Department of Economics- FEA/USP

Ordinary Least Squares Estimation for a Dynamic Game

FÁBIO A. MIESSI SANCHES, DANIEL SILVA JUNIOR AND SORAWOOT SRISUMA

WORKING PAPER SERIES Nº 2014-19

DEPARTMENT OF ECONOMICS, FEA-USP
WORKING PAPER Nº 2014-19

Ordinary Least Squares Estimation for a Dynamic Game

Fabio A. Miessi Sanches - fmiessi@gmail.com

Daniel Silva Junior - d.silva-junior@lse.ac.uk

Sorawoot Srisuma - s.srisuma@surrey.ac.uk

Abstract:

Estimation of dynamic games is known to be a numerically challenging task. A common form of the payoff functions employed in practice takes the linear-in-parameter specification. We show a least squares estimator taking a familiar OLS/GLS expression is available in such case. Our proposed estimator has a closed-form. It can be computed without any numerical optimization and always minimizes the least squares objective function. Our estimator is also asymptotically equivalent to the asymptotic least squares estimator of Pesendorfer and Schmidt-Dengler (2008). Our estimator appears to perform well in a simple Monte Carlo experiment.

Keywords: Closed-form Estimation, Dynamic Discrete Choice, Markovian Games.

JEL Codes: C14, C25, C61

Ordinary Least Squares Estimation of a Dynamic Game Model*

Fabio A. Miessi Sanches†
University of São Paulo

Daniel Silva Junior‡
London School of Economics

Sorawoot Srisuma§
University of Surrey

February 14, 2015

Abstract

Estimation of dynamic games is known to be a numerically challenging task. A common form of the payoff functions employed in practice takes the linear-in-parameter specification. We show a least squares estimator taking a familiar OLS/GLS expression is available in such case. Our proposed estimator has a closed-form. It can be computed without any numerical optimization and always minimizes the least squares objective function. Our estimator is also asymptotically equivalent to the asymptotic least squares estimator of Pesendorfer and Schmidt-Dengler (2008). Our estimator appears to perform well in a simple Monte Carlo experiment.

JEL Classification Numbers: C14, C25, C61

Keywords: Closed-form Estimation, Dynamic Discrete Choice, Markovian Games.

*This paper is an updated version of "An Alternative Asymptotic Least Squares Estimator for Dynamic Games" (November, 2013). We are grateful to Martin Pesendorfer for encouragement and support. We also thank Joachim Groeger, Emmanuel Guerre, Oliver Linton, Robert Miller, Pasquale Schiraldi, Richard Smith and Dimitri Szerman for useful advice and comments.

† *E-mail address*: fmiessi@gmail.com

‡ *E-mail address*: d.silva-junior@lse.ac.uk

§ *E-mail address*: s.srisuma@surrey.ac.uk

1 Introduction

We consider the computational aspect for estimating a popular class of dynamic games in an infinite time horizon, where players' private values enter the payoff function additively and are independent across players, under the conditional independence framework. Recent surveys for such model can be found in Aguirregabiria and Mira (2010) and Bajari, Hong and Nekipelov (2012). A variety of methods have been proposed to estimate these games in recent years; examples are given below. However, a common component of the methodologies in the literature is a nonlinear optimization problem that may act as a considerable deterrent for applied researchers to estimate dynamic games due to involved programming needs and/or long computational time.

In this note we propose a simple class of least squares estimators that have closed-form when the payoffs have a linear-in-parameter specification. Our estimator takes a familiar OLS expression in the simplest case, and the efficient version has the GLS form. The linear parameterization can be quite general. In games with finite states linear-in-parameter payoff can be interpreted as nonparametric, otherwise it can generally represent any nonlinear (basis) functions of observables. In any case payoff with the linear-in-parameter structure is the leading specification employed in empirical work.

Estimation of dynamic games can be challenging. Games with multiple equilibria give rise to incomplete models, where each parameter corresponds to multiple probability distributions (Tamer (2003)). Even without the multiplicity issue, a full solution approach is computationally demanding since the game has to be solved for every parameter value (Rust (1994)). A popular approach to estimate dynamic games is to perform a two-step estimation procedure. Its origin can be traced back to the novel work of Hotz and Miller (1993) in a single agent setting, whose insight is to perform inference on a model that is generated using the empirical decision rule that can be estimated in the first-step from the observed choice and transition probabilities. Their idea is also applicable in a game context, where the empirical equilibrium strategy is used to compute any expected discounted payoffs without solving the game even once. We call the collection of implied probability distributions generated in this way the empirical model.

The choice probabilities implied by the empirical model in a dynamic game are characterized by the cumulative distribution function of the normalized additive private values and the index of expected discounted payoffs (cf. McFadden (1974)). Many existing two-step methodologies use choice probabilities to construct objective functions for estimation. Examples include traditional criterions such as the pseudo-likelihood approach (Aguirregabiria and Mira (2007), Kasahara and Shimotsu

(2012)) and other moment and minimum distance based conditions (Pakes, Ostrovsky, Berry (2007), Pesendorfer and Schmidt-Dengler (2008, PSD hereafter)). However, in order to calculate these probabilities, the implied expected discounted payoffs first have to be calculated. Furthermore, choice probabilities are written in terms of integrals that are generally nonlinear mappings of the expected payoffs that have to be computed numerically outside the well-known conditional logit framework.

The main purpose of our work is to emphasize that the integration step used to obtain choice probabilities adds an unnecessary computational cost. We define a class of least squares estimators based on minimizing the distance of the payoffs observed from the data and those implied by the empirical model directly. In particular, when the payoffs have a linear-in-parameter specification the expected discounted payoffs inherit this structure[1] so that our objective function has an expression that resembles a familiar linear regression problem. Different norming of the distance gives different least squares estimator. When we do not impose the linear parameterization, our least squares problem becomes nonlinear and has no closed-form solution. Our approach mirrors the *asymptotic least squares* methodology of PSD, who instead minimize distances in terms of probabilities. The estimators obtained using our approach and PSD's are asymptotically equivalent. PSD's estimator provides a good theoretical benchmark as it includes the non-iterative likelihood estimator of Aguirregabiria and Mira (2007) and the moment estimator of Pakes, Ostrovsky and Berry (2007) as special cases. We refer the reader to the previous version of our work for the proof of this result. This note only focuses on developing a closed-form estimator for dynamic games and highlighting its practical simplicity.

