



City Research Online

City St George's, University of London

Citation: Miessi Sanches, F.A., Silva Junior, D. & Srisuma, S. (2016). Ordinary Least Squares Estimation of a Dynamic Game Model. *International Economic Review*, 57(2), pp. 623-634. doi: 10.1111/iere.12170

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/16624/>

Link to published version: <https://doi.org/10.1111/iere.12170>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Ordinary Least Squares Estimation of a Dynamic Game Model*

Fabio A. Miessi Sanches[†]

Daniel Silva Junior[‡]

University of São Paulo

London School of Economics

Sorawoot Srisuma[§]

University of Surrey

July 28, 2014

Abstract

The estimation of dynamic games is known to be a numerically challenging task. In this paper we propose an alternative class of asymptotic least squares estimators to Pesendorfer and Schmidt-Dengler's (2008), which includes several well known estimators in the literature as special cases. Our estimator can be substantially easier to compute. In the leading case with linear payoffs specification our estimator has a familiar OLS/GLS closed-form that does not require any optimization. When payoffs have partially linear form, we propose a sequential estimator where the parameters in the nonlinear term can be estimated independently of the linear components, the latter can then be obtained in closed-form. We show the class of estimators we propose and Pesendorfer and Schmidt-Dengler's are in fact asymptotically equivalent. Hence there is no theoretical cost in reducing the computational burden. 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.

*We are grateful to Martin Pesendorfer for encouragement and support. We thank Joachim Groeger, Oliver Linton, Robert Miller, Pasquale Schiraldi, Richard Smith and Dimitri Szerman for helpful 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 estimation problem for a class of dynamic games of incomplete information that generalizes the single agent discrete Markov decision models surveyed in Rust (1994); for a recent survey see Aguirregabiria and Mira (2010). The setup is in an infinite time horizon, where players' private values enter the payoff function additively and are independent across players, under the conditional independence framework. A Markov equilibrium of such game can be represented by a fixed point of nonlinear equations in the space of choice probabilities and has been shown to exist (e.g. see Aguirregabiria and Mira (2007) and Pesendorfer and Schmidt-Dengler (2008)). A variety of methods have been proposed by different authors to estimate the same class of games based on the equilibrium condition in recent years; examples are given below. However, a common component of these methodologies 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 paper we propose a class of asymptotic least squares estimators constructed based on the equilibrium condition of the game when represented in the space of payoffs. Our work is motivated by the well-received methodology developed in Pesendorfer and Schmidt-Dengler (2008), who propose an efficient estimator for a unifying class of estimators that includes the non-iterative pseudo-likelihood estimator of Aguirregabiria and Mira (2007) and the moment based estimators discussed in Pakes, Ostrovsky and Berry (2007) as special cases. In contrast to our work, Pesendorfer and Schmidt-Dengler use the choice probability representation of the equilibrium to construct their estimator. Our goal is to show there is much to gain computationally using our approach with no efficiency lost. Henceforth we use the abbreviation $ALSE_{PSD}$ when referring to a generic estimator of Pesendorfer and Schmidt-Dengler.

We claim our estimator can be substantially easier to compute than $ALSE_{PSD}$. In the leading case our estimator has a familiar OLS/GLS closed-form expression when the per-period payoff function takes a linear-in-parameter specification.¹ In an intermediate case when the payoff function has an additive partially linear form, Frisch-Waugh-Lovell theorem can be applied so the parameters in the nonlinear part can be estimated first (dimensional reduction), and the linear-in-parameter component

¹The linear payoffs structure may seem restrictive, but it is in fact quite general as it includes any nonlinear (basis) functions of observables; albeit perhaps with an atheoretic flavor. However, linear specification arises naturally in many applications, and/or does not cause much concern in terms of structural interpretability in other situations. A leading example for the latter is when the goal of an empirical analysis is to study market outcomes, such as competition study of market power. Some notable recent empirical applications of linear-in-parameter payoffs include Aguirregabiria and Mira (2007), Ryan (2012) and Collard-Wexler (2013).

can be obtained in closed-form in the second step.² Even in a more general nonlinear case, we argue that our estimator is still generally easier to compute than ALSE_{PSD} . ALSE_{PSD} also provides a good benchmark for a comparison with other estimators in the literature as it has a well-defined efficiency property. We establish a duality between our estimator and ALSE_{PSD} , in the sense that they can always be constructed to have the same asymptotic distribution. Therefore our efficient estimator is as efficient as the efficient ALSE_{PSD} .

The large sample properties of our estimator (and for asymptotic least squares generally) are easy to derive for discrete games. Technically, our estimation problem is a least squares problem with generated regressors and regressands, which are generally smooth functions of the finite dimensional first stage parameters that are nonparametrically identified. In addition, the number of square terms in the objective function does not grow with sample size but is determined by the cardinality of the action and state spaces. Therefore our estimator belongs to the class of asymptotic least squares estimators as defined in Gourieroux and Monfort (1985,1995) in the same sense as ALSE_{PSD} . The close connection between our estimator and ALSE_{PSD} goes even further given the smooth bijective relation between normalized expected payoffs and choice probabilities (Hotz and Miller (1993)'s inversion); ALSE_{PSD} is defined to minimize the distance between the probabilities implied by the pseudo-model and the data. We show that, locally around the true, using the inverse function theorem, our estimator can be constructed to have the same asymptotic distribution as any ALSE_{PSD} by choosing an appropriate weighting matrix and vice versa.

There are at least two reasons why the estimation of dynamic games can be non-trivial. First, as well-known from the single-agent problem, it involves value functions that generally do not have closed-form and need to be numerically evaluated so it is computationally demanding (see Rust (1996)). For games, there is also a potential issue of indeterminacy of multiple equilibria that gives rise to incomplete models (Tamer (2003)). A novel approach popularized by Hotz and Miller (1993) performs inference on the pseudo-model, generated from the observed data, by estimating the (policy) value functions that can significantly simplify the computational aspect. Pseudo-models are also generally easier to handle in a strategic environment as they have been shown to be complete for several classes of games (Srisuma (2013b)). Methodologies based on pseudo-models are often referred to as two-step estimators since they require estimation of value functions in the first stage. Many recently proposed estimators for dynamic games are two-step estimators.

However, despite the simplification of two-step methods, the numerical aspects for implementing existing estimators in the literature appear to remain a concern as they generally involve solving highly nonlinear optimization problems. It is not uncommon to see methodology papers using esti-

²Modeling of additive linear components in the payoffs often appear in games with entry/exit decisions, as fixed cost or scrap value, or more generally as fixed effects.

mation time, amongst other things, as a competing factor. Furthermore, it is also not unusual that the choice of players' per-period payoff specification is chosen with the ease of numerical implementation in mind. In particular there can be substantial benefits (in terms of computational time) in specifying player's payoff functions to be linear-in-parameters. As the action-specific expected payoffs can then be written as a linear transformation of the parameter, following from the linear structure that defines the expected payoffs using stationary Markovian beliefs; examples of such discussions can be found in Bajari, Benkard and Levin (2007, Section 3.3.1) and Pakes, Ostrovsky, Berry (2007, Section 3). As a result, a linear parameterization of the payoffs is a leading specification employed in empirical work (see Footnote 1 for examples).

The objective functions that are used to define many two-step estimators in the literature are constructed in terms of choice probabilities implied by the pseudo-model. These probabilities can be motivated by the equilibrium condition of the game, which can be stated in terms of consistent beliefs with probabilities of best responses. Choice probabilities are used to define traditional criterion functions such as pseudo-likelihood function (Aguirregabiria and Mira (2007), Kasahara and Shimotsu (2012)) or moment and minimum distance based conditions (Pakes, Ostrovsky, Berry (2007), Pesendorfer and Schmidt-Dengler (2008)). However, in order to calculate the probabilities implied by the pseudo-model, one must first compute the expected discounted payoffs that determine the region of integration to be integrated to compute the probabilities. Furthermore, the integral is generally a nonlinear map of the expected payoffs, and it typically has to be computed numerically outside the well-known conditional logit framework. The integral, following Hotz and Miller (1993)'s inversion result, in fact represents a one-to-one mapping between the probabilities and the normalized expected payoffs.

There are also other methodologies that use expected payoffs explicitly to define their objective functions. 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. 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 payoffs. First is the two-step estimator proposed by Bajari, Benkard and Levin (2007), who also use forward simulation like Hotz et al. However, generally no closed-form estimator is possible with Bajari, Benkard and Levin's methodology as they compare expected payoffs in the pseudo-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.^{3,4}

The rest of the paper is organized as follows. Section 2 begins with an illustrative example that motivates our estimator, and then describes the model and our estimator for games. Section 3 gives the main results. Section 4 presents results from Monte Carlo experiments that compare the statistical performance and relative speed of our estimator and ALSE_{PSD} . Section 5 concludes and provides a brief discussion on how our estimators can be adapted or applied to complement other recent results in the literature. All proofs can be found in the Appendix.

2 Methodology

We begin with an illustration that highlights the idea behind computational advantages of our estimation approach. Section 2.2 describes elements of the game. We define the pseudo-model in Section 2.3 and introduce our estimator in Section 2.4.

2.1 Least Squares in Probabilities vs Payoffs

Consider a model generated by the following binary choice variable:

$$a_t(\theta) = \mathbf{1}[v_\theta(x_t) \leq \varepsilon_t] \quad \text{for } \theta \in \Theta \subset \mathbb{R}^p,$$

where x_t and ε_t are independent. Let the cdf of ε_t be denoted by Q . For all x , let $P_\theta(x) = \Pr[a_t(\theta) = 1 | x_t = x]$, so that $P_\theta(x) = Q(v_\theta(x))$. Assume the support of x_t is finite, say $\{x^j\}_{j=1}^J$ for some $J < \infty$, so that we can define $\mathbf{P}_\theta = \Gamma(\mathbf{v}_\theta)$, where $\mathbf{P}_\theta = (P_\theta(x^1), \dots, P_\theta(x^J))^\top$, $\mathbf{v}_\theta = (v_\theta(x^1), \dots, v_\theta(x^J))^\top$ and $\Gamma(\mathbf{v}_\theta) = (Q(v_\theta(x^1)), \dots, Q(v_\theta(x^J)))^\top$.

