sum_{a_i \in A_i} \tilde{P}_i^{\tilde{\sigma}}(a_i|\omega, x_i) \Big\{ \Big[\pi_i(a_i, \omega, x_i, \overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega)) + e_i^{\tilde{P}} \left(a_i, \omega, x_i, \overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega) \right) \Big] \\ &\quad + \beta \sum_{\omega' \in \Omega_i} \sum_{x_i' \in X_i} \tilde{V}_i^{\tilde{\sigma}}(\omega', x_i')) p_x(x_i'|\omega, x_i, , \overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega), a_i) \times p_\omega(\omega'|\omega, x_i, , \overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega)) \Bigg\} \end{aligned} \tag{2.19}$$

where $e_i^{\tilde{P}} \left(a_i, \omega, x_i, \overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega) \right)$ now corresponds to the expected value of private information given observed states (ω, x_i), the competitor's steady-state function $\overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega)$ and the action chosen by player i being the oblivious strategy $\tilde{\sigma}_i(\omega, x_i, \varepsilon_i)$. As in the case of Markovian strategies, $e_i^{\tilde{P}} \left(a_i, \omega, x_i, \overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega) \right)$ can be shown to be a function only of player i's choice probability vector $\tilde{P}_i$ and the conditional private information density $g_i(.|.)$ by mimicking the proof of Aguirregabiria and Mira (2007, pp 9-10).

We now map the concept of oblivious equilibrium into probability space. We follow the same approach as in the definition of MPNE in probability space by stacking state-specific elements of **??** for each player i. The resulting system denining oblivious integrated Bellman equations given an oblivious profile $\tilde{\mathbf{P}}$ is

$$\tilde{V}_i^{\tilde{P}} = \sum_{a_i \in A_i} \tilde{P}_i(a_i) * \left[\tilde{\pi}_i(a_i) + e_i^{\tilde{P}}(a_i) + \beta \tilde{f}_i^{\tilde{P}}(a_i)\tilde{V}_i^{\tilde{P}}\right] \tag{2.20}$$

where $\tilde{P}_i(a_i)$, $\tilde{\pi}_i(a_i)$ and $e_i^{\tilde{P}}(a_i)$ are column vectors of dimension $G = |\Omega \times X_i|$ stacking state and action-specific elements of choice probability, oblivious payoff and conditional expected private information, respectively. Player i considers the probability matrix $\tilde{f}_i^{\tilde{P}}(a_i)$, which represents state transitions of (ω, x_i) given that action a_i is chosen and competitors' state vector is evaluated at $\overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega)$. Note that $G < R$ since player i now restricts attention to (ω, x_i) and uses this information and the probability profile to compute $\overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega)$. In addition to alleviating the curse of dimensionality through state space reduction, this simplification transforms the multi-agent dynamic game into a single-agent Markov decision process similar to the one of Aguirregabiria and Mira (2002). This fact will play a key role on showing uniqueness of oblivious equilibrium. The value function operator as a function of the oblivious choice probability is

$$\tilde{\varphi}_i(\tilde{\mathbf{P}}) = \left(I - \beta \tilde{F}_i^{\tilde{P}}\right)^{-1} \left\{ \sum_{a_i \in A_i} \tilde{P}_i(a_i) * \left[\tilde{\pi}_i(a_i) + e_i^{\tilde{P}}(a_i)\right] \right\} \tag{2.21}$$

where $\tilde{F}_i^{\tilde{P}} = \sum_{a_i \in A_i} \tilde{P}_i(a_i) * \tilde{f}_i^{\tilde{P}}(a_i)$.

We complete our definition of oblivious equilibrium (OE) in probability space by representing it as a fixed point in $\tilde{\mathbf{P}}$. As in the probabilistic representation of MPNE, we use the definition of choice probability in 2.4 to form an operator mapping the space of probability profiles into the space of individual choice probabilities. This operator consists of the probability that a given action solves the Oblivious Bellman equation 2.18. Since players restrict attention to oblivious strategies to solve this equation, each component of the operator is defined by

$$\begin{aligned}
\tilde{\Psi}_i(a_i|\omega, x_i; \tilde{\mathbf{P}}) &= \int \mathbf{1}(a_i = \tilde{\sigma}_i(\omega, x_i, \varepsilon_i)) g_i(d\varepsilon_i|\omega, x_i, \overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega)) && (2.22)\\
&= \int \mathbf{1}(a_i = \arg\max_{a_i \in A_i} \Big\{ \pi_i(a_i, \omega, x_i, \overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega)) + \varepsilon_i(a_i) && (2.23)\\
&\quad + \beta \sum_{\omega' \in \Omega_i} \sum_{x_i' \in X_i} \tilde{\varphi}_i(\omega', x_i'; \tilde{\mathbf{P}}) p_x(x_i'|\omega, x_i, , \overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega), a_i) \times \\
&\quad p_\omega(\omega'|\omega, x_i, \overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega)) \Big) g_i(d\varepsilon_i|\omega, x_i, \overline{\mathbf{s}}_{-i}^{\tilde{P}}(\omega))
\end{aligned}$$

where $\tilde{\varphi}_i(\omega', x_i'; \tilde{\mathbf{P}})$ regards the value function cell of $\tilde{\varphi}_i(\tilde{\mathbf{P}})$ associated with the combination of states (ω', x_i') computed given choice probability profile $\tilde{\mathbf{P}}$.