Other methodologies that use expected payoffs explicitly to construct objective functions also exist in the literature. The first such two-step estimator has been developed by Hotz, Miller, Sanders and Smith (1994), who estimate the expected payoffs by forward simulation, to estimate a dynamic decision problem for a single agent. Hotz et al. (1994) define their estimator using conditional moment restrictions. They also recognize it is possible to have a closed-form estimator when payoff functions have linear-in-parameter specification in the form of an IV estimator (see equation (5.8) in the Monte Carlo Study section of Hotz et al. (1994)). In the context of dynamic games, we are only aware of two other current methodologies that base their objective functions explicitly on expected

[1]Related discussions can be found in Bajari, Benkard and Levin (2007, Section 3.3.1) and Pakes, Ostrovsky, Berry (2007, Section 3).

payoffs. First is the two-step estimator proposed by Bajari, Benkard and Levin (2007), who also use forward simulation like Hotz et al. Although generally no closed-form estimator is possible with Bajari, Benkard and Levin's methodology as they compare expected payoffs in the empirical model and those generated by local perturbations. The other is Bajari, Chernozhukov, Hong and Nekipelov (2009), who provide nonparametric identification results for a more general game, with continuous state space, and propose an efficient one-step estimator.[2,3]

The remainder of this note is organized as follows. Section 2 defines the game and the empirical model. Section 3 defines our least squares estimator. Section 4 presents results from some Monte Carlo experiments that compare the statistical performance and relative speed of our estimator and that of PSD's. Section 5 concludes and discusses how our estimators can be used to complement other recent research in the literature. All proofs can be found in the Appendix.

2 Basic Setup

We consider a game with I players, indexed by $i \in \mathcal{I} = \{1, \ldots, I\}$. The elements of the game are:

Actions. The action set of each player is $A = \{0, 1, \ldots, K\}$. We denote the action variable for player i by a_{it}. Let $\mathbf{a}_t = (a_{1t}, \ldots, a_{It}) \in \mathbf{A} = \times_{i=1}^{I} A$. We will also occasionally abuse the notation and write $\mathbf{a}_t = (a_{it}, \mathbf{a}_{-it})$ where $\mathbf{a}_{-it} = (a_{1t}, \ldots, a_{i-1t}, a_{i+1t} \ldots, a_{It}) \in \mathbf{A} \backslash A$.

States. Player i's information set is represented by the state variables $s_{it} \in S$, where $s_{it} = (x_{it}, \varepsilon_{it})$ such that $x_{it} \in X$ is common knowledge to all players and $\varepsilon_{it} = (\varepsilon_{it}(1), \ldots, \varepsilon_{it}(K)) \in \mathcal{E}$ denotes private information only observed by player i. We define $\varepsilon_t = (\varepsilon_{1t}, \ldots, \varepsilon_{It})$. Note that we exclude the private value associated with action 0 for the purpose of normalization. We shall use s_{it} and (x_t, ε_{it}) interchangeably.

State Transition. Future states are uncertain. Players' actions and states today affect future states. The evolution of the states is summarize by a Markov transition law $P(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$.

Per Period Payoff Functions. Each player has a payoff function, $u_i : \mathbf{A} \times S \to \mathbb{R}$.

Discounting Factor. Future period's payoffs are discounted at the rate $\beta \in [0, 1)$.

[2] An earlier version of Bajari et al. (2009), Bajari and Hong (2006), proposes a two-step estimator that can be seen as the dynamic game version of Hotz et al. (1994).

[3] Another notable estimator that does not take a two-step approach is Egesdal, Lai and Su (2012). Although Egesdal et al. construct their objective functions in terms of choice probabilities.

We also impose the following assumptions, which are standard in the literature (e.g. see Aguirregabiria and Mira (2007), Pakes et al. (2008) and PSD).

Assumption M1 *(Additive Separability). $u_i(a_i, \mathbf{a}_{-i}, x, \varepsilon_i) = \pi_i(a_i, \mathbf{a}_{-i}, x) + \varepsilon_i(a_i)\mathbf{1}[a_i > 0]$ for all $i, a_i, \mathbf{a}_{-i}, x, \varepsilon_i$.*

Assumption M2 *(Conditional independence). The transition distribution of the states has the following factorization: $P(x_{t+1}, \varepsilon_{t+1}|x_t, \varepsilon_t, \mathbf{a}_t) = Q(\varepsilon_{t+1}) G(x_{t+1}|x_t, \mathbf{a}_t)$, where Q is the cumulative distribution function of ε_t and G denotes the transition law of x_{t+1} conditioning on $\mathbf{a}_t$ and x_t.*

Assumption M3 *(Independent private values). The private information is independently distributed across players, and each is absolutely continuous with respect to the Lebesgue measure whose density is bounded on $\mathbb{R}^K$.*

Assumption M4 *(Discrete public values). The support of x_t is finite so that $X = \{x^1, \ldots, x^J\}$ for some $J < \infty$.*

We consider an infinite time horizon game, where at time t, each player i observes s_{it} then chooses a_{it} simultaneously. Players are assumed to use stationary pure Markov strategies, α_i, so that $\alpha_i : S \to A$, $a_{it} = \alpha_i(s_{it})$ for all i, t, and whenever $s_{it} = s_{i\tau}$ then $\alpha_i(s_{it}) = \alpha_i(s_{i\tau})$ for any τ. Player i's beliefs, σ_i, is a distribution of $\mathbf{a}_t = (\alpha_1(s_{1t}), \ldots, \alpha_I(s_{It}))$ conditional on x_t for some strategy profile $(\alpha_1, \ldots, \alpha_I)$. The decision problem for each player is to solve:

$$\max_{a_i \in A_i}\{E_{\sigma_i}[u_i(a_{it}, \mathbf{a}_{-it}, s_i)|s_{it} = s_i, a_{it} = a_i] + \beta E_{\sigma_i}[W_i(s_{it+1}; \sigma_i)|s_{it} = s_i, a_{it} = a_i]\}, \tag{1}$$
$$\text{where } W_i(s_i; \sigma_i) = \sum_{\tau=0}^{\infty} \beta^\tau E_{\sigma_i}[u_i(\mathbf{a}_{t+\tau}, s_{it+\tau})|\, s_{it} = s_i],$$

for any s_i. Under M1 and M2, it is Player's i best response to choose action a_i if for all $a_i' \neq a_i$:

$$\begin{aligned} & E_{\sigma_i}[\pi_i(a_i, a_{-it}, x_t)|\, x_t = x] + \beta E_{\sigma_i}[W_i(s_{t+1}; \sigma_i)|\, x_t = x, a_{it} = a_i] + \varepsilon_i(a_i) \\ \geq \; & E_{\sigma_i}[\pi_i(a_i', a_{-it}, x_t)|\, x_t = x] + \beta E_{\sigma_i}[W_i(s_{t+1}; \sigma_i)|\, x_t = x, a_{it} = a_i'] + \varepsilon_i(a_i'). \end{aligned} \tag{2}$$

The subscript σ_i on the expectation operator makes explicit that present and future actions are integrated out with respect to the beliefs σ_i. $W_i(\cdot; \sigma_i)$ is a policy value function, where σ_i can be any beliefs, not necessarily equilibrium beliefs. Therefore the induced transition laws for future states are completely determined by the primitives and σ_i. Any strategy profile that solves the decision problems for all i, and is consistent with the beliefs, is an equilibrium strategy. It is well-known that players' best responses are pure strategies almost surely and Markov perfect equilibria for games under M1 - M4. Further details can be found in Aguirregabiria and Mira (2007) and PSD.