Suppose: we observe a random sample of $\{a_t, x_t\}$ where $a_t = a_t(\theta_0)$ for some $\theta_0 \in \Theta$, which is the parameter value of interest; v_θ is nonparametrically identified up to θ , and there exists a consistent estimator of $\mathbf{v}_\theta$, say $\hat{\mathbf{v}}_\theta$, for all θ ; and, Q is known and invertible. Let $\mathbf{P} = (P(x^1), \dots, P(x^J))^\top$ be a vector of choice probabilities identified from the data, so that $\mathbf{P} = \mathbf{P}_{\theta_0}$, then one may consider a class of estimators defined by

$$\hat{\theta}_p(\mathcal{V}) = \arg \min_{\theta \in \Theta} \left(\tilde{\mathbf{P}} - \hat{\mathbf{P}}_\theta \right)^\top \mathcal{V} \left(\tilde{\mathbf{P}} - \hat{\mathbf{P}}_\theta \right), \quad (1)$$

where $\tilde{\mathbf{P}}$ and $\hat{\mathbf{P}}_\theta$ are estimators for $\mathbf{P}$ and $\mathbf{P}_\theta$ respectively, and $\mathcal{V}$ be some positive definite matrix. Note that $\tilde{\mathbf{P}}$ and $\hat{\mathbf{P}}_{\theta_0}$ are generally different since the former is model-free while the latter is estimated

³An 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).

⁴Another notable estimator that does not take a two-step approach is Egesdal, Lai and Su (2012). However, Egesdal et al. construct their objective functions in terms of choice probabilities.

through $\widehat{\mathbf{v}}_\theta$. Similarly, we can define $\mathbf{v} = (Q^{-1}(P(x^1)), \dots, Q^{-1}(P(x^J)))^\top$, which is also identified from the data, so that $\mathbf{v} = \mathbf{v}_{\theta_0}$ by construction. Then one can also consider an alternative class of estimators:

$$\widehat{\theta}_v(\mathcal{W}) = \arg \min_{\theta \in \Theta} (\widetilde{\mathbf{v}} - \widehat{\mathbf{v}}_\theta)^\top \mathcal{W} (\widetilde{\mathbf{v}} - \widehat{\mathbf{v}}_\theta), \quad (2)$$

where $\widetilde{\mathbf{v}}$ is $\Gamma^{-1}(\widetilde{\mathbf{P}})$ and $\mathcal{W}$ is a positive definite matrix. As described previously, $\widetilde{\mathbf{v}}$ and $\widehat{\mathbf{v}}_{\theta_0}$ will also generally differ.

Equations (1) and (2) provide two different estimators for θ_0 . We argue the latter should generally be easier to compute than the former since it is more convenient to compute $(\widetilde{\mathbf{v}}, \widehat{\mathbf{v}}_\theta)$ relative $(\widetilde{\mathbf{P}}, \widehat{\mathbf{P}}_\theta)$ across different values of θ . This argument is most transparent when v_θ has a linear-in-parameter specification, i.e. $v_\theta(x_t) = \theta^\top v(x_t)$ for some p -dimensional vector $v(x_t)$. Then $\widehat{\mathbf{v}}_\theta$ can be written as $\widehat{\mathbf{X}}\theta$, where $\widehat{\mathbf{X}}$ is a J by p matrix such that its j -th row equals $\widehat{v}(x^j)^\top$. The solution to (2) is unique and has a closed-form, $(\widehat{\mathbf{X}}^\top \mathcal{W} \widehat{\mathbf{X}})^{-1} \widehat{\mathbf{X}}^\top \mathcal{W} \widetilde{\mathbf{v}}$, when $\widehat{\mathbf{X}}^\top \mathcal{W} \widehat{\mathbf{X}}$ is invertible. Even without the linear parameterization of v_θ , every evaluation of $\widehat{\mathbf{P}}_\theta$ requires the mapping of $v_\theta(x^j)$ by Q for all j , for every θ , where Q is generally a nonlinear function that may have to be computed numerically. In contrast, for (2), the potentially costly step of applying Q^{-1} has to be performed only once to estimate $\mathbf{v}$ that does not depend on θ . Regardless of the parameterization in v_θ , under some suitable regularity conditions, and appropriate choices of weighting matrices, the two estimators can be shown to be asymptotically equivalent near θ_0 in the sense that there exists $\mathcal{W}_\mathcal{V}$ and $\mathcal{V}_\mathcal{W}$ such that for any $\mathcal{V}$ and $\mathcal{W}$:

$$\begin{aligned} \sqrt{N} \left(\widehat{\theta}_v(\mathcal{W}_\mathcal{V}) - \theta_0 \right) &= \sqrt{N} \left(\widehat{\theta}_p(\mathcal{V}) - \theta_0 \right) + o_p(1), \\ \sqrt{N} \left(\widehat{\theta}_p(\mathcal{V}_\mathcal{W}) - \theta_0 \right) &= \sqrt{N} \left(\widehat{\theta}_v(\mathcal{W}) - \theta_0 \right) + o_p(1), \end{aligned}$$

where N denotes the sample size.

The estimator in (1) is closely related to ALSE_{PSD} and other Hotz and Miller (1993)'s type estimators that have been widely adopted in the dynamic game setting. In contrast the estimator based on (2) is the asymptotic least squares analog to the estimator proposed in Hotz et al. (1994). For the remainder of this section we develop an estimator based on (2) in the context of a dynamic game.

2.2 Framework

We consider a game with I players, indexed by $i \in \mathcal{I} = \{1, \dots, I\}$, over an infinite time horizon. The elements of the game in each period are as follows:

ACTIONS. For notational simplicity we assume all players have the same action space. The action set of each player is $A = \{0, 1, \dots, K+1\}$. We denote the action variable for player i by

a_{it} . Let $\mathbf{a}_t = (a_{1t}, \dots, a_{It}) \in \mathbf{A} = \times_{i=1}^I A$. We will also occasionally abuse the notation and write $\mathbf{a}_t = (a_{it}, \mathbf{a}_{-it})$ where $\mathbf{a}_{-it} = (a_{1t}, \dots, a_{i-1t}, a_{i+1t}, \dots, a_{It}) \in \mathbf{A} \setminus A$.

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 $\varepsilon_{it} \in \mathcal{E} = \mathbb{R}^{K+1}$ denotes private information only observed by player i . Note that common state space X is without any loss of generality. We shall use s_{it} and $(x_{it}, \varepsilon_{it})$ interchangeably. We define $(\mathbf{s}_t, \mathbf{s}_{-it}, \boldsymbol{\varepsilon}_t, \boldsymbol{\varepsilon}_{-it}, \mathcal{E})$ analogously to $(\mathbf{a}_t, \mathbf{a}_{-it}, A)$, and denote the support of $\mathbf{s}_t$ by $S = X \times \mathcal{E}$.

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, $u_i : A \times S \rightarrow \mathbb{R}$, which is time separable. The payoff function for player i can depend generally on $(\mathbf{a}_t, x_t, \varepsilon_{it})$ but not directly on $\boldsymbol{\varepsilon}_{-it}$.

DISCOUNTING FACTOR. Future period's payoffs are discounted at the rate $\beta_i \in (0, 1)$ for each player. For notational simplicity we take $\beta_i = \beta$ for all i .

We impose the following assumptions throughout the paper.

ASSUMPTION M1 (Additive Separability). $u_{i,\theta_i}(a_i, \mathbf{a}_{-i}, x, \varepsilon_i) = \pi_{i,\theta_i}(a_i, \mathbf{a}_{-i}, x) + \sum_{a' \in A} \varepsilon_i(a') \mathbf{1}[a_i = a']$ for all $i, \theta_i, a_i, \mathbf{a}_{-i}, x, \varepsilon_i$, where π_{i,θ_i} is known up to $\theta_i \in \Theta_i \subset \mathbb{R}^{p_i}$.

ASSUMPTION M2 (Conditional independence). The transitional distribution of the states has the following factorization: $P(x_{t+1}, \boldsymbol{\varepsilon}_{t+1} | x_t, \boldsymbol{\varepsilon}_t, \mathbf{a}_t) = Q(\boldsymbol{\varepsilon}_{t+1}) G(x_{t+1} | x_t, \mathbf{a}_t)$, where Q is the cumulative distribution function of $\boldsymbol{\varepsilon}_t$ 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 $\mathbb{R}^{K+1}$. So that $Q(\boldsymbol{\varepsilon}) = \prod_{i=1}^I Q_i(\varepsilon_i)$, where Q_i denotes the cumulative distribution function of ε_{it} .

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

M1 - M4 are standard in the modeling of dynamic discrete games in the literature. Note that M2 implies x_t and $\boldsymbol{\varepsilon}_t$ are independent, however, this can be relaxed slightly at the cost of more notation by changing all of our statements regarding Q and Q_i to be taken conditional on x_t . M4 is also

not essential for the general idea behind estimation of dynamic games. Although the complexity of the asymptotic theory and the practical aspects increase significantly when x_t includes continuous random variables; see Bajari et al. (2009) and Srisuma and Linton (2012).

At time t every player observes s_{it} , each then chooses a_{it} simultaneously. We consider a Markovian framework where players' behaviors are stationary across time and players are assumed to play pure strategies. More specifically, for some $\alpha_i : S \rightarrow A$, $a_{it} = \alpha_i(s_{it})$ for all i, t , so that whenever $s_{it} = s_{i\tau}$ then $\alpha_i(s_{it}) = \alpha_i(s_{i\tau})$ for any τ . The beliefs are also time invariant. Player i 's beliefs, σ_i , is a distribution of $\mathbf{a}_t = (\alpha_1(s_{1t}), \dots, \alpha_I(s_{It}))$ conditional on x_t for some pure Markov strategy profile $(\alpha_1, \dots, \alpha_I)$. The decision problem for each player is to solve

$$\max_{a_i \in A_i} \{E_{\sigma_i}[u_{i,\theta_i}(a_{it}, \mathbf{a}_{-it}, s_i) | s_{it} = s_i, a_{it} = a_i] + \beta E_{\sigma_i}[W_{i,\theta_i}(s_{it+1}; \sigma_i) | s_{it} = s_i, a_{it} = a_i]\}, \quad (3)$$

$$\text{where } W_{i,\theta_i}(s_i; \sigma_i) = \sum_{\tau=t}^{\infty} \beta^{\tau-t} E_{\sigma_i}[u_{i,\theta_i}(\mathbf{a}_\tau, s_{i\tau}) | s_{it} = s_i],$$

for any s_i . 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 ; in particular, player i forms an expectation for all players' future actions including herself, and today's actions of opposing players. $W_{i,\theta_i}(\cdot; \sigma_i)$ is a policy value function since the expected discounted return needs not be an optimal value from an optimization problem since σ_i can be any beliefs, not necessarily equilibrium beliefs. Note that the transition laws for future states are completely determined by the primitives and the beliefs. Any strategy profile that solves the decision problems for all i and is consistent with the beliefs satisfies 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 (e.g. see Aguirregabiria and Mira (2007) and Pesendorfer and Schmidt-Dengler (2008)). However, there may be multiple equilibria.