Letting $\tilde{\Psi}_i(\tilde{\mathbf{P}}) \equiv \left(\tilde{\Psi}_i(a_1|\omega, x_i; \tilde{\mathbf{P}}), ..., \tilde{\Psi}_i(a_{|A_i|}|\omega, x_i; \tilde{\mathbf{P}})\right)$, an OE representation in choice probabilities space consists of any profile $\tilde{\mathbf{P}}^*$ that is a solution the fixed point problem $\tilde{\mathbf{P}} = \tilde{\Psi}(\tilde{\mathbf{P}})$, where $\tilde{\Psi}(\tilde{\mathbf{P}}) = (\tilde{\Psi}_1(\tilde{\mathbf{P}}), ..., \tilde{\Psi}_N(\tilde{\mathbf{P}}))$. We use this representation of OE to examine the properties of OE and its relationship with MPNE in the next section.

3 Results on Oblivious Equilibrium

In this section we examine two topics that are at the core of our contributions. First, we discuss the existence and uniqueness of oblivious equilibrium. Second, we examine the approximation properties of oblivious equilibrium to MPNE when the number of players goes to infinity.

The steady-state industry vector is a key ingredient in the definition of oblivious equilibrium whose existence we have taken as given. For this reason, we start by discussing the existence and uniqueness of a steady-state solution for 2.14 for any oblivious choice probability profile $\tilde{\mathbf{P}}$.

Lemma 1 *For any oblivious choice probability profile $\tilde{\mathbf{P}}$, there exists an unique steady-state vector of competitor's states.*

Proof. Note that under an oblivious profile $\tilde{\mathbf{P}}$ the steady-state operator $T(\bar{\mathbf{s}}^P(\omega), \omega, \mathbf{s})$ defined in 2.15 simplifies to $T(\bar{\mathbf{s}}^P(\omega), \omega)$ since the industry state $\mathbf{s}$ is not an argument of oblivious strategies. From assumptions 2.2 and 2.3, it follows that the transition matrix for the industry state $\mathbf{s}$ implicit to 2.13 defines an irreducible Markov Chain for any oblivious profile $\tilde{\mathbf{P}}$ since every state of the industry vector $\mathbf{s}$ reaches other states with strictly positive probability. In addition, the transition matrix for $\mathbf{s}$ implicit to 2.13 is aperiodic by construction since all its diagonal elements are strictly positive[3]. Uniqueness of industry state vector follows from the well-known result that an irreducible and aperiodic Markov Chain has an unique steady-state distribution. ■

A direct implication of this lemma is that for any oblivious profile $\tilde{\mathbf{P}}$ there is only one steady-state vector of player i's competitors state, denoted $\bar{\mathbf{s}}^{\tilde{\mathbf{P}}}_{-i}(\omega)$. This simplifies considerably the analysis of OE existence and uniqueness, as it implies that every player considers at most one constant vector in lieu of $\mathbf{s}_{-i}$.

3.1 Existence and Uniqueness

As in the above discussion about MPNE, we establish the existence of oblivious equilibrium by direct application of Brower's fixed point theorem to the system $\tilde{\mathbf{P}} = \tilde{\Psi}(\tilde{\mathbf{P}})$. A more complicated yet interesting

[3] A Markov chain $\{Z_t, t \geq 0\}$ is said to be aperiodic if the greatest common divisor of the sequence $\{n : \Pr(Z_n = i|Z_0 = i)\}$ is 1 for every state i.

issue is the uniqueness of oblivious equilibrium. The following theorem addresses this issue:

Theorem 1 *There is an unique oblivious equilibrium for the dynamic discrete game.*

We discuss the intuition behind this theorem's proof, which we present in the Appendix. First, the fact that there exists an unique steady-state for any oblivious choice probability profile $\tilde{\mathbf{P}}$ implies that the function $\bar{\mathbf{s}}^{\tilde{P}}_{-i}(\omega)$ is unique regardless of whether $\tilde{\mathbf{P}}$ is an equilibrium or not. Since each agent takes $\bar{\mathbf{s}}^{\tilde{P}}_{-i}(\omega)$ as exogenously given, it is treated as a vector function of the aggregate shock ω. Thus, each agent i is solving a Markov decision process similar to the one of Aguirregabiria and Mira (2002) since $\varphi_i(\tilde{\mathbf{P}})$ only depends on $\tilde{P}_i$, not the whole profile $\tilde{\mathbf{P}}$. Thus, the oblivious profile $\tilde{\mathbf{P}}^*$ consisting of the collection of individual choice probabilities solving each player's Markov decision process given $\bar{\mathbf{s}}^{\tilde{\mathbf{P}}^*}_{-i}(\omega)$ must be an equilibrium. Aguirregabiria and Mira (2002) show that the solution to their Markov decision process is unique by exploiting the uniqueness of fixed point for the value function operator. Uniqueness of the oblivious equilibrium $\tilde{\mathbf{P}}^*$ is proven is established in a similar fashion by using the uniqueness of fixed points for both the value function and steady-state operators.

3.2 Asymptotic Results

The introduction of the oblivious equilibrium concept in the literature by Benkard, Weintraub and Van Roy (2008) is motivated by its ability to approximate MPNE when the number of players in the game is large. Here we address this issue in the context of the dynamic discrete game described above. We demonstrate that the distance between the value function vectors of MPNE and OE converges in probability to zero as the number of players goes to infinity.