An Empirical Model

The starting point is the structural assumption that we observe a random sample of $\{\mathbf{a}_t, x_t, x_{t+1}\}$ from a single equilibrium, where each a_{it} in $\mathbf{a}_t$ equals $\alpha_i(s_{it})$. Let $P_i(a_i|x) \equiv \Pr[a_{it} = a_i|x_t = x]$ for all a_i, x denote the equilibrium conditional choice probabilities. Then we have: (i) the equilibrium beliefs for all players is summarized by $\prod_{i=1}^{I} P_i$, and (ii) $\Pr[x_{t+1} = x'|x_t = x, \mathbf{a}_t = \mathbf{a}] = G(x'|x, \mathbf{a})$ for all $\mathbf{a}, x, x'$. In common with the related papers cited above, we shall also assume β and Q are known throughout. Therefore the knowledge of $(\prod_{i=1}^{I} P_i, G, Q)$ can be used to construct the stationary equilibrium decision rule that is consistent with the data generating process.

We next parameterize the payoff function. The payoff parameter for each player is denoted by $\theta_i \in \Theta_i \subset \mathbb{R}^{p_i}$, and we overwrite the payoff function associated with the observed variables in M1 by π_{i,θ_i}. Let $\theta_0 = (\theta_{10}^{\top}, \ldots, \theta_{I0}^{\top})^{\top} \in \Theta \equiv \times_{i=1}^{I}\Theta_i$ be the data generating parameter of interest.

The (conditional) probability distribution of the empirical model can be thought of as being derived from the following decision problem. For any θ_i, consider (cf. (1)):

$$\max_{a_i \in A_i}\{E[\pi_{i,\theta_i}(a_i, \mathbf{a}_{-it}, x_t)|\, x_t = x] + \varepsilon_i(a_i)\mathbf{1}[a_i > 0] + \beta E[V_{i,\theta_i}(s_{t+1})|\, x_t = x, a_{it} = a_i]\},$$
$$\text{where } V_{i,\theta_i}(s_i) = \sum_{\tau=0}^{\infty}\beta^{\tau}E[\pi_{i,\theta_i}(\mathbf{a}_{it+\tau}, x_{it+\tau}) + \sum_{a'>0}\varepsilon_{it+\tau}(a')\mathbf{1}[a_{it+\tau} = a']\,|s_t = s_i].$$

Here V_{i,θ_i} is the empirical policy value function, where all players use the equilibrium strategy observed in the data. Note that we have omitted the dependence on the beliefs for notational convenience. Then we can define the implied choice specific expected discounted payoffs as:

$$v_{i,\theta_i}(a_i, x) = E[\pi_{i,\theta_i}(a_i, \mathbf{a}_{-it}, x_t)|\, x_t = x] + \beta E[V_{i,\theta_i}(s_{t+1})|\, x_t = x, a_{it} = a_i]. \tag{3}$$

The implied choice probabilities can also be written in terms of differences in choice specific expected payoffs. Let $\Delta v_{i,\theta_i}(a_i, x)$ denote $v_{i,\theta_i}(a_i, x) - v_{i,\theta_i}(0, x)$ for $a_i > 0$ and any x, then we define:

$$P_{i,\theta_i}(a_i|x) = \Pr[\Delta v_{i,\theta_i}(a_i, x_t) + \varepsilon_{it}(a_i) > \Delta v_{i,\theta_i}(a_i', x_t) + \varepsilon_{it}(a_i') \text{ for all } a_i' > 0|x_t = x],$$

and $P_{i,\theta_i}(0|x) = 1 - \sum_{a_i>0} P_{i,\theta_i}(a_i|x)$.

The empirical model can now be defined as $\{P_\theta\}_{\theta\in\Theta}$ such that $P_\theta = \prod_{i=1}^{I} P_{i,\theta_i}$. By the structural assumption, that we observe outcomes of an equilibrium play, $P_{i,\theta_{i0}}$ must equal P_i for all i (see equation (2)). Therefore the empirical model can be useful for the purpose of estimating θ_0.[4] In particular the form of P_{i,θ_i} is familiar from the classical random utility model (e.g. see McFadden (1974)) with a normalized index mean utility of $\Delta v_{i,\theta_i}$.

We shall focus on the form of v_{i,θ_i} when π_{i,θ_i} has a linear-in-parameter specification.

Assumption M5 *(Linear-in-parameter payoffs). For all* $(i, \theta_i, a_i, \mathbf{a}_{-i}, x)$,

$$\pi_{i,\theta_i}(a_i, \mathbf{a}_{-i}, x) = \theta_i^\top \pi_{i0}(a_i, \mathbf{a}_{-i}, x),$$

for some p_i-*dimensional vector* $\pi_{i0}(a_i, \mathbf{a}_{-i}, x) = \left(\pi_{i0}^1(a_i, \mathbf{a}_{-i}, x), \ldots, \pi_{i0}^{p_i}(a_i, \mathbf{a}_{-i}, x)\right)^\top$, *where* $p_i < J$.

The requirement $p_i < J$ ensures π_{i,θ_i} satisfies a necessary order condition on the payoffs for identification as the game under consideration is generally under-identified (Proposition 2 in PSD).

The term v_{i,θ_i} appears complicated as it is written in terms of expectations of present and future payoffs. It shall be helpful to re-write a version of equation (3) here, where V_{i,θ_i} is expressed explicitly in terms of the sum of future discounted payoffs:

$$\begin{aligned} v_{i,\theta_i}(a_i, x) &= \underline{v}_{i,\theta_i}(a_i, x) + \underline{v}_i(a_i, x), \text{ where} \\ \underline{v}_{i,\theta_i}(a_i, x) &= \sum_{\tau=0}^{\infty} \beta^\tau E[\pi_{i,\theta_i}(a_{it}, \mathbf{a}_{-it+\tau}, x_{it+\tau}) | x_t = x, a_{it} = a_i], \\ \underline{v}_i(a_i, x) &= \sum_{\tau=0}^{\infty} \beta^{\tau+1} E[\sum_{a'>0} \varepsilon_{it+\tau+1}(a') \mathbf{1}[a_{it+\tau+1} = a'] | x_t = x, a_{it} = a_i]. \end{aligned} \tag{4}$$

Since expectations and summations are linear operations, $\underline{v}_{i,\theta_i}$ can be written as some linear combination of π_{i,θ_i}, so that under M5, $\underline{v}_{i,\theta_i}(a_i, x) = \theta_i^\top \underline{v}_{i0}(a_i, x)$ for some p_i-dimensional vector $\underline{v}_{i0}(a_i, x) = \left(\underline{v}_{i0}^1(a_i, x), \ldots, \underline{v}_{i0}^{p_i}(a_i, x)\right)^\top$. Therefore the linear-in-parameter structure of π_{i,θ_i} is inherited by v_{i,θ_i}. Furthermore, since the support of (a_{it}, x_t) is finite, we have a matrix representation for $\{\Delta v_{i,\theta_i}(a_i, x)\}_{a_i>0, x\in X}$ which we now state as a lemma.

[4] Since all of the expectations in v_{i,θ_i} are calculated using the same equilibrium beliefs observed from the data, there is no need to solve the game for any θ_i .