2.3 Pseudo-Model

We now define the pseudo-model that plays a central role in two-step estimation methods. The starting point is the structural assumption that we observe random sample of $\{\alpha_1^*(s_{1t}), \dots, \alpha_I^*(s_{It}), x_t, x_{t+1}\}$ from a single equilibrium, where $\alpha_i^* = \alpha_{i,\theta_{i0}}$ for some $\theta_{i0} \in \Theta_i \subset \mathbb{R}^{p_i}$ for all i . Let $P_i^*(a_i|x) = \Pr[\alpha_i^*(s_{it}) = a_i | x_t = x]$ for all a_i, x . Then we have: (i) the equilibrium beliefs for all players is summarized by $\prod_{i=1}^I P_i^*$; (ii) $\Pr[a_{it} = a_i | x_t = x] = P_i^*(a_i|x)$ and $\Pr[x_{t+1} = x' | x_t = x, a_t = a] = G(x'|x, a)$ for all a, x, x' . For notational simplicity, for this section and the next, we shall: omit $*$; let α_i and P_i denote the equilibrium strategy and choice probability function for player i ; and, without any ambiguity let $a_{it} = \alpha_i(s_{it})$ for all i, t . Then the pseudo-model can be defined as a collection of joint

conditional distributions indexed by $\theta = (\theta_1^\top, \dots, \theta_I^\top)^\top \in \times_{i=1}^I \Theta_i = \Theta \subset \mathbb{R}^p$. Also let θ_0 denote $(\theta_{10}^\top, \dots, \theta_{I0}^\top)^\top$.

DEFINITION: The pseudo-model is $\{P_\theta\}_{\theta \in \Theta}$ such that $P_\theta = \prod_{i=1}^I P_{i,\theta_i}$ and for all i, θ_i, a_i, x :

$$\begin{aligned} P_{i,\theta_i}(a|x) &= \Pr[\alpha_{i,\theta_i}(s_{it}) = a | x_t = x] \quad \text{a.s., where} \\ \alpha_{i,\theta_i}(s_{it}) &= \arg \max_{a_i \in A} \{E[\pi_{i,\theta_i}(a_i, a_{-it}, x_t) | x_t] + \varepsilon_{it}(a_i) + \beta E[V_{i,\theta_i}(s_{t+1}) | x_t, a_{it} = a_i]\}, \\ V_{i,\theta_i}(s_{it}) &= E[\pi_{i,\theta_i}(a_{it}, a_{-it}, x_t) + \sum_{a'=0}^K \varepsilon_{it}(a') \mathbf{1}[a_{it} = a'] | s_{it}] + \beta E[V_{i,\theta_i}(s_{it+1}) | s_{it}]. \end{aligned}$$

By construction $P_{i,\theta_i} = P_i$ for all i when $\theta_i = \theta_{i0}$ for all i , and V_{i,θ_i} also equals $W_{i,\theta_i}(\cdot; \sigma_i)$ (as defined in (3)), when $\sigma_i = \prod_{j=1}^I P_j$. Let $v_{i,\theta_i}(a_i, x) = E[\pi_{i,\theta_i}(a_i, a_{-it}, x_t) | x_t = x] + \beta E[V_{i,\theta_i}(s_{t+1}) | x_t = x, a_{it} = a_i]$ then we can write

$$P_{i,\theta_i}(a|x) = \Pr[v_{i,\theta_i}(a_i, x_t) + \varepsilon_{it}(a_i) > v_{i,\theta_i}(a'_i, x_t) + \varepsilon_{it}(a'_i) \text{ for all } a'_i \neq a_i | x_t = x], \quad (4)$$

which is familiar from the classical random utility model (e.g. see McFadden (1974)) with mean utility v_{i,θ_i} . The numerical advantage in working with the *pseudo*-model, as opposed to the *actual* model, is that v_{i,θ_i} is relatively straightforward to compute for different θ_i , since all expectations that define v_{i,θ_i} are calculated independent of θ_i ; all with respect to $P(s_{t+1} | s_t, a_t)$ for all players that is equivalent to earlier notation using E_{σ_i} when $\sigma_i = \prod_{j=1}^I P_j$ for all i .

We shall heavily exploit the fact that v_{i,θ_i} is a linear transformation of π_{i,θ_i} . To see this, first look at the choice-specific expected return:

$$\begin{aligned} E[V_{i,\theta_i}(s_{t+1}) | x_t, a_{it} = a_i] &= E[E[V_{i,\theta_i}(s_{t+1}) | x_{t+1}] | x_t, a_{it} = a_i], \text{ and} \\ E[V_{i,\theta_i}(s_t) | x_t] &= E[\pi_{i,\theta_i}(a_{it}, a_{-it}, x_t) + \sum_{a'=0}^K \varepsilon_{it}(a') \mathbf{1}[a_{it} = a'] | x_t] + \beta E[E[V_{i,\theta_i}(s_{t+1}) | x_{t+1}] | x_t]. \end{aligned}$$

Let $m_{i,\theta_i} = E[V_{i,\theta_i}(s_{it}) | x_t = \cdot]$ and $g_{i,\theta_i} = E[V_{i,\theta_i}(s_{it+1}) | x_t = \cdot, a_{it} = \cdot]$. Then, using a linear functional notation, we have

$$\begin{aligned} g_{i,\theta_i} &= \mathcal{H}_i m_{i,\theta_i}, \\ m_{i,\theta_i} &= r_{i,\theta_i} + \underline{r}_i + \mathcal{L} m_{i,\theta_i}, \text{ where for all } a, x \\ r_{i,\theta_i}(x) &= E[\pi_{i,\theta_i}(a_{it}, a_{-it}, x_t) | x_t = x], \\ \underline{r}_i(x) &= E[\sum_{a'=0}^K \varepsilon_{it}(a') \mathbf{1}[a_{it} = a'] | x_t = x], \\ \mathcal{L} m(x) &= \beta E[m(x_{t+1}) | x_t = x], \\ \mathcal{H}_i m(a, x) &= E[m(x_{t+1}) | x_t = x, a_{it} = a], \end{aligned}$$

where $\mathcal{L}$ and $\mathcal{H}_i$ are linear maps and r_{i,θ_i} is a linear transformation of π_{i,θ_i} . Since $(I - \mathcal{L})^{-1}$ is also generally a well-defined linear map, as $\mathcal{L}$ is a contraction as its norm is strictly less than 1, then

$$v_{i,\theta_i} = (\mathcal{R}_i + \beta \mathcal{H}_i (I - \mathcal{L})^{-1} \mathcal{R}) \pi_{i,\theta_i} + \underline{v}_i,$$

where $\mathcal{R}_i$ and $\mathcal{R}$ are conditional expectation operators, conditioning on x_t , integrating over a_{-it} and a_t respectively, and $\underline{v}_i = \beta \mathcal{H}_i (I - \mathcal{L})^{-1} \underline{r}_i$.

The 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 $a_i > 0$, then (4) becomes

$$P_{i,\theta_i}(a|x) = \Pr[\Delta v_{i,\theta_i}(a_i, x_t) + \varepsilon_{it}(a_i) > \Delta v_{i,\theta_i}(a'_i, x_t) + \varepsilon_{it}(a'_i) \text{ for all } a'_i > 0 | x_t = x]. \quad (5)$$

Since A and X are finite, the relationship between $\{\Delta v_{i,\theta_i}(a_i, x)\}_{a_i > 0, x \in X}$ and $\{\pi_{i,\theta_i}(\mathbf{a}, x)\}_{\mathbf{a} \in \mathbf{A}, x \in X}$ can be represented through a matrix equation. We state this representation as a lemma.

LEMMA R: Under M1 - M4 $\{\Delta v_{i,\theta_i}(a_i, x)\}_{a_i > 0, x \in X}$ can then be represented by a JK -vector, $\Delta \mathbf{v}_{i,\theta_i}$:

$$\Delta \mathbf{v}_{i,\theta_i} = \mathbf{D}(\mathbf{R}_i + \beta \mathbf{H}_i \mathbf{M} \mathbf{R}) \pi_{i,\theta_i} + \Delta \underline{\mathbf{v}}_i, \quad (6)$$

where π_{i,θ_i} is a $J(K+1)^I$ -vector of $\{\pi_{i,\theta_i}(a, x)\}_{a \in A, x \in X}$ so that elements in: $\mathbf{R}_i \pi_{i,\theta_i}$ are $\{E[\pi_{i,\theta_i}(a_i, a_{-it}, x_t) | x_t = x]\}_{a_i \in A, x \in X}$; $\mathbf{R} \pi_{i,\theta_i}$ are $\{E[\pi_{i,\theta_i}(a_{it}, a_{-it}, x_t) | x_t = x]\}_{x \in X}$; $\mathbf{M}$ involve $\{\Pr[x_{t+1} = x' | x_t = x]\}$; $\mathbf{H}_i$ are $\Pr[x_{t+1} = x' | x_t = x, a_{it} = a_i]$; and, $\mathbf{D}$ is a difference matrix with respect to the expected payoffs from playing action 0; and, $\Delta \underline{\mathbf{v}}_i$ is the differenced vector form of the transformation of $\underline{r}_i$ by $\beta_i \mathcal{H}_i (I - \mathcal{L})^{-1}$ normalized by action 0. The detailed constructions of $\Delta \underline{\mathbf{v}}_i$, $\mathbf{D}$, $\mathbf{R}_i$, $\mathbf{R}$, $\mathbf{H}_i$ and $\mathbf{M}$ are provided in the Appendix.

In what follows, we let $\Delta \mathbf{v}_i$ denote $\Delta \mathbf{v}_{i,\theta_{i0}}$. And, similarly, it shall be convenient to vectorize the probabilities. In particular, we let $\mathbf{P}_{i,\theta_i}$ and $\mathbf{P}_i$ denote the JK -vector that represent $\{P_{i,\theta_i}(a_i|x)\}_{a_i > 0, x \in X}$ and $\{P_i(a_i|x)\}_{a_i > 0, x \in X}$ respectively.

2.4 Estimation

Many objective functions proposed in the literature often can be written directly in terms of the probabilities from the pseudo-model, such as pseudo-likelihood and GMM, based on the construction that $\mathbf{P}_{i,\theta_i}$ coincides with $\mathbf{P}_i$ when $\theta_i = \theta_{i0}$. However, from a numerical perspective, computing the pseudo-probabilities requires a costly additional step of computation, namely the integration with respect to the distribution of ε_{it} that maps $\Delta \mathbf{v}_{i,\theta_i}$ into $\mathbf{P}_{i,\theta_i}$ (see (5)). These integrals generally do not have closed-form in the expected payoffs outside the well-known exception when private values are i.i.d. extreme value. Even if the integrals have closed-form, the integration is generally a nonlinear

mapping of $\Delta \mathbf{v}_{i,\theta_i}$ into $\mathbf{P}_{i,\theta_i}$. In order to preserve the linear structure outlined previously, we propose to construct objective functions based directly on $\Delta \mathbf{v}_{i,\theta_i}$.