There are three key differences between our results and the ones of Benkard, Weintraub and Van Roy (2008). First, the metric considered in the convergence result of Benkard, Weintraub and Van Roy (2008) is the expected difference between value functions under MPNE and OE, where expectations are taken considering the steady-state distribution under OE. This metric does not rule out the possibility that value functions differ considerably under the two equilibrium concepts while having the same expected value. The metric considered in our work is convergence in probability, which by definition rules out that caveat. Second, Benkard, Weintraub and Van Roy (2008) consider a game without aggregate shocks ω. This condition is required to validate their convergence proof. Even though simulation results presented in Benkard, Weintraub and Van Roy (2009) indicate that the presence of aggregate shocks does not compromise the quality of OE approximations to MPNE, no concrete proof is presented. In our famework we explicitly allow for aggregate states without compromising our covergence proofs. Finally, Benkard, Weintraub and Van Roy (2008) require what they name as a "light-tail" condition to proof their asymptotic result. In the Ericson-Pakes model of investment in quality, the "light-tail" condition consists of a restriction on the players' profit function aimed at ruling out equilibria where a small portion of players holds a high fraction of total market profits. This would be the case of markets where a few firms have most of market share (e.g., software industry). In these situations, OE could result in a very poor approximation to MPNE despite the large number of players since the distance between steady-state and current industry vector is large due to the presence of dominant firms. This type of assumption is not required for proving that OE approximates well MPNE when the number of players is large. There are two intuitive reasons for this fact. First, we have shown that OE is unique for any number of game players. So imposing a "light-tail" condition will not lead to any improvement of the approximation properties of OE. Second, the observed state space is discrete and finite. Consequently, the fraction of players who are at a given state converges in probability to a constant by invoking a suitable law of large numbers. This fact avoids the need of any "light-tail" condition to ensure convergence under a given metric. However, we need an auxiliary assumption on the flow payoff function to ensure that the concept of convergence as $N \to \infty$ is well-defined. The following assumption ensures that property:

Assumption 4 *The flow payoff function $\pi_i(a_{it}, \omega_t, x_{it}, \mathbf{s}_{-i,t})$ can be written as a function continuous in $\boldsymbol{\eta}_t$, denoted $\pi_{i,N}(a_{it}, \omega_t, x_{it}, \boldsymbol{\eta}_t)$, where $\boldsymbol{\eta}_t$ is the relative frequency of each state of X_i for a game with N players.*

This assumption is fairly weak and applies to most applied work on dynamic discrete games, including the market entry and exit example described above. We are now ready to establish our main asymptotic result:

Theorem 2 *The distance between player value functions using Markovian and Oblivious strategies converges to zero in probability as the number of game players goes to infinity.*

Here we discuss the proof main steps and intuition while outlining the complete version in the Appendix. First, we invoke the uniform law of large numbers of Vapnik and Chervonenkis (1971) to claim that the relative frequency vector for X_i, (i.e. $\boldsymbol{\eta}_t$), converges under any Markovian or Oblivious strategy profiles to the same vector as $N \to \infty$. That is, the limiting steady-state distribution of players under MPNE and OE is the same. Second, we note that strategy functions defined over the Markovian state space $\Theta \equiv \Omega \times X_i \times \mathbf{S}_{-i} \times \Sigma_i$ do not improve upon strategies defined over the oblivious space set $\tilde{\Theta} \equiv \Omega \times X_i \times \Sigma_i$ for solving the dynamic programing problem of an oblivious player described in 2.16. This is a feature of oblivious strategies that our framework shares with the one of Weintraub, Benkard and Van Roy (2008). We use the triangular inequality to bound the difference between value functions by the sum of two components which are shown to converge to zero in probability. The convergence in probability of the latter to zero follows from our assumption on flow payoff and the uniform convergence in probability of $\boldsymbol{\eta}_t$.

Corollary 1 *The distance between an MPNE and the unique OE converges in probability to zero as $N \to \infty$.*

Proof. This results follows from the convergence of state distributions $\boldsymbol{\eta}_t$ established using the uniform law of large numbers of Vapnik and Chervonenkis (1971) and the continuity of the choice probability operators $\Lambda_i(.)$ and $\tilde{\Lambda}_i(.)$. By the uniform law of large numbers of the uniform law of large numbers of Vapnik and Chervonenkis (1971) the relative frequency vector $\boldsymbol{\eta}_t(\omega)$ converges in probability uniformily to a constant vector $\boldsymbol{\eta}(\omega)$ as $N \to \infty$.Using the triangular inequality, we have

$$\left\|\Lambda_i(V_{i,N}^{\mathbf{P}}) - \tilde{\Lambda}_i(\tilde{V}_{i,N}^{\tilde{\mathbf{P}}})\right\|_\infty \leq \left\|\Lambda_i(V_{i,N}^{\mathbf{P}}) - \Lambda_i(V_{i,N}^{\mathbf{P}})\right\|_\infty + \left\|\Lambda_i(V_{i,N}^{\mathbf{P}}) - \tilde{\Lambda}_i(\tilde{V}_{i,N}^{\tilde{\mathbf{P}}})\right\|_\infty \tag{3.1}$$

where $V_i^{\mathbf{P}}$ is player i's value function under MPNE profile $\mathbf{P}$ when the competitor's state is calculated using the limit $\boldsymbol{\eta}(\omega)$. It follows from Assumption 4 that value function vectors are continuous in $\boldsymbol{\eta}_t$. In addition, $\Lambda_i(.)$ and $\tilde{\Lambda}_i(.)$ are continuous functionals mapping value function vectors into choice probabilities. Therefore, it follows from the Mann-Wald Continuous Mapping theorem that $\lim_{N\to\infty} \Pr\left\{\left\|\Lambda_i(V_{i,N}^{\mathbf{P}}) - \Lambda_i(V_{i,N}^{\mathbf{P}})\right\|_\infty > \varepsilon\right\} = 0$ and that $\lim_{N\to\infty} \Pr\left\{\left\|\Lambda_i(V_{i,N}^{\mathbf{P}}) - \tilde{\Lambda}_i(\tilde{V}_{i,N}^{\tilde{\mathbf{P}}})\right\|_\infty > \varepsilon\right\} = 0, \forall\varepsilon > 0$. Since each parcel of the right-hand side of 3.1 converges in probability, the proof is complete. ■