Lemma R: Under M1 - M5 $\{\Delta v_{i,\theta_i}(a_i, x)\}_{a_i>0, x\in X}$ can be represented by a $JK-$vector $\Delta \mathbf{v}_{i,\theta_i}$:

$$\Delta \mathbf{v}_{i,\theta_i} = \mathcal{X}_i \theta_i + \Delta \underline{\mathbf{v}}_i, \tag{5}$$

for some JK by p_i matrix $\mathcal{X}_i$ and a $JK-$vector $\Delta \underline{\mathbf{v}}_i$.

We provide the detailed compositions of $\mathcal{X}_i$ and $\Delta \underline{\mathbf{v}}_i$ in Appendix A. For the moment it suffices to say they are known in terms of $(\beta, \prod_{i=1}^{I} P_i, G, Q)$.

3 Closed-Form Least Squares Estimation

Under the continuity of the distribution of ε_{it} with large support (M3), there is an invertible map relating $\{P_{i,\theta_i}(a_i|x)\}_{a_i>0, x\in X}$ and $\{\Delta v_{i,\theta_i}(a_i, x)\}_{a_i>0, x\in X}$ (e.g. Proposition 1 of Hotz and Miller (1993)). Let $\mathbf{P}_{i,\theta_i}$ denote a $JK-$vector of $\{P_{i,\theta_i}(a_i|x)\}_{a_i>0, x\in X}$. We denote the invertible map by Φ_i so that $\Delta \mathbf{v}_{i,\theta_i} = \Phi_i(\mathbf{P}_{i,\theta_i})$ for every i, θ_i. Similarly we can define the vectors of choice probabilities and expected discounted payoffs observed from the data. Let $\mathbf{P}_i$ denote a $JK-$vector of $\{P_i(a_i|x)\}_{a_i>0, x\in X}$ and $\Delta \mathbf{v}_i = \Phi_i(\mathbf{P}_i)$ be a vector of the same dimension. Then we can define a $JK-$vector $\mathcal{Y}_i$, where

$$\mathcal{Y}_i = \Phi_i(\mathbf{P}_i) - \Delta \underline{\mathbf{v}}_i.$$

Therefore by construction:

$$\mathcal{Y}_i = \mathcal{X}_i \theta_i \;\text{ when }\; \theta_i = \theta_{i0}. \tag{6}$$

Let $\mathcal{Y} = \left(\mathcal{Y}_1^\top, \ldots, \mathcal{Y}_I^\top\right)^\top$, $\theta = \left(\theta_1^\top, \ldots, \theta_I^\top\right)^\top$ and define a block diagonal matrix $\mathcal{X} = diag(\mathcal{X}_1, \ldots, \mathcal{X}_I)$. A natural estimator of θ_0 can be motivated from minimizing the sample counterpart of the following least squares criterion:

$$\mathcal{S}(\theta; \mathcal{W}) = (\mathcal{Y} - \mathcal{X}\theta)^\top \mathcal{W} (\mathcal{Y} - \mathcal{X}\theta),$$

where $\mathcal{W}$ is some positive definite (p.d.) weighting matrix.

It is also worth emphasizing that $\mathcal{X}$ and $\mathcal{Y}$ are *known* functions of $(\beta, \prod_{i=1}^{I} P_i, G, Q)$. Then, given a sample from a single equilibrium, $(\prod_{i=1}^{I} P_i, G)$ can be identified from the data under weak conditions. Consequently we consider an objective function where $(\mathcal{X}, \mathcal{Y})$ is replaced by some consistent estimator $(\widehat{\mathcal{X}}, \widehat{\mathcal{Y}})$ in the first-step. We denote the sample counterpart of $\mathcal{S}(\theta; \mathcal{W})$ by $\widehat{\mathcal{S}}(\theta; \widehat{\mathcal{W}})$, where for some p.d. matrix $\widehat{\mathcal{W}}$ that converges in probability to $\mathcal{W}$,

$$\widehat{\mathcal{S}}(\theta; \widehat{\mathcal{W}}) = (\widehat{\mathcal{Y}} - \widehat{\mathcal{X}}\theta)^\top \widehat{\mathcal{W}} (\widehat{\mathcal{Y}} - \widehat{\mathcal{X}}\theta).$$

Our estimator is defined to minimize $\widehat{\mathcal{S}}(\theta;\widehat{\mathcal{W}})$. If $\widehat{\mathcal{X}}$ has full column rank we obtain a closed-form least squares solution:

$$\begin{aligned}\widehat{\theta}(\widehat{\mathcal{W}}) &= \arg\min_{\theta\in\Theta}\widehat{\mathcal{S}}(\theta;\widehat{\mathcal{W}}) \\ &= (\widehat{\mathcal{X}}^{\top}\widehat{\mathcal{W}}\widehat{\mathcal{X}})^{-1}\widehat{\mathcal{X}}^{\top}\widehat{\mathcal{W}}\widehat{\mathcal{Y}}.\end{aligned} \tag{7}$$

The simplest estimator can be obtained by using the identity weighting, and the expression above simplifies to an OLS estimator: $(\widehat{\mathcal{X}}^{\top}\widehat{\mathcal{X}})^{-1}\widehat{\mathcal{X}}^{\top}\widehat{\mathcal{Y}}$. Under some mild regularity conditions our estimator is consistent and asymptotically normal. We provide some large sample results as well as a discussion of efficient estimation in Appendix B.

4 Numerical Illustration

We illustrate the performance of our estimator using the Monte Carlo design in Section 7 of PSD. Consider a two-firm dynamic entry game. In each period t, each firm $i(=1,2)$ has two possible choices, $a_{it}\in\{0,1\}$. Observed state variables are previous period's actions, $x_t=(a_{1t-1},a_{2t-1})$. Firm $1'$s period payoffs are described as follows:

$$\pi_{1,\theta}\left(a_{1t},a_{2t},x_t\right)=a_{1t}\left(\mu_1+\mu_2 a_{2t}\right)+a_{1t}\left(1-a_{1t-1}\right)F+\left(1-a_{1t}\right)a_{1t-1}W, \tag{8}$$

where $\theta=(\mu_1,\mu_2,F,W)$ denote respectively the monopoly profit, duopoly profit, entry cost and scrap value. Each firm also receives additive private shocks that are i.i.d. $\mathcal{N}(0,1)$. The game is symmetric and Firm's 2 payoffs are defined analogously. We also provide a detailed construction of $\mathcal{X}_i$ for this simple model in Appendix A.

We set $(\mu_{10},\mu_{20},F_0,W_0)=(1.2,-1.2,-0.2,0.1)$. PSD show there are three distinct equilibria for this game, one of which is symmetric. We generate the data using the symmetric equilibrium and estimate (μ_{10},μ_{20},F_0) while W_0 is assumed known for the purpose of identification. For each sample size $T=100,500,1000,5000$, using 1000 simulations, we report the same set of statistics as PSD for our OLS and GLS estimators, as well as their identity weighted and efficient asymptotic least squares estimators (denoted by PSD-I and PSD-E respectively). The results are collected in Table 1.[5] The estimators are consistent and their performance is similar across the two asymptotic least squares

[5] Our Table 1 corresponds to equilibria (iii) in PSD, and it can be compared directly with Table 3 in their paper on page 922.

approaches. We find similar results with data generated from other (non-symmetric) equilibria. We do not report these results for the sake of space.