The validity of such objective functions, to identify θ_0 , follows from the bijective relation between $\Delta \mathbf{v}_{i,\theta_i}$ and $\mathbf{P}_{i,\theta_i}$ for each i . This well-known result follows from Proposition 1 of Hotz and Miller (1993), which we shall refer to as Hotz and Miller's inversion in this paper (also see Lemma 8 of Matzkin (1991), Lemma 1 of Pesendorfer and Schmidt-Dengler (2008), and, for a recent generalization of these results, Norets and Takahashi (2013)).⁵ In particular, it immediately follows that for any θ_i , $\mathbf{P}_{i,\theta_i}$ coincides with $\mathbf{P}_i$ if and only if $\Delta \mathbf{v}_{i,\theta_i}$ coincides with $\Delta \mathbf{v}_i$, where $\Delta \mathbf{v}_i$ is identifiable from the data by Hotz and Miller's inversion. Then we can construct a class of estimators based on minimizing the distance between $\{\Delta \mathbf{v}_{i,\theta_i}\}_{i=1}^I$ and $\{\Delta \mathbf{v}_i\}_{i=1}^I$.

Using Lemma R, we can write $\Delta \mathbf{v}_{i,\theta_i} = \mathcal{X}_i(\theta_i) + \Delta \underline{\mathbf{v}}_i$, where

$$\mathcal{X}_i(\theta_i) = \mathbf{D}(\mathbf{R}_i + \beta \mathbf{H}_i \mathbf{M} \mathbf{R}) \boldsymbol{\pi}_{i,\theta_i}. \quad (7)$$

Note that θ_i enters $\mathcal{X}_i(\theta_i)$ through a matrix transform of the vector $\boldsymbol{\pi}_{i,\theta_i}$, where the former does not depend on θ_i and the latter is completely known and specified by the researcher. By Hotz and Miller's inversion, we also have $\Delta \mathbf{v}_i = \Phi_i(\mathbf{P}_i)$ for some nonlinear, but known, function Φ_i that only depends on the distributional assumption of ε_{it} . Then we can define a JK -vector, $\mathcal{Y}_i$, where

$$\mathcal{Y}_i = \Phi_i(\mathbf{P}_i) - \Delta \underline{\mathbf{v}}_i. \quad (8)$$

Note that $\mathcal{Y}_i$ is defined independently of θ_i . So that, by construction:

$$\mathcal{Y}_i = \mathcal{X}_i(\theta_i) \quad \text{when } \theta_i = \theta_{i0}.$$

Let $\mathcal{Y} = (\mathcal{Y}_1^\top, \dots, \mathcal{Y}_I^\top)^\top$, $\theta = (\theta_1^\top, \dots, \theta_I^\top)^\top$ and define a block diagonal matrix $\mathcal{X}(\theta) = \text{diag}(\mathcal{X}_1(\theta_1), \dots, \mathcal{X}_I(\theta_I))$. In the next section we analyze the asymptotic properties for a class of estimators that are motivated from minimizing

$$\mathcal{S}(\theta; \mathcal{W}) = (\mathcal{Y} - \mathcal{X}(\theta))^\top \mathcal{W} (\mathcal{Y} - \mathcal{X}(\theta)), \quad (9)$$

over Θ , for some weighting matrix $\mathcal{W}$.

It is also worth emphasizing that, through $\{\Delta \underline{\mathbf{v}}_i\}_{i=1}^I$, $\{\mathbf{R}_i\}_{i=1}^I$, $\mathbf{R}$, $\mathbf{L}$ and $\{\mathbf{H}_i\}_{i=1}^I$, for any θ : $\mathcal{X}(\theta)$ and $\mathcal{Y}$ are explicit functions, say $\mathcal{T}_{\mathcal{X}}(\theta; \gamma_0)$ and $\mathcal{T}_{\mathcal{Y}}(\gamma_0)$ respectively, of a finite-dimensional vector, γ_0 , that consists of choice and transition probabilities. However, optimization with $\mathcal{S}(\theta; \mathcal{W})$ is infeasible since $\mathcal{X}(\theta)$ and $\mathcal{Y}$ are not observed, as γ_0 is unknown. Given a sample from a single equilibrium,

⁵Pesendorfer and Schmidt-Dengler (2008) also show equilibrium condition can be characterized in terms of expected payoffs; see details of their Lemma 1 for further discussions.

$\{\alpha_1^*(s_{1t}), \dots, \alpha_I^*(s_{It}), x_t, x_{t+1}\}, \gamma_0$ can be identified from the data under weak conditions, hence $\mathcal{X}(\theta)$ and $\mathcal{Y}$ can also be estimated directly from the data for all θ . Consequently we consider a feasible estimation criterion where $\mathcal{X}$ and $\mathcal{Y}$ are replaced by $\hat{\mathcal{X}}(\theta) = \mathcal{T}_{\mathcal{X}}(\theta; \hat{\gamma})$ and $\hat{\mathcal{Y}} = \mathcal{T}_{\mathcal{Y}}(\hat{\gamma})$ respectively, for some preliminary estimator, $\hat{\gamma}$, of γ_0 . We denote the sample counterpart of $\mathcal{S}$ by $\hat{\mathcal{S}}$, so that

$$\hat{\mathcal{S}}(\theta; \widehat{\mathcal{W}}) = (\hat{\mathcal{Y}} - \hat{\mathcal{X}}(\theta))^\top \widehat{\mathcal{W}} (\hat{\mathcal{Y}} - \hat{\mathcal{X}}(\theta)), \quad (10)$$

where $\widehat{\mathcal{W}}$ can be random and depend on the sample size. We define our estimator, $\hat{\theta}(\widehat{\mathcal{W}})$, to be the minimizer of $\hat{\mathcal{S}}(\theta; \widehat{\mathcal{W}})$:

$$\hat{\theta}(\widehat{\mathcal{W}}) = \arg \min_{\theta \in \Theta} \hat{\mathcal{S}}(\theta; \widehat{\mathcal{W}}).$$

Therefore $\hat{\theta}(\widehat{\mathcal{W}})$ is generally a nonlinear least square estimator with generated regressors and regressands. Note that $\hat{\mathcal{S}}(\theta; \widehat{\mathcal{W}})$ is easy to evaluate for different values of θ , following (7) and (8), $\hat{\mathcal{X}}_i(\theta)$ can be computed by a matrix multiplication of $\boldsymbol{\pi}_{i, \theta_i}$ by the estimator of $\mathbf{D}(\mathbf{R}_i + \beta \mathbf{H}_i \mathbf{M} \mathbf{R}_i)$, which does not depend on θ_i , and $\hat{\mathcal{Y}}_i$ is also independent of θ_i .

3 Main Results

We give large sample properties of our estimator in full generality in Section 3.1. We consider special cases when payoffs have linear-in-parameter and partially linear specifications in Section 3.2 and 3.3 respectively. We discuss the relationship between our estimator and ALSE_{PSD} in Section 3.4. In what follows we denote the matrix norm by $\|\cdot\|$, so that $\|B\| = \sqrt{\text{trace}(B^\top B)}$ for any real matrix B , and we let “ $\xrightarrow{p}$ ” and “ $\xrightarrow{d}$ ” denote convergence in probability and distribution respectively.

3.1 General Case

From the previous section, we see that $\mathcal{T}_{\mathcal{X}}(\theta; \cdot)$ and $\mathcal{T}_{\mathcal{Y}}(\cdot)$ are deterministic and smooth functions in γ for any θ . To analyze the asymptotic properties of $\hat{\theta}(\widehat{\mathcal{W}})$, it will be useful to keep separate the sampling distribution of the preliminary estimator and the corresponding generated regressors and regressands. We begin with a preliminary requirement for $\hat{\gamma}$.

ASSUMPTION P: (i) $\hat{\gamma} \xrightarrow{p} \gamma_0$; and, (ii) $\sqrt{N}(\hat{\gamma} - \gamma_0) \xrightarrow{d} \mathcal{N}(0, \Xi)$.

There are several choices for $\hat{\gamma}$ in practice that satisfy P under very weak conditions. The simplest options are perhaps the empirical choice and transition probabilities, otherwise kernel estimators can be employed (Li and Racine (2006)). We now present our regularity conditions and main results in terms of $(\mathcal{X}(\theta), \mathcal{Y})$ and their estimators $(\hat{\mathcal{X}}(\theta), \hat{\mathcal{Y}})$.

ASSUMPTION A1: $\theta_0 \in \text{int}(\Theta)$ where Θ is a compact subset of $\mathbb{R}^p$, and $\mathcal{X}(\theta) = \mathcal{X}(\theta_0)$ if and only if $\theta = \theta_0$.

ASSUMPTION A2: $\widehat{\mathcal{W}} \xrightarrow{p} \mathcal{W}$, where $\mathcal{W}$ is a non-stochastic positive definite matrix.

ASSUMPTION A3: $\sup_{\theta \in \Theta} \|\mathcal{X}(\theta)\|$ and $\|\mathcal{Y}\|$ are finite, and $\sup_{\theta \in \Theta} \|\widehat{\mathcal{X}}(\theta) - \mathcal{X}(\theta)\| \xrightarrow{p} 0$ and $\widehat{\mathcal{Y}} \xrightarrow{p} \mathcal{Y}$.

ASSUMPTION A4: $\mathcal{X}(\theta)$ is continuously differentiable at θ_0 and $\nabla_{\mathcal{X}} = \left. \frac{\partial \mathcal{X}(\theta)}{\partial \theta^\top} \right|_{\theta=\theta_0}$ has full column rank.

ASSUMPTION A5: $\sup_{\theta \in B_\delta(\theta_0)} \left\| \frac{\partial \widehat{\mathcal{X}}(\theta)}{\partial \theta^\top} - \frac{\partial \mathcal{X}(\theta)}{\partial \theta^\top} \right\| \xrightarrow{p} 0$, where $B_\delta(\theta_0)$ denotes some neighborhood of θ_0 .

Define $\widehat{\mathcal{U}} = \widehat{\mathcal{Y}} - \widehat{\mathcal{X}}(\theta_0)$.

ASSUMPTION A6: $\sqrt{N}\widehat{\mathcal{U}} \xrightarrow{d} \mathcal{N}(0, \Sigma)$ for some non-stochastic positive definite matrix Σ .

COMMENTS ON ASSUMPTIONS A1 - A3.