The intuition behind this result relies on the fact that, as $N \to \infty$, $\boldsymbol{\eta}_t(\omega)$ converges to a constant vector. This means that $\bar{\mathbf{s}}^{\tilde{P}}(\omega)$ becomes approximately constant when the game has many players. This implies that, for each player, their rival state vector is provides approximately the same information as the steady-state computed using OE. Since the key difference between MPNE and OE is that the former considers competitors' state vector $\mathbf{S}_{-i}$ in the strategy set, it comes at no surprise that, that the choice probabilities converge. This argument is confirmed using the continuity of value function and choice probability operators and the Continuous Mapping theorem.

4 Conclusion

This paper builds upon and extends the oblivious equilibrium concept of Benkard, Weintraub and Van Roy (2008a, 2009a) to Aguirregabiria-Mira-style dynamic discrete games. We find that under assumptions commonly posed in both empirical and theoretical literature on dynamic games, the resulting oblivious equilibrium is unique and can approximate Markov-Prefect equilibria in games with many players. Our approximation metric is stronger than the one considered by Weintraub and Van Roy (2008). Unlike this latter pioneer work, our proofs do not require absence of aggregate shocks as state variables. These results allow researchers to compute equilibria for policy experiments without facing equilibrium multiplicity issues frequently present at dynamic games.

The calculation of the unique oblivious equilibrium is computationally affordable even for large state space dimensions for two reasons. First, the oblivious equilibrium computation essentially requires the researcher to compute an unique fixed point in probability space representing a single-agent problem. The computational costs of this type of fixed-point have been minimized courtesy of improved computing resources and advances in equilibrium calculation in probability spaces (e.g. Aguirregabiria and Mira 2002). Second, the oblivious equilibrium also involves computing a fixed point for industry steady-state vector, whose dimension is equal to the player state space X_i. The probability-generating function approach of Wang and Yang (1995) for sums of Markov chains avoids the cumbersome and possibly prohibitive task of deriving the transition matrix for the industry state vector.

The fact that the distance between value functions of MPNE and OE shrinks as $N \to \infty$ by no means guarantees that OE is well-defined for infinitely many players. Weintraub, Benkard and Van Roy (2009b) examine this issue in the context of their Ericson-Pakes-like model. In our context, the existence and properties of OE when $N = \infty$ depends on the concrete assumptions on game primitives. In particular, it depends on which (and how many) of the cells of the industry steady-state vector $\bar{\mathbf{s}}^{\tilde{P}}(\omega)$ converge to a constant as $N \to \infty$. In this latter case, the asymptotic distribution results of Wang and Yang (1995) can be used to show that the OE in dynamic discrete games has well-defined equilibrium choice probabilities when $N = \infty$, where the cells of $\bar{\mathbf{s}}^{\tilde{P}}(\omega)$ converging to constants as $N \to \infty$ are given by the Compound Poisson distribution parameters.Since the fulfillment of those conditions requires a more extensive examination of the model primitives at hand, further work on this direction is necessary and is left for future research.