We also study the computational time taken to construct the estimators. We introduce an additive market fixed effect to the per-period payoff in the game described above. We use the number of markets, denoted by M, to control the complexity of the game.[6] For each M, we solve the model for the symmetric equilibrium and simulate it five times. Table 2 reports the average central processing unit (CPU) time in seconds taken to compute our OLS and GLS estimators and for PSD-I and PSD-E. The standard errors of the computing time are reported in parentheses.[7] Our closed-form estimators are substantially faster to compute, which is not surprising, and the distinction grows exponentially with more parameters in the model. The reported CPU time also includes the construction of the optimal weighting matrices. Since the procedure to compute the optimal weighting matrices are similar for both estimators, its contribution in this setting can be approximated by comparing the CPU time taken to estimate OLS and GLS as M varies. More generally, we also expect the computation time for PSD's estimator to grow at a faster rate with larger action and/or state spaces for any fixed M relative to our closed-form approach.

Another numerical property of our estimator that is not quantified above is it trivially obtains the global minimizer. In contrast, a numerical solution to a general nonlinear optimization routine may be sensitive to the search algorithm and initial values.

5 Conclusions and Possible Extensions

There can be a substantial computational advantage in defining objective functions in terms of payoffs instead of probabilities for the estimation of dynamic games. We propose a class of closed-form least squares estimators when the commonly used linear-in-parameter payoff is employed. Closed-form estimation is attractive for its simplicity and stability compared to any search algorithm. Our estimators are asymptotically equivalent to those proposed by Pesendorfer and Schmidt-Dengler (2008), which include other well-known estimators in the literature.

[6] There are other ways to vary the complexity of the game, e.g. by changing the number of potential actions and states. However, the difficulty to solve and estimate such game increases significantly as the game becomes more complex. Our design is chosen for its simplicity as it only requires us to solve a simple game multiple times.

[7] The simulation was performed using MATLAB (R2012a, 64 bit version) on a standard PC running on an Intel Core (TM) 2 Duo 3.16 GHz processor with 4 GB RAM.

T	Estimator	F_0		μ_{10}		μ_{20}		MSE
100	OLS	-0.304	(0.475)	0.997	(0.398)	-0.895	(0.558)	0.840
	GLS	-0.436	(0.356)	1.015	(0.352)	-0.88	(0.446)	0.641
	PSD-I	-0.241	(0.514)	1.102	(0.471)	-1.023	(0.624)	0.917
	PSD-E	-0.397	(0.445)	1.081	(0.381)	-0.975	(0.526)	0.722
500	OLS	-0.225	(0.244)	1.149	(0.187)	-1.118	(0.282)	0.184
	GLS	-0.260	(0.229)	1.159	(0.185)	-1.122	(0.278)	0.175
	PSD-I	-0.201	(0.258)	1.200	(0.222)	-1.176	(0.304)	0.208
	PSD-E	-0.230	(0.239)	1.177	(0.189)	-1.157	(0.287)	0.178
1000	OLS	-0.214	(0.177)	1.169	(0.134)	-1.158	(0.204)	0.093
	GLS	-0.227	(0.170)	1.179	(0.136)	-1.166	(0.206)	0.092
	PSD-I	-0.202	(0.180)	1.193	(0.147)	-1.187	(0.211)	0.099
	PSD-E	-0.207	(0.186)	1.191	(0.148)	-1.188	(0.220)	0.105
5000	OLS	-0.203	(0.082)	1.194	(0.062)	-1.190	(0.093)	0.019
	GLS	-0.205	(0.076)	1.197	(0.060)	-1.192	(0.090)	0.017
	PSD-I	-0.201	(0.083)	1.200	(0.066)	-1.196	(0.095)	0.020
	PSD-E	-0.201	(0.078)	1.199	(0.061)	-1.197	(0.094)	0.018

Table 1: *Monte Carlo results. OLS and GLS are our closed-form estimators. PSD-I and PSD-E are respectively the identity weighted and efficient estimators of Pesendorfer and Schmidt-Dengler (2008).*

M	1	10	20	30	100	200
OLS	0.0021	0.0125	0.0245	0.0366	0.1241	0.2654
	(0.0010)	(0.0000)	(0.0000)	(0.0001)	(0.0004)	(0.0004)
GLS	0.0180	0.1542	0.3091	0.4658	1.8504	5.6084
	(0.0038)	(0.0001)	(0.0013)	(0.0002)	(0.0023)	(0.0069)
PSD-I	0.2084	4.9957	28.6415	73.3173	1171.5137	5657.6393
	(0.0089)	(0.0351)	(0.1805)	(0.0846)	(1.9478)	(0.9183)
PSD-E	0.3564	10.4140	52.0471	109.5519	1607.2349	7621.5963
	(0.0079)	(0.0359)	(0.1824)	(0.1049)	(2.6654)	(1.2093)

Table 2: *Computation time. OLS and GLS are our closed-form estimators. PSD-I and PSD-E are respectively the identity weighted and efficient estimators of Pesendorfer and Schmidt-Dengler (2008).*

The computational gain from closed-form estimation accumulates beyond point estimation. Any iteration or resampling algorithms (e.g. to compute standard errors) would clearly benefit. In particular, for the former, the bias reduction procedures in Aguirregabiria and Mira (2007) and Kasahara and Shimotsu (2012) can use OLS/GLS estimator at each step of iteration instead of a pseudo-likelihood estimator. It would be interesting to verify if the asymptotic equivalence still holds with such iteration procedure.

Analogous closed-form estimation is also possible in some models with common unobserved heterogeneity and/or for empirical games with multiple equilibria. This follows since, in principle, Hotz and Miller's (1993) two-step approach can be used whenever a nonparametric estimator is available to construct an empirical model that is consistent with the observed data in the first step. For example, see Aguirregabiria and Mira (2007, Section 3.5), where nonparametric identification results of Kasahara and Shimotsu (2009) can be applied. Therefore we are hopeful that closed-form estimation based on minimizing expected payoffs is generally possible beyond the basic setup of our game, particularly given recent identification results for games with multiple equilibria (Aguirregabiria and Mira (2013), Xiao (2014)) and other dynamic models with latent state variables (e.g. Hu and Shum (2012)).