These conditions are sufficient for the consistency of our estimator. A1 - A2 constitute to a high level identification condition as it ensures (9) has a unique solution at θ_0 . There has been little work on more primitive conditions for parametric identification of payoff functions in dynamic games. Most identification results in the literature are nonparametric that build on the work of Magnac and Thesmar (2002); see Pesendorfer and Schmidt-Dengler (2008) and Bajari et al. (2009). However, using Hotz and Miller's inversion, it follows that the condition for identification of the pseudo-model at θ_0 , in the sense that $P_{i,\theta_i} = P_{i,\theta_{i0}}$ for all i if and only if $\theta_i = \theta_{i0}$ for all i , is precisely the identification condition required in A1. Furthermore, by inspecting Lemma R more closely, for each i , we see that the necessary and sufficient condition for the unique parameterization of $\mathcal{X}_i(\theta_i)$ at θ_{i0} is for the intersection between the $\{\boldsymbol{\pi}_{i,\theta_i} - \boldsymbol{\pi}_{i,\theta_{i0}} : \theta_i \in \Theta_i \setminus \{\theta_{i0}\}\}$ and the null space of $\mathbf{D}(\mathbf{R}_i + \beta \mathbf{H}_i \mathbf{M} \mathbf{R})$ to be empty. Although, without any restriction on $\boldsymbol{\pi}_{i,\theta_i}$, A1 generally does not hold since $\mathbf{D}(\mathbf{R}_i + \beta \mathbf{H}_i \mathbf{M} \mathbf{R})$ is always rank-deficient. For a closely related discussion see Srisuma (2013a), who provides constructive conditions for parametric identification results in a single agent model that can be generalized directly to the games considered in this paper. Also see the identification condition and comments of B1 in Section 3.2 when linear-in-parameter restriction is imposed. The uniform boundedness and consistency conditions essentially depend on $\{\boldsymbol{\pi}_{i,\theta_i}\}_{i=1}^I$. In particular, if $\mathbf{D}(\mathbf{R}_i + \beta \mathbf{H}_i \mathbf{M} \mathbf{R})$ is finite then continuity of $\{\boldsymbol{\pi}_{i,\theta_i}\}_{i=1}^I$ ensures $\sup_{\theta \in \Theta} \|\mathcal{X}(\theta)\|$ is finite since Θ is compact. Then uniform consistency also follows if there exists a consistent estimator for $\mathbf{D}(\mathbf{R}_i + \beta \mathbf{H}_i \mathbf{M} \mathbf{R})$, which is implied by P(i).

COMMENTS ON ASSUMPTIONS A4 - A6.

For the distribution theory, additional local conditions around θ_0 are required. A4 - A5 are standard smoothness and regularity conditions for an asymptotic normality of an extremum estimator that optimizes a smooth objective function. Similar to the discussion of sufficient conditions for A3, using Lemma R, a sufficient condition for continuous differentiability of $\mathcal{X}(\theta)$ in A4 is continuous differentiability of π_{i,θ_i} at θ_{i0} for all i , then A5 will also follow if P(i) holds. Furthermore, if P(ii) holds, so that the elements in $\widehat{\mathcal{X}}(\theta_0)$ and $\widehat{\mathcal{Y}}$ have asymptotically normal distribution, then by applying a delta-method A6 also holds with $\Sigma = \nabla_\gamma \Xi \nabla_\gamma^\top$, where $\nabla_\gamma = \frac{\partial}{\partial \gamma^\top} (\mathcal{T}_\mathcal{Y}(\gamma) - \mathcal{T}_\mathcal{X}(\theta_0; \gamma))|_{\gamma=\gamma_0}$.

Our estimators are consistent and asymptotically normal under these assumptions.

THEOREM 1 (CONSISTENCY): *Under assumptions A1 - A3, $\widehat{\theta}(\widehat{\mathcal{W}}) \xrightarrow{p} \theta_0$.*

THEOREM 2 (ASYMPTOTIC NORMALITY): *Under assumptions A1 - A6,*

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

where $\Omega_\mathcal{W} = (\nabla_\mathcal{X}^\top \mathcal{W} \nabla_\mathcal{X})^{-1} \nabla_\mathcal{X}^\top \mathcal{W} \Sigma \mathcal{W} \nabla_\mathcal{X} (\nabla_\mathcal{X}^\top \mathcal{W} \nabla_\mathcal{X})^{-1}$.

In large sample, the estimators that uniquely solve (10) are distinguishable up to the first order by $\Omega_\mathcal{W}$. The efficient estimator in this class can be found by choosing the optimal weighting matrix, $\mathcal{W}^*$, that minimizes $\Omega_\mathcal{W}$ over the set of all possible positive definite matrices (i.e. efficiency gain in the spirit of Chamberlain (1982) and Hansen (1982) for instance).

THEOREM 3 (EFFICIENCY): *Under assumptions A1 - A6, (i) the asymptotic variance of $\sqrt{N} (\widehat{\theta}(\widehat{\mathcal{W}}) - \theta_0)$ is bounded below by $\Omega_{\Sigma^{-1}} = (\nabla_\mathcal{X}^\top \Sigma^{-1} \nabla_\mathcal{X})^{-1}$; and, (ii) if $\widehat{\mathcal{W}} \xrightarrow{p} \Sigma^{-1}$ then $\sqrt{N} (\widehat{\theta}(\widehat{\mathcal{W}}) - \theta_0) \xrightarrow{d} \mathcal{N}(0, \Omega_{\Sigma^{-1}})$.*

The first part of Theorem 3 says that the lower variance bound for the class of estimators we consider is $(\nabla_\mathcal{X}^\top \Sigma^{-1} \nabla_\mathcal{X})^{-1}$. The second part states that any consistent estimator of Σ^{-1} is sufficient to produce an efficient estimator. In practice, consistent estimator for Σ^{-1} will typically require a preliminary consistent estimator for θ_0 . The simplest choice is to choose $\mathcal{W}$ to be an identity matrix, I_d . In this case the estimator for θ_{i0} can be computed individually for each player. We state this in the following corollary.

COROLLARY A (IDENTITY WEIGHTED ESTIMATOR): *Under assumptions A1, A3 - A6, $\sqrt{N} (\widehat{\theta}(I_d) - \theta_0) \xrightarrow{d} \mathcal{N}(0, \Omega_{I_d})$, where $\widehat{\theta}(I_d) = \left(\widehat{\theta}_1(I_d)^\top, \dots, \widehat{\theta}_I(I_d)^\top \right)^\top$. Furthermore, for all i : $\widehat{\theta}_i(I_d) = \arg \min_{\theta_i \in \Theta_i} (\widehat{\mathcal{Y}}_i - \widehat{\mathcal{X}}_i(\theta_i))^\top (\widehat{\mathcal{Y}}_i - \widehat{\mathcal{X}}_i(\theta_i))$ such that $\sqrt{N} (\widehat{\theta}_i(I_d) - \theta_{i0}) \xrightarrow{d} \mathcal{N}(0, (\nabla_{\mathcal{X}_i}^\top \nabla_{\mathcal{X}_i})^{-1} \nabla_{\mathcal{X}_i}^\top \Sigma_i \nabla_{\mathcal{X}_i} (\nabla_{\mathcal{X}_i}^\top \nabla_{\mathcal{X}_i})^{-1})$ with $\nabla_{\mathcal{X}_i} = \frac{\partial \mathcal{X}_i(\theta)}{\partial \theta_i^\top} \Big|_{\theta_i=\theta_{i0}}$ and $\Sigma_i = \lim_{N \rightarrow \infty} \text{Var}(\sqrt{N}(\widehat{\mathcal{Y}}_i - \widehat{\mathcal{X}}_i(\theta_{i0})))$.*

3.2 Linear-in-Parameter Specification

We now consider the leading special case when payoff functions have a linear-in-parameter specification.

ASSUMPTION M5 (*Linear-in-parameter payoffs*). For all $(i, \theta_i, a_i, a_{-i}, x)$,

$$\pi_{i, \theta_i}(a_i, a_{-i}, x) = \theta_i^\top \pi_i(a_i, a_{-i}, x),$$

for some p -dimensional vector $\pi_i(a_i, a_{-i}, x) = (\pi_i^1(a_i, a_{-i}, x), \dots, \pi_i^p(a_i, a_{-i}, x))^\top$.

We assume M1 - M5 hold throughout this subsection. Then, with a slight abuse of notation, $\mathcal{X}_i(\theta_i)$ in (7) simplifies to $\mathcal{X}_i \theta_i$, where

$$\mathcal{X}_i = \mathbf{D}(\mathbf{R}_i + \beta \mathbf{H}_i \mathbf{M} \mathbf{R}) \mathbf{\Pi}_i, \quad (11)$$

and $\mathbf{\Pi}_i$ is a $J(K+1)^I$ by p matrix of $\{\pi_i(a_i, a_{-i}, x)\}_{a_i \in A, x \in X}$. Let $\mathcal{X} = \text{diag}(\mathcal{X}_1, \dots, \mathcal{X}_I)$. The limiting and sample objective functions defined in (9) and (10) respectively become

$$\begin{aligned} \mathcal{S}^{lip}(\theta; \mathcal{W}) &= (\mathcal{Y} - \mathcal{X}\theta)^\top \mathcal{W}(\mathcal{Y} - \mathcal{X}\theta), \text{ and} \\ \widehat{\mathcal{S}}^{lip}(\theta; \widehat{\mathcal{W}}) &= (\widehat{\mathcal{Y}} - \widehat{\mathcal{X}}\theta)^\top \widehat{\mathcal{W}}(\widehat{\mathcal{Y}} - \widehat{\mathcal{X}}\theta). \end{aligned}$$

If $\widehat{\mathcal{X}}^\top \widehat{\mathcal{W}} \widehat{\mathcal{X}}$ is non-singular, then $\widehat{\mathcal{S}}^{lip}(\theta; \widehat{\mathcal{W}})$ is globally convex. The solution to the minimization problem has a well-known closed-form expression of a weighted least squares estimator, namely

$$\widehat{\theta}^{lip}(\widehat{\mathcal{W}}) = \left(\widehat{\mathcal{X}}^\top \widehat{\mathcal{W}} \widehat{\mathcal{X}} \right)^{-1} \widehat{\mathcal{X}}^\top \widehat{\mathcal{W}} \widehat{\mathcal{Y}}. \quad (12)$$

Although the large sample properties for $\widehat{\theta}^{lip}(\widehat{\mathcal{W}})$ follow immediately from Section 3.1, they can be specialized substantially to incorporate M5. Since the results in this subsection may be most relevant for empirical applications we provide some details here.

ASSUMPTION B1: $\mathcal{X}$ has full column rank.

ASSUMPTION B2: $\widehat{\mathcal{W}} \xrightarrow{p} \mathcal{W}$, where $\mathcal{W}$ is a non-stochastic positive definite matrix.

ASSUMPTION B3: $\|\mathcal{X}\|$ and $\|\mathcal{Y}\|$ are finite, and $\widehat{\mathcal{X}} \xrightarrow{p} \mathcal{X}$ and $\widehat{\mathcal{Y}} \xrightarrow{p} \mathcal{Y}$.

Define $\widehat{\mathcal{U}}^{lip} = \widehat{\mathcal{Y}} - \widehat{\mathcal{X}}\theta_0$.