Appendix: Proofs of Theorems

Proof of Theorem 1. First, note that the oblivious value function operator of each player i, $\varphi_i(\tilde{\mathbf{P}})$, depends only of individual choice probabilities $\tilde{P}_i$, not the whole profile $\tilde{\mathbf{P}}$. This follows from the fact that each player i takes the competitor's state $\bar{\mathbf{s}}^{\tilde{\mathbf{P}}}_{-i}(\omega)$ as an exogenously given function of ω and the components $\tilde{F}_i^{\tilde{P}}$, $\tilde{\pi}_i(a_i)$ and $e_i^{\tilde{P}}(a_i)$ feeding in $\varphi_i(\tilde{\mathbf{P}})$ depend only on $\tilde{P}_i$.Thus, without loss of generality, consider the individual value function operator $\varphi_i(\tilde{P}_i)$. Second, note that the right-hand side of 2.20 defines an operator on the space of oblivious value functions, which we denote by $\tilde{\Gamma}_i(\tilde{V}_i^{\tilde{P}})$. As in Aguirregabiria and Mira (2002), this operator can be shown to be a contraction mapping by invoking the results of Rust, Traub and Wozniakowski (2002). Thus, for every player i there is an unique oblivious value function vector $\tilde{V}_i^{\tilde{P}}$ solving the fixed point problem $\tilde{V}_i^{\tilde{P}} = \tilde{\Gamma}_i(\tilde{V}_i^{\tilde{P}})$. Third, recall that under an oblivious profile $\tilde{\mathbf{P}}$ the steady-state operator $T(\bar{\mathbf{s}}^P(\omega), \omega, \mathbf{s})$ defined in 2.15 simplifies to $T(\bar{\mathbf{s}}^P(\omega), \omega)$. The uniqueness of industry steady-state vector under oblivious strategy profiles implies that there is an unique $\bar{\mathbf{s}}^P(\omega)$ solving the fixed-point problem $\bar{\mathbf{s}}^P(\omega) = T(\bar{\mathbf{s}}^P(\omega), \omega)$. We use these facts to demonstrate that if the oblivious profile $\tilde{\mathbf{P}}^* = \left(\tilde{P}_1^*, ..., \tilde{P}_N^*\right)$ is a fixed point of the operator $\tilde{\mathbf{\Psi}}(\tilde{\mathbf{P}})$, then for every player i, $\varphi_i(\tilde{P}_i^*)$ is a fixed point of $\tilde{\Gamma}_i(.)$. Our proof strategy follows an approach similar to the one of Aguirregabiria and Mira (2002) with necessary modifications for our setup. From the definition of oblivious equilibrium, the fixed point condition $\tilde{\mathbf{P}}^* = \tilde{\mathbf{\Psi}}(\tilde{\mathbf{P}}^*)$ implies that for every player i we have $\tilde{P}_i^* = \tilde{\Psi}_i(\tilde{P}_i^*)$. Defining the operator $\tilde{\Lambda}_i(\tilde{V}_i^{\tilde{P}})$ as the vector stacking all choice probabilities defined in the right-hand side of 2.23, we have that $\tilde{P}_i^* = \tilde{\Psi}_i(\tilde{P}_i^*) = \tilde{\Lambda}_i(\varphi_i(\tilde{P}_i^*))$ for every player i. Using the definition of $\tilde{\Gamma}_i(\tilde{V}_i^{\tilde{P}})$ for the particular case $\tilde{V}_i^{\tilde{P}} = \varphi_i(\tilde{P}_i^*)$, we have

$$\begin{aligned}
\tilde{\Gamma}_i(\varphi_i(\tilde{P}_i^*)) &= \sum_{a_i \in A_i} \tilde{\Lambda}_i(a_i|\varphi_i(\tilde{P}_i^*)) * \left[\tilde{\pi}_i(a_i, \bar{\mathbf{s}}_{-i}^{\tilde{\mathbf{\Psi}}(\tilde{\mathbf{P}})}) + e_i^{\tilde{P}}\left(a_i, \tilde{\Lambda}_i(\varphi_i(\tilde{P}_i^*))\right) + \beta \tilde{f}_i^{\tilde{P}}(a_i)\varphi_i(\tilde{P}_i^*)\right] \\
&= \sum_{a_i \in A_i} \tilde{P}_i^*(a_i) * \left[\tilde{\pi}_i(a_i, \bar{\mathbf{s}}_{-i}^{\tilde{\mathbf{P}}^*}) + e_i^{\tilde{P}}\left(a_i, \tilde{P}_i^*\right) + \beta \tilde{f}_i^{\tilde{P}}(a_i)\varphi_i(\tilde{P}_i^*)\right] = \varphi_i(\tilde{P}_i^*)
\end{aligned}$$

where we use the terms $\tilde{\pi}_i(a_i, \bar{\mathbf{s}}_{-i}^{\tilde{\mathbf{\Psi}}(\tilde{\mathbf{P}})})$ and $e_i^{\tilde{P}}\left(a_i, \tilde{\Lambda}_i(\varphi_i(\tilde{P}_i^*))\right)$ instead of the original notation $\tilde{\pi}_i(a_i)$ and $e_i^{\tilde{P}}(a_i)$ to highlight the role of the (exogenously given) competitor's steady-state vector and choice probability vector, respectively.

We now show that there is an unique oblivious profile $\tilde{\mathbf{P}}^*$ such that $\tilde{\mathbf{P}}^* = \tilde{\mathbf{\Psi}}(\tilde{\mathbf{P}}^*)$. Suppose that there are two different oblivious equilibria, $\tilde{\mathbf{P}}^*$ and $\tilde{\mathbf{P}}^{**}$. Then it follows that $\tilde{P}_i^* = \tilde{\Psi}_i(\tilde{P}_i^*)$ and $\tilde{P}_i^{**} = \tilde{\Psi}_i(\tilde{P}_i^{**})$ for each player i. Consequently, both $\varphi_i(\tilde{P}_i^*)$ and $\varphi_i(\tilde{P}_i^{**})$ are fixed points of $\tilde{\Gamma}_i(.)$. As noted above, the fixed points of the operators $\tilde{\Gamma}_i(.)$ and $T(., \omega)$ are unique. This implies that $\varphi_i(\tilde{P}_i^*) = \varphi_i(\tilde{P}_i^{**})$ for every player i. Then, it follows that, for every player i, we have $\tilde{P}_i^* = \tilde{\Psi}_i(\tilde{P}_i^*) = \tilde{\Lambda}_i(\varphi_i(\tilde{P}_i^*)) = \tilde{\Lambda}_i(\varphi_i(\tilde{P}_i^{**})) = \tilde{\Psi}_i(\tilde{P}_i^{**}) = \tilde{P}_i^{**}$. But this means that $\tilde{\mathbf{P}}^* = (\tilde{P}_1^*, .., \tilde{P}_N^*) = (\tilde{P}_1^{**}, .., \tilde{P}_N^*) = \tilde{\mathbf{P}}^{**}$, which is a contradiction$\rightarrow\leftarrow$. ■