Appendix A - Representation Lemma

PROOF OF LEMMA R. We first write v_{i,θ_i} in equation (4) in a matrix form. The conditional expectations of discrete random variables are just weighted sums they can be represented using matrices. In particular we can vectorize $\{v_{i,\theta_i}(a_i,x)\}_{a_i\in A,x\in X}$ into the following form,

$$\mathbf{v}_{i,\theta_i} = (\mathbf{R}_i + \beta\mathbf{H}_i\mathbf{MR})\,\mathbf{\Pi}_i\theta_i + \underline{\mathbf{v}}_i.$$

Note that $(\mathbf{R}_i + \beta\mathbf{H}_i\mathbf{MR})\,\mathbf{\Pi}_i$ is $\mathcal{X}_i$ in equation (5), which is central for estimation. Then $\mathbf{\Pi}_i\theta_i$ and $\underline{\mathbf{v}}_i$ represent $\{\pi_{i,\theta_i}(\mathbf{a},x)\}_{\mathbf{a}\in\mathbf{A},x\in X}$ and $\{\underline{v}_i(a_i,x)\}_{a_i\in A,x\in X}$ respectively, and:

Matrix:	Representing	Vector:	Representing
$\mathbf{R}_i$	$E[\phi(\mathbf{a}_{-it})\,\vert x_t=\cdot,a_{it}=\cdot]$	$\mathbf{R}_i\mathbf{\Pi}_i\theta_i$	$\{E[\pi_{i,\theta_i}(a_{it},\mathbf{a}_{-it},x_t)\,\vert x_t=x,a_{it}=a_i]\}_{a_i\in A,x\in X}$
$\mathbf{R}$	$E[\phi(\mathbf{a}_t)\,\vert x_t=\cdot]$	$\mathbf{R\Pi}_i\theta_i$	$\{E[\pi_{i,\theta_i}(\mathbf{a}_t,x_t)\,\vert x_t=x]\}_{x\in X}$
$\mathbf{M}$	$\sum_{\tau=0}^{\infty}\beta^\tau E[\phi(\mathbf{a}_{t+\tau},x_{it+\tau})\,\vert x_t=\cdot]$	$\mathbf{MR\Pi}_i\theta_i$	$\{\sum_{\tau=0}^{\infty}\beta^\tau E[\pi_{i,\theta_i}(a_{it},\mathbf{a}_{-it+\tau},x_{it+\tau})\,\vert x_t=x]\}_{x\in X}$
$\mathbf{H}_i$	$E[\phi(x_{t+1})\,\vert x_t=\cdot,a_{it}=\cdot]$	$\mathbf{H}_i\mathbf{MR\Pi}_i\theta_i$	$\{\sum_{\tau=1}^{\infty}\beta^\tau E[\pi_{i,\theta_i}(a_{it},\mathbf{a}_{-it+\tau},x_{it+\tau})\,\vert x_t=x,a_{it}=a_i]\}_{a_i\in A,x\in X}$

for any generic function ϕ. For the details of the matrices and vectors above, we need additional notations to those already defined in the main text. The representation of the choice specific expected payoffs in this paper stacks the vector in a repeating sequence of $\{x^j\}$ for each action. By writing $v^a_{i,\theta_i} = \left(v_{i,\theta_i}(a,x^1),\ldots,v_{i,\theta_i}\left(a,x^J\right)\right)$ for all $a\in A$, then $\mathbf{v}_{i,\theta_i} = \left(v^0_{i,\theta_i},\ldots,v^K_{i,\theta_i}\right)^\top$ is a $J(K+1)-$vector. Let $\pi_i^{a_1\ldots a_I} = \left(\pi_{i0}(a_1,\ldots,a_I,x^1),\ldots,\pi_{i0}\left(a_1,\ldots,a_I,x^J\right)\right)$ for all $a_1,\ldots,a_I$, and $\mathbf{\Pi}_i = \left(\pi_i^{0\ldots0},\ldots,\pi_i^{K\ldots K}\right)^\top$, so that $\mathbf{\Pi}_i$ is a $J(K+1)^I$ by p_i matrix. Then: $\mathbf{H}_i$ is a block-diagonal matrix $\mathbf{diag}\left(H_i^0,H_i^1,\ldots,H_i^K\right)$, where H_i^a is a $J\times J$ matrix such that $(H_i^a)_{jj'} = \Pr\left[x_{t+1}=x^{j'}\vert x_t=x^j,a_{it}=a\right]$; $\mathbf{M} = \left(I_{(K+1)^I}\otimes M\right)$, where $M=(I_J-L)^{-1}$ and L denotes a $J\times J$ matrix such that $(L)_{jj'} = \beta\Pr\left[x_{t+1}=x^{j'}\vert x_t=x^j\right]$ and I_d denotes an identity matrix of size d; $\mathbf{R} = \left(\iota_{(K+1)^I}\otimes R\right)$ is a $J(K+1)^I$ by $J(K+1)^I$ matrix, where ι_d denotes a $d-$column vector of ones, $R = \left[\begin{array}{ccc} P^{0\ldots0} & \cdots & P^{K\ldots K}\end{array}\right]$ is a J by $J(K+1)^I$ matrix so that $P^{a_1\ldots a_I} = diag(P(a_1,\ldots,a_I\vert x^1),\ldots,P\left(a_1,\ldots,a_I\vert x^J\right))$ with $P(a_1,\ldots,a_I\vert x) = \prod_{j=1}^I P_j(a_j\vert x)$; and $\mathbf{R}_i$ is a $J(K+1)$ by $J(K+1)^I$ matrix such that $\mathbf{R}_i\mathbf{\Pi}_i = \left[\begin{array}{ccc}(R_i^0\mathbf{\Pi}_i)^\top & \cdots & \left(R_i^K\mathbf{\Pi}_i\right)^\top\end{array}\right]^\top$ gives a $J(K+1)$ by p_i matrix with the first J rows is $R_i^0\mathbf{\Pi}_i = \left(E[\pi_{i0}(0,\mathbf{a}_{-it},x_t)\,\vert x_t=x^1],\ldots,E\left[\pi_{i0}(0,\mathbf{a}_{-it},x_t)\,\vert x_t=x^J\right]\right)^\top$, and the next J rows is $R_i^1\mathbf{\Pi}_i = \left(E[\pi_{i0}(1,\mathbf{a}_{-it},x_t)\,\vert x_t=x^1],\ldots,E\left[\pi_{i0}(1,\mathbf{a}_{-it},x_t)\,\vert x_t=x^J\right]\right)^\top$ and so on. Define $\Delta v^a_{i,\theta_i} = \left(v_{i,\theta_i}(a,x^1)-v_{i,\theta_i}(0,x^1),\ldots,v_{i,\theta_i}\left(a,x^J\right)-v_{i,\theta_i}\left(0,x^J\right)\right)$ for all $a>0$, and $\Delta\mathbf{v}_\theta = \left(\Delta v^1_{i,\theta_i},\ldots,\Delta v^K_{i,\theta_i}\right)^\top$. Then let $\mathbf{D}$ denote the $JK\times J(K+1)$ matrix that performs the transformation

$\mathbf{D}\mathbf{v}_\theta = \Delta\mathbf{v}_\theta$. Finally $\underline{v}_i$can be constructed similarly. Let $\underline{v}_i^a = \left(\underline{v}_i\left(a, x^1\right), \ldots, \underline{v}_i\left(a, x^J\right)\right)$ for all a, so that $\underline{v}_i^a = \beta H_i^a M \underline{e}_i$ where $\underline{e}_i = (E[\sum_{a'>0} \varepsilon_{it}(a') \mathbf{1}[a_{it} = a'] | x_t = x^1], \ldots, E[\sum_{a'>0} \varepsilon_{it}(a') \mathbf{1}[a_{it} = a'] | x_t = x^J])$. We define $\underline{\mathbf{v}}_i = \left(\underline{v}_i^0, \ldots, \underline{v}_i^K\right)^\top$, so that $\Delta\underline{\mathbf{v}}_i = \mathbf{D}\underline{\mathbf{v}}_i$ is also a $JK-$vector. Then the expression in equation (5) immediately follows.■