ASSUMPTION B4: $\sqrt{N}\widehat{\mathcal{U}}^{lip} \xrightarrow{d} \mathcal{N}(0, \Sigma^{lip})$ for some non-stochastic positive definite matrix Σ^{lip} .

COMMENTS ON ASSUMPTIONS B1 - B4.

Similar to A1 - A2, B1 and B2 ensure $\mathcal{S}^{lip}(\theta; \mathcal{W})$ has a unique solution at θ_0 . In this case, the full rank condition of $\mathcal{X}$ is a necessary and sufficient condition for the identification of the pseudo-model (for more details see Srisuma (2013)). The sample counterpart of B1, namely the rank condition of $\widehat{\mathcal{X}}$, also has a finite sample significance. If $\widehat{\mathcal{W}}$ is positive definite, then the full column rank condition of $\widehat{\mathcal{X}}$ is necessary and sufficient for $\widehat{\mathcal{S}}^{lip}(\theta; \widehat{\mathcal{W}})$ to have a unique solution, which equals to $\widehat{\theta}^{lip}(\widehat{\mathcal{W}})$ as defined in (12). Assumptions B3 and B4 are immediate specializations of A3 - A6.

We state the large sample properties for $\widehat{\theta}^{lip}(\widehat{\mathcal{W}})$ as corollaries without proofs.

COROLLARY 1 (CONSISTENCY): *Under assumptions B1 - B3, $\widehat{\theta}^{lip}(\widehat{\mathcal{W}}) \xrightarrow{p} \theta_0$.*

COROLLARY 2 (ASYMPTOTIC NORMALITY): *Under assumptions B1 - B4,*

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

where $\Omega_{\mathcal{W}}^{lip} = (\mathcal{X}^\top \mathcal{W} \mathcal{X})^{-1} \mathcal{X}^\top \mathcal{W} \Sigma^{lip} \mathcal{W} \mathcal{X} (\mathcal{X}^\top \mathcal{W} \mathcal{X})^{-1}$.

COROLLARY 3 (EFFICIENCY): *Under assumptions B1 - B4, (i) the asymptotic variance of $\sqrt{N} \left(\widehat{\theta}^{lip}(\widehat{\mathcal{W}}) - \theta_0 \right)$ is bounded below by $\Omega_{\Sigma^{lip-1}}^{lip} = \left(\mathcal{X}^\top \Sigma^{lip-1} \mathcal{X} \right)^{-1}$; and, (ii) if $\widehat{\mathcal{W}} \xrightarrow{p} \Sigma^{lip-1}$ then $\sqrt{N} \left(\widehat{\theta}^{lip}(\widehat{\mathcal{W}}) - \theta_0 \right) \xrightarrow{d} \mathcal{N} \left(0, \Omega_{\Sigma^{lip-1}}^{lip} \right)$.*

Similarly to the general case, consistent estimator for Σ^{lip-1} requires a preliminary consistent estimator for θ_0 . We have the counterpart to Corollary A when we choose $\mathcal{W}$ to be an identity matrix $\mathbf{I}$.

COROLLARY B (IDENTITY WEIGHTED ESTIMATOR): *Under assumptions B1, B3 and B4, $\sqrt{N} \left(\widehat{\theta}^{lip}(\mathbf{I}) - \theta_0 \right) \xrightarrow{d} \mathcal{N} \left(0, \Omega_{\mathbf{I}}^{lip} \right)$, where $\widehat{\theta}^{lip}(\mathbf{I}) = \left(\widehat{\mathcal{X}}^\top \widehat{\mathcal{X}} \right)^{-1} \widehat{\mathcal{X}}^\top \widehat{\mathcal{Y}}$ and $\Omega_{\mathbf{I}}^{lip} = (\mathcal{X}^\top \mathcal{X})^{-1} \mathcal{X}^\top \Sigma \mathcal{X} (\mathcal{X}^\top \mathcal{X})^{-1}$. Furthermore, for all i : $\widehat{\theta}^{lip}(\mathbf{I}) = \left(\widehat{\theta}_1^{lip}(\mathbf{I})^\top, \dots, \widehat{\theta}_I^{lip}(\mathbf{I})^\top \right)^\top$ such that $\widehat{\theta}_i^{lip}(\mathbf{I}) = \left(\widehat{\mathcal{X}}_i^\top \widehat{\mathcal{X}}_i \right)^{-1} \widehat{\mathcal{X}}_i^\top \widehat{\mathcal{Y}}_i$ and $\sqrt{N} \left(\widehat{\theta}_i^{lip}(\mathbf{I}) - \theta_{i0} \right) \xrightarrow{d} \mathcal{N} \left(0, (\mathcal{X}_i^\top \mathcal{X}_i)^{-1} \mathcal{X}_i^\top \Sigma_i^{lip} \mathcal{X}_i (\mathcal{X}_i^\top \mathcal{X}_i)^{-1} \right)$ with $\Sigma_i^{lip} = \lim_{N \rightarrow \infty} \text{Var}(\sqrt{N}(\widehat{\mathcal{Y}}_i - \widehat{\mathcal{X}}_i \theta_{i0}))$.*

We have shown here that once we have $(\widehat{\mathcal{Y}}, \widehat{\mathcal{X}})$, under some regularity conditions, a consistent estimator for θ_0 can be obtained by an OLS estimator, $\widehat{\theta}^{lip}(\mathbf{I}) = \left(\widehat{\mathcal{X}}^\top \widehat{\mathcal{X}} \right)^{-1} \widehat{\mathcal{X}}^\top \widehat{\mathcal{Y}}$ (Corollary B), which can be used to construct an efficient estimator using a familiar a feasible GLS formulation, $\widehat{\theta}^{lip} \left(\widehat{\Sigma}^{lip-1} \right) = \left(\widehat{\mathcal{X}}^\top \widehat{\Sigma}^{lip-1} \widehat{\mathcal{X}} \right)^{-1} \widehat{\mathcal{X}}^\top \widehat{\Sigma}^{lip-1} \widehat{\mathcal{Y}}$ where $\widehat{\Sigma}^{lip-1}$ is a consistent estimator of Σ^{lip-1} .

Our closed-form estimators also readily accommodate linear restrictions. For instance, sometimes there are a priori restrictions one may wish to impose on θ_0 such as symmetry. More formally, suppose

θ_0 is known to satisfy $\mathcal{D}^\top \theta_0 = \delta$ for some known p by q matrix $\mathcal{D}$ that has full row rank $q < p$ and some q -dimensional vector δ . Then a restricted estimator $\tilde{\theta}^{lip}(\widehat{\mathcal{W}})$ that minimizes (10) subject to $\mathcal{D}^\top \tilde{\theta}^{lip}(\widehat{\mathcal{W}}) = \delta$, has the following closed-form expression

$$\tilde{\theta}^{lip}(\widehat{\mathcal{W}}) = \hat{\theta}^{lip}(\widehat{\mathcal{W}}) - \left(\hat{\mathcal{X}}^\top \widehat{\mathcal{W}} \hat{\mathcal{X}} \right)^{-1} \mathcal{D} \left(\mathcal{D}^\top \left(\hat{\mathcal{X}}^\top \widehat{\mathcal{W}} \hat{\mathcal{X}} \right)^{-1} \mathcal{D} \right)^{-1} \left(\mathcal{D}^\top \hat{\theta}^{lip}(\widehat{\mathcal{W}}) - \delta \right),$$

where $\hat{\theta}^{lip}(\widehat{\mathcal{W}})$ is the unrestricted estimator defined in (12). The expression above can be derived using Lagrangean method or through matrix manipulations (see Amemiya (1985, Section 1.4)). And, since $\tilde{\theta}^{lip}(\widehat{\mathcal{W}})$ is an affine transformation of $\hat{\theta}^{lip}(\widehat{\mathcal{W}})$, it is easy to verify that the optimal weighting matrices for $\tilde{\theta}^{lip}(\widehat{\mathcal{W}})$ are the same as those described in Corollary 3, i.e. any $\widehat{\mathcal{W}} \xrightarrow{p} \Sigma^{lip^{-1}}$.

3.3 Partially Linear Specification

One may argue that, in some situations, Assumption M5 is at odds with the spirit of structural estimation if the functions in the vector π_i are interpreted as basis functions. We relax the linear-in-parameter requirement and instead consider a partially linear structure that may arise naturally by ways of additive fixed effects, or, frequently in modeling of entry/exit games, as fixed costs or scrap value. Now suppose $\theta_i = (\theta_i^{A\top}, \theta_i^{B\top})^\top$ for all i .

ASSUMPTION M6 (*Partially linear payoffs*). For all $(i, \theta_i, a_i, a_{-i}, x)$,

$$\pi_{i,\theta_i}(a_i, a_{-i}, x) = \theta_i^{A\top} \pi_i^A(a_i, a_{-i}, x) + \pi_{i,\theta_i^B}^B(a_i, a_{-i}, x),$$

for some p -dimensional vector $\pi_i^A(a_i, a_{-i}, x) = \left(\pi_i^{A1}(a_i, a_{-i}, x), \dots, \pi_i^{Ap}(a_i, a_{-i}, x) \right)^\top$.

We assume M1 - M4 and M6 hold throughout this subsection. Then it is easy to see that the RHS of equation (6) in Lemma R becomes

$$\mathbf{D}(\mathbf{R}_i + \beta \mathbf{H}_i \mathbf{M} \mathbf{R}) \pi_{i,\theta_i}^A + \mathbf{D}(\mathbf{R}_i + \beta \mathbf{H}_i \mathbf{M} \mathbf{R}) \pi_{i,\theta_i^B}^B + \Delta \mathbf{v}_i,$$

and, we define, analogously to (7) and (11), $\mathcal{X}_i^A = \mathbf{D}(\mathbf{R}_i + \beta \mathbf{H}_i \mathbf{M} \mathbf{R}) \Pi_i^A$, and $\mathcal{X}_i^B(\theta_i^B) = \mathbf{D}(\mathbf{R}_i + \beta \mathbf{H}_i \mathbf{M} \mathbf{R}) \pi_{i,\theta_i^B}^B$. Once again, stacking up the vectors from all players, the limiting and sample objective functions defined in (9) and (10) respectively become

$$\begin{aligned} \mathcal{S}^{pl}(\theta; \mathcal{W}) &= (\mathcal{Y} - \mathcal{X}^A \theta^A - \mathcal{X}^B(\theta^B))^\top \mathcal{W} (\mathcal{Y} - \mathcal{X}^A \theta^A - \mathcal{X}^B(\theta^B)), \text{ and} \\ \widehat{\mathcal{S}}^{pl}(\theta; \widehat{\mathcal{W}}) &= (\widehat{\mathcal{Y}} - \widehat{\mathcal{X}}^A \theta^A - \widehat{\mathcal{X}}^B(\theta^B))^\top \widehat{\mathcal{W}} (\widehat{\mathcal{Y}} - \widehat{\mathcal{X}}^A \theta^A - \widehat{\mathcal{X}}^B(\theta^B)), \end{aligned}$$

where the terms in the above display should by now be familiar. In order to avoid repetition we only provide a brief discussion of how θ can be (efficiently) estimated.