Proof of Theorem 2. *The distance between player value functions using Markovian and Oblivious strategies converges to zero in probability as the number of game players goes to infinity.*

Without loss of generality, consider that player i can choose a Markovian strategy defined over the set $\Theta \equiv \Omega \times X_i \times \mathbf{S}_{-i} \times \Sigma_i$ to solve his optimization problem given that all other players restrict attention to oblivious strategies defined over the set $\tilde{\Theta} \equiv \Omega \times X_i \times \Sigma_i$. Then his expected value function vector from using a Markovian strategy given that the other $N-1$ players follow an oblivious profile $\tilde{\mathbf{P}}_{-i}$ is

$$V_{i,N}^{P_i,\tilde{\mathbf{P}}_{-i}} = \sum_{a_i \in A_i} P_i(a_i) * \left[\pi_i(a_i) + e_i^{P_i}(a_i) + \beta F_i^{P_i,\tilde{\mathbf{P}}_{-i}}(a_i) V_i^{P_i,\tilde{\mathbf{P}}_{-i}}\right] \tag{4.1}$$

In addition, the expected value function if player i instead considers an oblivious strategy set $\tilde{\Theta}$ to maximize his discounted sum of payoffs is

$$V_{i,N}^{\tilde{\mathbf{P}}} = \sum_{a_i \in A_i} \tilde{P}_i(a_i) * \left[\pi_i(a_i) + e_i^{\tilde{P}_i}(a_i) + \beta F_i^{\tilde{\mathbf{P}}}(a_i) V_i^{\tilde{\mathbf{P}}}\right] \tag{4.2}$$

We establish our result by proving that each element of the vector $V_{i,N}^{P_i,\tilde{\mathbf{P}}_{-i}} - V_{i,N}^{\tilde{\mathbf{P}}}$ converges in probability to zero as $N \to \infty$. Letting $\tilde{V}_{i,N}^{\tilde{\mathbf{P}}}$ be player i's oblivious value function when all N players choose under the profile $\tilde{\mathbf{P}}$, it follows by triangular inequality that

$$\begin{aligned}
\left\| V_{i,N}^{P_i,\tilde{\mathbf{P}}_{-i}} - V_{i,N}^{\tilde{\mathbf{P}}} \right\|_\infty &= \left\| V_{i,N}^{P_i,\tilde{\mathbf{P}}_{-i}} - \tilde{V}_{i,N}^{\tilde{\mathbf{P}}} + \tilde{V}_{i,N}^{\tilde{\mathbf{P}}} - V_{i,N}^{\tilde{\mathbf{P}}} \right\|_\infty & (4.3)\\
&\leq \left\| V_{i,N}^{P_i,\tilde{\mathbf{P}}_{-i}} - \tilde{V}_{i,N}^{\tilde{\mathbf{P}}} \right\|_\infty + \left\| \tilde{V}_{i,N}^{\tilde{\mathbf{P}}} - V_{i,N}^{\tilde{\mathbf{P}}} \right\|_\infty & (4.4)\\
&\leq \left\| V_{i,N}^{P_i,\tilde{\mathbf{P}}_{-i}} - \tilde{V}_{i,N}^{P_i,\tilde{\mathbf{P}}_{-i}} \right\|_\infty + \left\| \tilde{V}_{i,N}^{\tilde{P}^*} - V_{i,N}^{\tilde{P}^*} \right\|_\infty & (4.5)
\end{aligned}$$

where the last inequality follows from the fact that an oblivious value function cannot be improved upon by considering Markovian strategies over Θ. Since each of the norms in the right-hand side of the inequality consider value function differences given the same probability profile, we have

$$\begin{aligned}
\left\| V_{i,N}^{P_i,\tilde{\mathbf{P}}_{-i}} - V_{i,N}^{\tilde{\mathbf{P}}} \right\|_\infty &\leq \left\| V_{i,N}^{P_i,\tilde{\mathbf{P}}_{-i}} - \tilde{V}_{i,N}^{P_i,\tilde{\mathbf{P}}_{-i}} \right\|_\infty + \left\| \tilde{V}_{i,N}^{\tilde{P}^*} - V_{i,N}^{\tilde{P}^*} \right\|_\infty & (4.6)\\
&\leq \left\| \left(I - \beta F_i^{P_i,\tilde{\mathbf{P}}_{-i}}\right)^{-1} \left\{ \sum_{a_i \in A_i} P_i(a_i) * [\pi_i(a_i) - \tilde{\pi}_i(a_i)] \right\} \right\|_\infty & (4.7)\\
&\quad + \left\| \left(I - \beta F_i^{\tilde{\mathbf{P}}}\right)^{-1} \left\{ \sum_{a_i \in A_i} \tilde{P}_i(a_i) * [\tilde{\pi}_i(a_i) - \pi_i(a_i)] \right\} \right\|_\infty & (4.8)
\end{aligned}$$

Thus, it suffices to have convergence in probability for payoff functions $\pi_i(.)$ as $N \to \infty$ to establish our result. By the uniform law of large numbers of the uniform law of large numbers of Vapnik and Chervonenkis (1971) the relative frequency vector $\boldsymbol{\eta}_t(\omega)$ converges in probability uniformily to a constant vector $\boldsymbol{\eta}(\omega)$ as