Construction of $\mathcal{X}_i$ used in the Simulation Study. Here we provide some explicit details of $\mathcal{X}_i$ for the game we have described in Section 4. We only show $\mathcal{X}_1$ to avoid repetition. $\mathcal{X}_2$ can be constructed similarly. From (8), note that the payoff function of player 1 satisfies M5:

$$\pi_{1,\theta}(a_{1t}, a_{2t}, x_t) = a_{1t} \cdot \mu_1 + a_{1t}a_{2t} \cdot \mu_2 + a_{1t}(1 - a_{1t-1}) \cdot F + (1 - a_{1t})\, a_{1t-1} \cdot W. \tag{9}$$

So we can write $\pi_{1,\theta}(a_{1t}, a_{2t}, x_t) = \theta^\top \pi_{10}(a_{1t}, a_{2t}, x_t)$ with $\theta = (\mu_1, \mu_2, F, W)^\top$ and $\pi_{10}(a_{1t}a_{2t}, x_t) = (a_{1t}, a_{1t}a_{2t}, a_{1t}(1 - a_{1t-1}), (1 - a_{1t})\, a_{1t-1})^\top$. Then, following equation (4) and its subsequent discussion, we have $\underline{v}_{1,\theta}(a, x) = \theta^\top \underline{v}_{10}(a, x)$ for some $4-$dimensional vector $\underline{v}_{10}(a, x)$ for any a, x. With two actions $\mathcal{X}_1$ is just a vectorization of $\{\underline{v}_{10}(1, x) - \underline{v}_{10}(0, x)\}_{x \in X}$. In terms of the notation used above we have: $(\mathbf{R}_1 + \beta\mathbf{H}_1\mathbf{M}\mathbf{R})\,\mathbf{\Pi}_1 = \begin{bmatrix} R_1^0\mathbf{\Pi}_1 + \beta H_1^0 (I_4 - L)^{-1} R\mathbf{\Pi}_1 \\ R_1^1\mathbf{\Pi}_1 + \beta H_1^1 (I_4 - L)^{-1} R\mathbf{\Pi}_1 \end{bmatrix} = \begin{bmatrix} \{\underline{v}_{10}(0, x)\}_{x \in X} \\ \{\underline{v}_{10}(1, x)\}_{x \in X} \end{bmatrix}$. We order the elements in the state vector according to $(a_{1t-1}, a_{2t-1}) = ((0,0), (0,1), (1,0), (1,1))^\top$. Then let

$p_i(x)$ stand for $\Pr[a_{it}=1|x_t=x]$ and $q_i(x)=1-p_i(x)$, we have:

$$\begin{aligned}
R_1^0\mathbf{\Pi}_1 &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \; R_1^1\mathbf{\Pi}_1 = \begin{bmatrix} 1 & p_2((0,0)) & 1 & 0 \\ 1 & p_2((0,1)) & 1 & 0 \\ 1 & p_2((1,0)) & 0 & 0 \\ 1 & p_2((1,1)) & 0 & 0 \end{bmatrix}, \\
R\mathbf{\Pi}_1 &= \begin{bmatrix} p_1((0,0)) & p_1((0,0))\,p_2((0,0)) & p_1((0,0)) & 0 \\ p_1((0,1)) & p_1((0,1))\,p_2((0,1)) & p_1((0,1)) & 0 \\ p_1((1,0)) & p_1((1,0))\,p_2((1,0)) & 0 & q_1((1,0)) \\ p_1((1,1)) & p_1((1,1))\,p_2((1,1)) & 0 & q_1((1,1)) \end{bmatrix}, \\
L &= \beta \begin{bmatrix} q_1((0,0))\,q_2((0,0)) & q_1((0,0))\,p_2((0,0)) & p_1((0,0))\,q_2((0,0)) & p_1((0,0))\,p_2((0,0)) \\ q_1((0,1))\,q_2((0,1)) & q_1((0,1))\,p_2((0,1)) & p_1((0,1))\,q_2((0,1)) & p_1((0,1))\,p_2((0,1)) \\ q_1((1,0))\,q_2((1,0)) & q_1((1,0))\,p_2((1,0)) & p_1((1,0))\,q_2((1,0)) & p_1((1,0))\,p_2((1,0)) \\ q_1((1,1))\,q_2((1,1)) & q_1((1,1))\,p_2((1,1)) & p_1((1,1))\,q_2((1,1)) & p_1((1,1))\,p_2((1,1)) \end{bmatrix}, \\
H_1^0 &= \begin{bmatrix} q_2((0,0)) & p_2((0,0)) & 0 & 0 \\ q_2((0,1)) & p_2((0,1)) & 0 & 0 \\ q_2((1,0)) & p_2((1,0)) & 0 & 0 \\ q_2((1,1)) & p_2((1,1)) & 0 & 0 \end{bmatrix} \text{ and } H_1^1 = \begin{bmatrix} 0 & 0 & q_2((0,0)) & p_2((0,0)) \\ 0 & 0 & q_2((0,1)) & p_2((0,1)) \\ 0 & 0 & q_2((1,0)) & p_2((1,0)) \\ 0 & 0 & q_2((1,1)) & p_2((1,1)) \end{bmatrix}.
\end{aligned}$$

We do not write out R_1^a, R and $\mathbf{\Pi}_1$ separately since they are cumbersome. (The number of columns of R_1^a and R, and the number of rows in $\mathbf{\Pi}_1$ are 2^4 that equals to the number of all distinct possibilities of $(a_{1t}, a_{2t}, a_{1t-1}, a_{2t-1})$.) Although it is obvious from (9) how the expressions for $R_1^a\mathbf{\Pi}_1$ and $R\mathbf{\Pi}_1$ respectively vectorize $\{E[\pi_0(a_{1t}, a_{2t}, x_{it})\,|x_t = x, a_{it} = a]\}_{x\in X}$ and $\{E[\pi_0(a_{1t}, a_{2t}, x_{it})\,|x_t = x]\}_{x\in X}$. The contents of H_1^a are simply conditional choice probabilities of player 2's action, and those in L are products of the choice probabilities of both players since their actions are conditionally independent. Then given the data the conditional choice probabilities can be estimated, and the sample counterpart of $\mathcal{X}_1$ can be constructed for the purpose of estimation.