The structural identifying condition in this setting is:

$$\mathcal{Y} = \mathcal{X}^A \theta^A + \mathcal{X}^B (\theta^B) \quad \text{if and only if} \quad (\theta^A, \theta^B) = (\theta_0^A, \theta_0^B).$$

The additively linear structure allows us to use a Frisch-Waugh-Lovell type argument to estimate θ_0^A and θ_0^B sequentially in two stages. In particular, θ_0^A and θ_0^B satisfy the following identities:

$$\mathcal{M}_{\mathcal{W}A} \mathcal{Y} = \mathcal{M}_{\mathcal{W}A} \mathcal{X}^B (\theta_0^B), \quad (13)$$

where $\mathcal{M}_{\mathcal{W}A} = I - \mathcal{X}^A (\mathcal{X}^{A\top} \mathcal{W} \mathcal{X}^A)^{-1} \mathcal{X}^{A\top} \mathcal{W}$ is an oblique projection matrix (e.g. see Davidson and MacKinnon (1993)), so that $\mathcal{M}_{\mathcal{W}A} \mathcal{X}^A$ is a matrix of zeros, and

$$\mathcal{Y} - \mathcal{X}^B (\theta_0^B) = \mathcal{X}^A \theta_0^A. \quad (14)$$

An asymptotic least squares estimator that minimizes $\widehat{\mathcal{S}}^{pl}(\theta; \widehat{\mathcal{W}})$ can then be constructed sequentially in two stages. Let

$$\widehat{\mathcal{S}}_1^{pl}(\theta^B; \widehat{\mathcal{W}}) = (\mathcal{M}_{\widehat{\mathcal{W}}A} \widehat{\mathcal{Y}} - \mathcal{M}_{\widehat{\mathcal{W}}A} \widehat{\mathcal{X}} (\theta^B))^\top \widehat{\mathcal{W}} (\mathcal{M}_{\widehat{\mathcal{W}}A} \widehat{\mathcal{Y}} - \mathcal{M}_{\widehat{\mathcal{W}}A} \widehat{\mathcal{X}} (\theta^B)),$$

where $\mathcal{M}_{\widehat{\mathcal{W}}A} = I - \widehat{\mathcal{X}}^A (\widehat{\mathcal{X}}^{A\top} \widehat{\mathcal{W}} \widehat{\mathcal{X}}^A)^{-1} \widehat{\mathcal{X}}^{A\top} \widehat{\mathcal{W}}$. In the first stage we obtain $\widehat{\theta}^{plB}(\widehat{\mathcal{W}}) = \arg \min_{\theta^B} \widehat{\mathcal{S}}_1^{pl}(\theta^B; \widehat{\mathcal{W}})$. For the second stage, let

$$\widehat{\mathcal{S}}_2^{pl}(\theta^A; \widehat{\mathcal{W}}) = (\widehat{\mathcal{Y}} - \widehat{\mathcal{X}}^B(\widehat{\theta}^B) - \widehat{\mathcal{X}}^A \theta^A)^\top \widehat{\mathcal{W}} (\widehat{\mathcal{Y}} - \widehat{\mathcal{X}}^B(\widehat{\theta}^B) - \widehat{\mathcal{X}}^A \theta^A).$$

Then $\widehat{\theta}^{plA}(\widehat{\mathcal{W}}) = \arg \min_{\theta^A} \widehat{\mathcal{S}}_2^{pl}(\theta^A; \widehat{\mathcal{W}}) = (\widehat{\mathcal{X}}^{A\top} \widehat{\mathcal{W}} \widehat{\mathcal{X}}^A)^{-1} \widehat{\mathcal{X}}^{A\top} \widehat{\mathcal{W}} (\widehat{\mathcal{Y}} - \widehat{\mathcal{X}}^B(\widehat{\theta}^B))$. It is easy to verify the first order conditions that $\widehat{\theta}^{plA}(\widehat{\mathcal{W}})$ and $\widehat{\theta}^{plB}(\widehat{\mathcal{W}})$ individually solve are identical to the ones obtained from jointly minimizing $\widehat{\mathcal{S}}^{pl}(\theta; \widehat{\mathcal{W}})$.

The practical advantage of the sequential approach is purely numerical, in the same spirit as the well-known partition regression methods described since the work of Frisch and Waugh (1933). Specifically, we only need to perform nonlinear optimization routine to search over a reduced parameter space for $\widehat{\theta}^{plB}(\widehat{\mathcal{W}})$ in the first stage, as $\widehat{\theta}^{plA}(\widehat{\mathcal{W}})$ has a closed-form expression in terms of $\widehat{\theta}^{plB}(\widehat{\mathcal{W}})$. Note also that the optimal weighting matrix for $\widehat{\mathcal{S}}_1^{pl}$ and $\widehat{\mathcal{S}}_2^{pl}$ is the same, and is identical to the one described in Theorem 3.

3.4 An Equivalent ALSE

Generally it is not possible to directly compare asymptotic efficiency of different estimators in the literature, although they estimate the same model, since many of the estimators are defined using non-nesting objective functions. An exception can be found in Pesendorfer and Schmidt-Dengler (2008), who show ALSE_{PSD} includes some estimators of Aguirregabiria and Mira (2007) and Pakes,

Ostrovsky and Berry (2007) as special cases. Similar to our general estimator defined in Section 2, the class of $ALSE_{PSD}$ is also indexed by a positive definite matrix and optimal weights can be found to define an efficient estimator (cf. Theorem 3). As implied by the Proposition E below, our efficient estimator is asymptotically equivalent to the efficient $ALSE_{PSD}$. In fact, more is true, the class of estimators we consider and that of Pesendorfer and Schmidt-Dengler are asymptotically equivalent in the sense that one can choose appropriate weighting matrices so that the two estimators always have the same asymptotic distribution.

PROPOSITION E. *$ALSE_{PSD}$ and our estimator are asymptotically equivalent.*

The equivalence follows from the existence of a smooth bijective relation between the choice probabilities and the normalized expected payoffs, i.e. essentially by Hotz and Miller’s inversion and an application of the inverse function theorem. The precise relationship between the two estimators are summarized by the equations in display (17) that can be found in the Appendix.

We end this section with a remark on the relationship between asymptotic least squares estimators and GMM estimators. $ALSE_{PSD}$ and our estimator are defined using objective functions that look at the differences between the data and pseudo-model implied probabilities and payoffs respectively at every possible actions and observed states. These differences can also be written as moment conditions, thus asymptotic least squares estimators can also equivalently be defined as GMM estimators (see Chamberlain (1987)). As a consequence, it follows from Proposition E that the GMM estimators of Hotz and Miller (1993) and Hotz et al. (1994) are also asymptotically equivalent for a stationary single agent decision model (a special case of our game when $I = 1$).⁶

4 Monte Carlo Experiments

We illustrate the performance of our closed-form estimator using the Monte Carlo design in Section 7 of Pesendorfer and Schmidt-Dengler (2008); who also provide further comparison with other estimators in the literature.

SETUP

Consider a symmetric two-firm dynamic entry game. In each period t , each firm $i (= 1, 2)$ has two possible choices: be active or not active, $a_{it} \in \{0, 1\}$, where 0 corresponds to “not active” and 1 to “active”. Publically observed state variable has four elements, and can be represented by the

⁶The estimator of Hotz et al. (1994) has an additional source of sampling error since they estimate the discounted expected payoffs, $E[V_{i,\theta_i}(s_{t+1})|x_t, a_{it}]$, by forward simulation. However, under suitable conditions, the error from forward simulation does not affect the asymptotic distribution of their estimator.

actions made by both firms in period $t - 1$, so that $x_t = (a_{1t-1}, a_{2t-1})$. The vector of states evolves over time according to the transition $s_{t+1} = a_t$. Firm 1's period payoffs are described as follows:

$$\pi_{1,\theta}(a_{1t}, a_{2t}, x_t) = \mathbf{1}[a_{1t} = 1] \cdot [\theta_1 + \theta_2 a_{2t}] + \mathbf{1}[a_{1t} = 1, a_{1t-1} = 0] \cdot F + \mathbf{1}[a_{1t} = 0, a_{1t-1} = 1] \cdot W,$$

where $(\theta_1, \theta_2, F, W)$ denote respectively the monopoly profit, duopoly profit, entry cost and scrap value that firm 1 may obtain. 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 set $(\theta_{10}, \theta_{20}, F_0, W_0) = (1.2, -1.2, -0.2, 0.1)$. Pesendorfer and Schmidt-Dengler (2008, p.920) show that there are three distinct equilibria (five if we permute the identity of the players as there is one symmetric equilibrium). We generate the data using different equilibria of the game and provide estimates for $(\theta_{10}, \theta_{20}, F_0)$ for each equilibrium. W_0 is taken as known, since it is not separately identified (see Aguirregabiria and Suzuki (2013)). For each sample size $T = 100, 500, 1000, 5000$, we report the same statistics as Pesendorfer and Schmidt-Dengler (mean and standard deviation of the estimator for each parameter, and the averaged mean squared error across the three parameters) from 1000 simulations of four estimators: OLS, GLS, PSD-I and PSD-E, for each equilibrium. OLS and GLS estimators correspond to our inefficient and efficient estimators that have closed-form (see Corollary B and Corollary 3 respectively). PSD-I and PSD-E are the inefficient and efficient versions of $ALSE_{PSD}$ respectively; the former uses identity weighting matrix. Our Tables 1 - 3 below correspond respectively to equilibria 1 - 3 in Pesendorfer and Schmidt-Dengler (2008), thus are directly comparable to their Tables 1 - 3 on p.921-922.