$N \to \infty$. regardless of the distribution from which players are sampled from. Using assumption 4 and the Mann-Wald Continuous Mapping theorem, the continuity of the payoff function $\pi_{i,N}(a_{it}, \omega_t, x_{it}, \boldsymbol{\eta}_t)$ in $\boldsymbol{\eta}_t$ ensures the uniform convergence in probability condition $\lim_{N\to\infty} \Pr\{|\pi_{i,N}(a_{it}, \omega_t, x_{it}, \boldsymbol{\eta}_t(\omega_t)) - \pi_{i,N}(a_{it}, \omega_t, x_{it}, \boldsymbol{\eta}(\omega_t))| > \varepsilon\} = 0$, $\forall \varepsilon > 0$, for any possible combination $(a_{it}, \omega_t, x_{it})$. By extending this argument for every element in $[\tilde{\pi}_i(a_i) - \pi_i(a_i)]$, it follows that $\lim_{N\to\infty} \Pr\left\{\left\|V_{i,N}^{P_i,\tilde{\mathbf{P}}_{-i}} - V_{i,N}^{\tilde{\mathbf{P}}}\right\|_\infty > \varepsilon\right\} = 0$, $\forall \varepsilon > 0$, which completes our proof. ■

References

[1] Abbring, J. and J. Campbell (2009): "Last-In First-Out Oligopoly Dynamics" forthcoming *Econometrica.*

[2] Abbring, J., J. Campbell and N. Yang (2010): "Simple Markov-perfect industry dynamics" , Tilburg University, mimeo.

[3] Ackerberg, D., C. L. Benkard, S. Berry and A. Pakes (2006): "Econometric Tools for Analyzing Market Outcomes", working paper prepared for J. Heckman and E. Leamer (eds), *Handbook of Econometrics*, Volume 6.

[4] Aguirregabiria V. (2009): "A method of Implementing Counterfactual Experiments in Models with Multiple Equilibria," University of Toronto, mimeo.

[5] Aguirregabiria, V. and C. Ho (2009a): "A Dynamic Oligopoly Game of the US Airline Industry: Estimation and Policy Experiments," University of Toronto, mimeo.

[6] Aguirregabiria, V. and C. Ho (2009b): "A Dynamic Game of Airline Network Competition: Hub-and-Spoke Networks and Entry Deterrence," forthcoming *International Journal of Industrial Organization.*

[7] Aguirregabiria, V. and P. Mira (2002): "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models", *Econometrica*, Vol.70, No. 1, pp 1519-1543.

[8] Aguirregabiria, V. and P. Mira (2007): "Sequential Estimation of Dynamic Discrete Games", *Econometrica*, Vol.75, No. 1, pp 1-53.

[9] Aguirregabiria, V. and P. Mira (2008): "Structural Estimation og Games with Multiple Equilibria in the Data," University of Toronto, mimeo.

[10] Aguirregabiria, V. and P. Mira (2009): "Dynamic Discrete Choice Structural Models: a Survey," forthcoming *Journal of Econometrics.*

[11] Aguirregabiria, V., P. Mira and H. Roman (2007): "An Estimable Dynamic Model of Entry, Exit and Growth in Oligopoly Retail Markets," *American Economic Review*, Vol.97, No. 2, pp 449-454.

[12] Aguirregabiria, V. and G. Vicentini (2006): "Dynamic Spatial Competition Between Multi-Store Firms," University of Toronto, mimeo.

[13] Bajari, P., C. L. Benkard and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition", *Econometrica*, Vol. 75, No. 5, pp 1331-1370.

[14] Bajari, P., J. Fox and S. Ryan (2007): "Linear Regression Estimation of Discrete Choice Models with Nonparametric Distributions of Random Coefficients," *American Economic Review*, Vol. 97 Issue 2, 459-463.

[15] Bajari, P., J. Fox, K. Kim and S. Ryan 2009: "A Simple Nonparametric Estimator for the Distribution of Random Coefficients," mimeo, University of Minnesota.

[16] Bajari, P., H. Hong and S. Ryan (2009): "Identification and Estimation of Dynamic Discrete Games of Complete Information," forthcoming, *Econometrica.*

[17] Bajari, P., V. Chernozhukov, H. Hong and D. Nekipelov (2009): "Semiparametric Estimation of Dynamic Discrete Game of Incomplete Information," mimeo, University of Minnesota.

[18] Benkard, C.L. (2004): "A Dynamic Analysis of the Market for Wide-bodied commercial aircraft", *The Review of Economic Studies*, Vol. 71 pp 581-611.

[19] Beresteanu, A. and P. Ellickson (2009): "The Dynamics of Retail Oligopoly", Duke University, mimeo.

[20] Berry, S., M. Ostrovsky and A. Pakes (2007): "Simple Estimators for the Parameters of Dynamic Games with Entry and Exit Examples)", *RAND Journal of Economics,* Vol..38, No. 2, pp.373-399.

[21] Bresnahan, T. and P. Reiss (1993): "Measuring the importance of sunk costs," *Annales d'Economie et de Statistique* Vol. 31, pp.181–217.

[22] Collard-Wexler, A. (2009): "Plant Turnover and Demand Fluctuations in the Ready-Mix Concrete Industry", Stern School of Business, New York University, mimeo.

[23] Doraszelski, U. (2003): "A R&D Race with Knowledge Accumulation", *The RAND Journal of Economics*, Vol. 34, No. 1, 2003, pp. 19-41.

[24] Doraszelski, U. and K. Judd (2009): "Avoiding the Curse of Dimensionality in Dynamic Stochastic Games", Harvard University, mimeo.