Appendix B - Large Sample Properties

In what follows we denote the matrix norm by $\|\cdot\|$, so that $\|B\| = \sqrt{trace(B^\top B)}$ for any real matrix B, and we let "$\xrightarrow{p}$" and "$\xrightarrow{d}$" denote convergence in probability and distribution respectively. Suppose the following hold:

Assumption B1: *$\mathcal{X}$ has full column rank and $\mathcal{W}$ is p.d.*

Assumption B2: *$\|\mathcal{W}\|, \|\mathcal{X}\|$ and $\|\mathcal{Y}\|$ are finite, and $\widehat{\mathcal{W}} \xrightarrow{p} \mathcal{W}, \widehat{\mathcal{X}} \xrightarrow{p} \mathcal{X}$ and $\widehat{\mathcal{Y}} \xrightarrow{p} \mathcal{Y}$.*

Assumption B3: *Let $\widehat{\mathcal{U}} = \widehat{\mathcal{Y}} - \widehat{\mathcal{X}}\theta_0$, $\sqrt{N}\widehat{\mathcal{U}} \xrightarrow{d} \mathcal{N}(0, \Sigma)$ where Σ is p.d. and non-stochastic.*

B1 assumes θ_0 is the unique minimizer of $\mathcal{S}(\theta; \mathcal{W})$. When $\mathcal{W}$ is p.d., the full rank condition of $\mathcal{X}$ is necessary and sufficient condition for the identification of θ_0. Analogously, if $\widehat{\mathcal{W}}$ is p.d., $\widehat{\mathcal{X}}$ has full column rank if and only if $\widehat{\mathcal{S}}(\theta; \widehat{\mathcal{W}})$ has a unique solution (in (7)). B2 and B3 are standard high level conditions that can be verified under weak conditions since $(\widehat{\mathcal{X}}, \widehat{\mathcal{Y}})$ are smooth mappings of the choice and transition probabilities. Then:

Proposition 1(Consistency): *Under assumptions A1 - A2, $\widehat{\theta}(\widehat{\mathcal{W}}) \xrightarrow{p} \theta_0$.*

Proposition 2(Asymptotic Normality): *Under assumptions A1 - A3,*

$$\sqrt{N}(\widehat{\theta}(\widehat{\mathcal{W}}) - \theta_0) \xrightarrow{d} \mathcal{N}(0, \Omega_{\mathcal{W}}),$$

where $\Omega_{\mathcal{W}} = \left(\mathcal{X}^\top \mathcal{W} \mathcal{X}\right)^{-1} \mathcal{X}^\top \mathcal{W} \Sigma \mathcal{W} \mathcal{X} \left(\mathcal{X}^\top \mathcal{W} \mathcal{X}\right)^{-1}$. Furthermore, $\Omega_{\mathcal{W}} - \Omega_{\Sigma^{-1}}$ is positive semi-definite for any $\mathcal{W}$.

Note that efficient estimation requires a consistent estimator of Σ, which can be constructed using any preliminary consistent estimator of θ_0 such as $(\widehat{\mathcal{X}}^\top \widehat{\mathcal{X}})^{-1} \widehat{\mathcal{X}}^\top \widehat{\mathcal{Y}}$.

Proof of Proposition 1. Under A1 and A2 $\widehat{\mathcal{W}}\widehat{\mathcal{X}}$ has full column rank with probability approaching (w.p.a.) 1. Consistency immediately follows by repeated applications of continuous mapping theorem.■

Proof of Proposition 2. Using the definitions of the estimator in (7) and $\widehat{\mathcal{U}}$, we have w.p.a. 1:

$$\begin{aligned} \widehat{\theta} &= \theta_0 + (\widehat{\mathcal{X}}^\top \widehat{\mathcal{W}} \widehat{\mathcal{X}})^{-1} \widehat{\mathcal{X}}^\top \widehat{\mathcal{W}} \widehat{\mathcal{U}} \\ &= \theta_0 + (\mathcal{X}^\top \mathcal{W} \mathcal{X})^{-1} \mathcal{X}^\top \mathcal{W} \widehat{\mathcal{U}} + o_p(||\widehat{\mathcal{U}}||), \end{aligned}$$

where the second equality follows from continuous mapping theorem. Asymptotic normality follows from Assumption A3 and an application of Slutsky's theorem. The efficiency proof for this type of variance structure is well-known (e.g. see Hansen (1982, Theorem 3.2)).■

References

[1] Aguirregabiria, V. and P. Mira (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, **75**, 1-53.

[2] Aguirregabiria, V., and P. Mira (2010): “Dynamic Discrete Choice Structural Models: A Survey,” *Journal of Econometrics,* **156**, 38-67

[3] Aguirregabiria, V. and P. Mira (2013): “Identification of Games of Incomplete Information with Multiple Equilibria and Common Unobserved Heterogeneity,” *Working paper*, University of Toronto.

[4] Bajari, P. and H. Hong (2006): “Semiparametric Estimation of a Dynamic Game of Incomplete Information,” *NBER Technical Working Paper 320.*

[5] Bajari, P., C.L. Benkard, and J. Levin (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, **75**, 1331-1370.

[6] Bajari, P., V. Chernozhukov, H. Hong and D. Nekipelov (2009): “Identification and Efficient Estimation of a Dynamic Discrete Game,” *Working paper*, University of Minnesota.

[7] Bajari, P., H. Hong and D. Nekipelov (2012): *Econometrics for Game Theory*, Advances in Economics and Econometrics: Theory and Applications, 10th World Congress.

[8] Egesdal, M., Z. Lai and C. Su (2013): “Estimating Dynamic Discrete-Choice Games of Incomplete Information,” *Working Paper*, University of Chicago Booth School of Business.

[9] Hansen, L.P. (1982): “Large Sample Properties of Generalized Method of Moments Estimators,” *Econometrica*, **50**, 1029 -1054.

[10] Hotz, V., and R.A. Miller (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, **60**, 497-531.

[11] Hotz, V., R.A. Miller, S. Sanders and J. Smith (1994): “A Simulation Estimator for Dynamic Models of Discrete Choice,” *Review of Economic Studies*, **61**, 265-289.

[12] Hu, Y., and M. Shum (2012): “Nonparametric Identification of Dynamic Models with Unobserved State Variables,” *Journal of Econometrics*, **171**, 32-44.

[13] Kasahara, H. and K. Shimotsu (2009), "Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices," *Econometrica*, **77**, 135-175.

[14] Kasahara, H. and K. Shimotsu (2012), "Sequential Estimation of Structural Models with a Fixed Point Constraint," *Econometrica*, **80**, 2303-2319

[15] Pakes, A., M. Ostrovsky, and S. Berry (2007): "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Example)," *RAND Journal of Economics*, **38**, 373-399.

[16] Pesendorfer, M., and P. Schmidt-Dengler (2008): "Asymptotic Least Squares Estimator for Dynamic Games," *Review of Economics Studies*, **75**, 901-928.

[17] Rust, J. (1994): " Structural Estimation of Markov Decision Process," *Handbook of Econometrics*, vol. 4, eds. R. Engle and D. McFadden. North Holland.

[18] Srisuma, S. (2013): "Minimum Distance Estimators for Dynamic Games," *Quantitative Economics*, **4**, 549-583.

[19] Tamer, E. (2003), "Incomplete Simultaneous Discrete Response Model with Multiple Equilibria," *Review of Economic Studies*, **70**, 147–165.

[20] Xiao, R. (2014): "Identification and Estimation of Incomplete Information Games with Multiple Equilibria," *Working paper*, Johns Hopkins University.