[25] Doraszelski, U. and M. Satherthwaite (2009): "Computable Markov-Perfect Industry Dynamics" forthcoming *RAND Journal of Economics.*

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

[27] Gowrisankaran, G. (1999): "A Dynamic Model of Endogenous Horizontal Mergers", *RAND Journal of Economics*, Vol. 30 pp 56-83.

[28] Gowrisankaran, G. and R. Town (1997): "Dynamic Equilibrium in the Hospital Industry", *Journal of Economics and Management Strategy*, Vol. 6, No. 1 pp 45-64.

[29] Hotz, J., and R. Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models", *The Review of Economic Studies*, Vol. 60, No. 2, pp 497-529.

[30] Hotz, J., R. Miller, S. Sanders and J. Smith (1994): "A Simulator Estimator for Dynamic Models of Discrete Choice", *The Review of Economic Studies*, Vol. 61, No. 2, pp 265-289.

[31] Macieira, J. (2009a): "A Dynamic Model of Innovation and Technological Competition in the Supercomputer Industry" Virginia Tech, mimeo.

[32] Macieira, J. (2009b): "Mergers with Innovation in Differentiated Product Industries: the Case of Supercomputers" Virginia Tech, mimeo.

[33] Magnac, T. and D. Thesmar (2002): "Identifying Dynamic Discrete Decision Processes", *Econometrica*, Vol. 70, No. 2, pp. 801-816.

[34] Maskin, E. and J. Tirole (1988a): "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs," *Econometrica*, Vol. 56, 549-570.

[35] Maskin, E. and J. Tirole (1988b): "A Theory of Dynamic Oligopoly, II: Price Competition," *Econometrica*, Vol. 56, 571-600.

[36] Maskin, E. and J. Tirole (2001), "Markov Perfect Equilibrium: Observable Actions," *Journal of Economic Theory*, Vol. 100, pp. 191-219.

[37] Nevo, A. and V. Aguirregabiria (2009): "Recent Developments in the Econometrics of Dynamic Games of Oligopoly Competition", Northwestern University, mimeo.

[38] Nevo, A. and F. Rossi (2007): "An approach for extending dynamic models to settings with multi-product firms", Northwestern University, mimeo.

[39] Pakes, A. and P. McGuire (1994): "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model", *The RAND Journal of Economics*, Vol. 25 pp 555-589.

[40] Pakes, A. and P. McGuire (2001): "Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'Curse' of Dimensionality", *Econometrica*, Vol. 69, No. 5, pp 1261-1281.

[41] Pesendorfer, M. and M. Jofre-Bonet (2003):"Sequential Estimation of Dynamic Discrete Games," *Econometrica*, Vol.71, No. 5, pp 1443-1489.

[42] Pesendorfer, M. and P. Schmidt-Dengler (2003): "Identification and Estimation of Dynamic Games," NBER working paper 9276.

[43] Pesendorfer, M. and P. Schmidt-Dengler (2008): "Asymptotic Least-Squares Estimators for Dynamic Discrete Games," *The Review of Economic Studies*, Vol. 78, No. 3, pp 901-928.

[44] Pesendorfer, M. and P. Schmidt-Dengler (2009): "Sequential Estimation of Dynamic Discrete Games," forthcoming *Econometrica*.

[45] Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher", *Econometrica*, Vol. 55, No. 5, pp 999-1033.

[46] Rust, J, J. Traub and H. Wozniakowski (2002): "Is there a Curse of Dimensionality for Contraction Fixed Points in the Worst Case?,", *Econometrica*, Vol. 70, No. 1, pp 285-329.

[47] Ryan, S. (2009): "The Costs of Environmental Regulation in a Concentrated Industry," MIT, mimeo.

[48] Ryan, S., E. Duflo and R. Hanna (2009): "Incentives Work: Getting Teachers to Come to School", MIT, mimeo.

[49] Ryan, S. and K. Tucker (2009): "Estimation of Demand with Dynamics, Network Effects and Heterogeneity", MIT, mimeo.

[50] Song, Minjae (2006): "A Dynamic Analysis of Cooperative Research in the Semiconductor Industry", School of Economics, Georgia Institute of Technology, mimeo.

[51] Sweeting, A. (2009): "Dynamic Product Repositioning in Differentiated Product Industries: The Case of Format Switching in the Commercial Radio Industry", Duke University, mimeo.

[52] Wang, H. and Z. Yang (1995): "On a Multinomial Markov Distribution", *Mathematics Scientist*, Vol. 20, pp 40-49.

[53] Weintraub, G. C.L Benkard and B. Van Roy (2006): "Oblivious Equilibrium: A Mean Field Approximation for Large Scale Dynamic Games," in *Advances in Neural Information Processing Systems*, Vol. 18.

[54] Weintraub, G. C.L Benkard and B. Van Roy (2008a): "Markov Perfect Industry Dynamics with Many Firms," *Econometrica*, Vol. 55, No. 5, pp 999-1033.

[55] Weintraub, G. C.L Benkard and B. Van Roy (2008b): "Nonstationary Oblivious Equilibria," Stanford University, mimeo.

[56] Weintraub, G. C.L Benkard and B. Van Roy (2009a):"Computational Methods for Oblivious Equilibrium," forthcoming *Operations Research.*

[57] Weintraub, G. C.L Benkard and B. Van Roy (2009b): "Industry Dynamics: Foundations for Models with an Infinite Number of Firms," Stanford University, mimeo.

[58] Xu, Y. (2008): "A Structural Empirical Model of R&D Investment, Firm Heterogeneity, and Industry Evolution", New York University, mimeo.