T	Estimator	F		θ_{10}		θ_{20}		MSE
100	OLS	-0.244	(0.328)	1.071	(0.330)	-1.087	(0.385)	0.396
	GLS	-0.210	(0.136)	1.227	(0.276)	-1.230	(0.255)	0.161
	PSD-I	-0.262	(0.316)	1.083	(0.341)	-1.094	(0.390)	0.395
	PSD-E	-0.175	(0.155)	1.292	(0.303)	-1.327	(0.301)	0.231
500	OLS	-0.213	(0.151)	1.169	(0.141)	-1.161	(0.179)	0.077
	GLS	-0.197	(0.048)	1.213	(0.133)	-1.209	(0.096)	0.029
	PSD-I	-0.220	(0.148)	1.176	(0.144)	-1.167	(0.186)	0.079
	PSD-E	-0.188	(0.047)	1.232	(0.129)	-1.223	(0.102)	0.031
1000	OLS	-0.206	(0.105)	1.184	(0.090)	-1.182	(0.125)	0.035
	GLS	-0.200	(0.030)	1.200	(0.081)	-1.197	(0.062)	0.011
	PSD-I	-0.209	(0.102)	1.186	(0.090)	-1.185	(0.130)	0.036
	PSD-E	-0.195	(0.029)	1.212	(0.077)	-1.204	(0.064)	0.011
5000	OLS	-0.204	(0.079)	1.194	(0.061)	-1.190	(0.093)	0.019
	GLS	-0.206	(0.074)	1.196	(0.059)	-1.192	(0.089)	0.017
	PSD-I	-0.201	(0.079)	1.199	(0.064)	-1.196	(0.094)	0.019
	PSD-E	-0.203	(0.077)	1.198	(0.061)	-1.195	(0.092)	0.018

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

T	Estimator	F		θ_{10}		θ_{20}		MSE
100	OLS	-0.317	(0.472)	0.971	(0.380)	-0.891	(0.543)	0.822
	GLS	-0.428	(0.333)	0.998	(0.328)	-0.892	(0.438)	0.598
	PSD-I	-0.264	(0.495)	1.065	(0.434)	-1.006	(0.592)	0.843
	PSD-E	-0.422	(1.098)	1.073	(0.488)	-0.976	(0.588)	1.903
500	OLS	-0.221	(0.236)	1.147	(0.192)	-1.120	(0.280)	0.181
	GLS	-0.262	(0.210)	1.153	(0.180)	-1.116	(0.261)	0.157
	PSD-I	-0.201	(0.242)	1.192	(0.205)	-1.171	(0.284)	0.182
	PSD-E	-0.232	(0.214)	1.172	(0.182)	-1.154	(0.265)	0.153
1000	OLS	-0.216	(0.168)	1.166	(0.135)	-1.155	(0.196)	0.088
	GLS	-0.233	(0.144)	1.171	(0.123)	-1.157	(0.180)	0.072
	PSD-I	-0.205	(0.171)	1.189	(0.142)	-1.182	(0.201)	0.090
	PSD-E	-0.220	(0.150)	1.177	(0.126)	-1.173	(0.187)	0.075
5000	OLS	-0.205	(0.076)	1.192	(0.058)	-1.189	(0.091)	0.018
	GLS	-0.203	(0.037)	1.196	(0.039)	-1.195	(0.050)	0.005
	PSD-I	-0.202	(0.076)	1.197	(0.061)	-1.196	(0.092)	0.018
	PSD-E	-0.200	(0.043)	1.197	(0.040)	-1.201	(0.058)	0.007

Table 2: *Monte Carlo results (Equilibrium 2). OLS and GLS are our closed-form estimators that are inefficient and efficient respectively. PSD-I and PSD-E are asymptotic least squares estimators of Pesendorfer and Schmidt-Dengler (2008) that are inefficient (identity weighted) and efficient respectively.*

T	Estimator	F		θ_{10}		θ_{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 3: *Monte Carlo results (Equilibrium 3). OLS and GLS are our closed-form estimators that are inefficient and efficient respectively. PSD-I and PSD-E are asymptotic least squares estimators of Pesendorfer and Schmidt-Dengler (2008) that are inefficient (identity weighted) and efficient respectively.*

The results are as expected from the theory. At smaller sample sizes the estimators are genuinely different regardless of the choice of weight matrices. Since the model is fully parametric both efficient estimators generally perform better than the inefficient ones even at $T = 100$ across all equilibria. With larger sample sizes the inefficient and efficient estimators seem to have similar properties for both methods. Although, in theory, the inefficient estimators need not be asymptotically equivalent as both are weighed by the same identity matrix (see equation (17) in the Appendix).

We now abstract away from the statistical properties and consider the numerical aspects. To illustrate the potential for computational advantages of our estimator, 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.⁷ For each M , we solve the model once and simulated five times using the symmetric equilibrium. We report in Table 4, the average central processing unit (CPU) times in seconds to compute our estimators and $ALSE_{PSD}$ that minimize their respective limiting objective functions (no sampling error, using true choice and transition probabilities); standard errors are in parentheses.⁸

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

Table 4: *Computation time. OLS and GLS are our closed-form estimators that are inefficient and efficient respectively. PSD-I and PSD-E are asymptotic least squares estimators of Pesendorfer and Schmidt-Dengler (2008) that are inefficient (identity weighted) and efficient respectively.*

Our estimators are substantially faster to compute, and the distinction grows exponentially with more parameters in the model. The reported CPU times also include the construction of the optimal

⁷There 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 games increases significantly as the games become more complexed. Our design is chosen for its simplicity as it only requires us to solve a simple game multiple times.

⁸The 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.

weighting matrices, using numerical derivatives, for GLS and PSD-E. The procedure to compute the optimal weighting matrices are similar for both (asymptotic least squares) estimators, so its contribution in this setting can be approximated by comparing the CPU times of OLS and GLS as M varies. Our results are model specific and we precaution against extrapolations as different designs, as well as algorithms and softwares, will have different convergence properties for $ALSE_{PSD}$. Although a claim that closed-form estimation is generally a much simpler task is quite innocuous. We also expect the computation time for $ALSE_{PSD}$ to grow at a faster rate with larger action and/or state spaces for any fixed M . Indeed another, perhaps even more important, numerical property of our closed-form estimators is they are always global minimizers. In contrast, a numerical solution to a general nonlinear optimization routine can be sensitive to the search algorithm, initial values, and as well as the nature of the objective function.⁹

5 Conclusions and Possible Extensions

We have shown there can be some non-trivial computational gains in defining estimators that optimize objective functions constructed in terms of expected payoffs instead of choice probabilities for the estimation of structural dynamic discrete choice problems. The most transparent advantages of our approach follow from an opportunity to utilize familiar linear regression techniques, which arise when the period payoff functions are modeled to have fully or partially linear-in-parameter structure. Since the class of estimators we propose is asymptotically equivalent to the unifying class of estimators developed by Pesendorfer and Schmidt-Dengler (2008), there appears to be no theoretical costs associated with our approach to simplify and improve the numerical aspects of the estimation problem. Our estimators also perform well in Monte Carlo exercises in terms of speed and statistical properties.

The computation advantages we describe in this paper accumulates beyond the procedure to obtain a point estimate. For instance, resampling methods that are often used in practice to obtain standard errors (or perhaps to improve finite sample properties) clearly would benefit. The type of objective functions we propose also naturally complements other research in the literature that aims to improve the performance and/or scope of two-step methodologies. Two traditional criticisms of two-step estimators are large finite sample bias (from the first stage nonparametric estimation of choice probabilities), and the inability to accommodate unobserved heterogeneity and state variables that are persistent over time. For the former, Aguirregabiria and Mira (2002,2007) propose an

⁹It is easy to construct a game where the (limiting) objective function defined using pseudo-probabilities has multiple local minima such that some popular built-in optimization package produces different minimizers that depend on the initial search value.

iteration scheme that can improve the finite sample properties by imposing some structure for the first stage estimators; see Kasahara and Shimotsu (2008,2012) for further discussions and some theoretical justifications. At each iteration, the structural estimator can update the choice probabilities implied by the pseudo-model that are then used to define a new pseudo-likelihood function. To incorporate our estimator, alternatively one can use the updated probabilities to construct an objective function that defines the distance between the (updated) observed and implied expected payoffs. For the latter, the recent nonparametric identification results of Kasahara and Shimotsu (2009) and Hu and Shum (2012) show any two-step approach can also be readily applied to estimate a more general dynamic model than the one considered in this paper.

Appendix

Proofs of Theorems

PROOF OF THEOREM 1. Under A1 to A3, $\mathcal{S}(\theta; \mathcal{W})$ has a well-separated minimum at θ_0 . Let $\psi(\theta) = \mathcal{Y} - \mathcal{X}(\theta)$ and $\widehat{\psi}(\theta) = \widehat{\mathcal{Y}} - \widehat{\mathcal{X}}(\theta)$. Under A4, it follows that $\sup_{\theta \in \Theta} \|\psi(\theta)\| < \infty$ and $\sup_{\theta \in \Theta} \|\widehat{\psi}(\theta) - \psi(\theta)\| = o_p(1)$. Then through some tedious algebra, of repeatedly adding nulls and using properties of the matrix norm:

$$\begin{aligned} \widehat{\mathcal{S}}(\theta; \widehat{\mathcal{W}}) - \mathcal{S}(\theta; \mathcal{W}) &= \widehat{\psi}(\theta)^\top \widehat{\mathcal{W}} \widehat{\psi}(\theta) - \psi(\theta)^\top \mathcal{W} \psi(\theta) \\ &= 2\psi(\theta)^\top \mathcal{W} (\widehat{\psi}(\theta) - \psi(\theta)) + o_p(\|\widehat{\psi}(\theta) - \psi(\theta)\|), \end{aligned}$$

where the smaller order terms are uniform over Θ under A2 - A3. Therefore $\sup_{\theta \in \Theta} |\widehat{\mathcal{S}}(\theta; \widehat{\mathcal{W}}) - \mathcal{S}(\theta; \mathcal{W})| = o_p(1)$, and consistency follows from a standard argument (e.g. see Newey and McFadden (1994)). ■

PROOF OF THEOREM 2. Under our assumptions, $\widehat{\theta}(\widehat{\mathcal{W}})$ satisfies the first order condition from differentiating (10) with respect to θ with probability tending to 1, i.e.

$$0 = \left(\frac{\partial \widehat{\mathcal{X}}(\theta)}{\partial \theta^\top} \bigg|_{\theta = \widehat{\theta}(\widehat{\mathcal{W}})} \right)^\top \widehat{\mathcal{W}} (\widehat{\mathcal{Y}} - \widehat{\mathcal{X}}(\widehat{\theta}(\widehat{\mathcal{W}})))$$

holds with probability tending to 1. Since $\mathcal{Y} - \mathcal{X}(\theta_0) = 0$, by adding nulls, we have

$$\begin{aligned} \widehat{\mathcal{Y}} - \widehat{\mathcal{X}}(\widehat{\theta}) &= \widehat{\mathcal{U}} + E_1 + E_2 \\ &= \widehat{\mathcal{U}} - \nabla_{\mathcal{X}} (\widehat{\theta}(\widehat{\mathcal{W}}) - \theta_0) + o_p(\|\widehat{\theta}(\widehat{\mathcal{W}}) - \theta_0\|), \end{aligned}$$

where $E_1 = -(\mathcal{X}(\widehat{\theta}(\widehat{\mathcal{W}})) - \mathcal{X}(\theta_0))$ and $E_2 = \widehat{\mathcal{X}}(\widehat{\theta}(\widehat{\mathcal{W}})) - \widehat{\mathcal{X}}(\theta_0) - (\mathcal{X}(\widehat{\theta}(\widehat{\mathcal{W}})) - \mathcal{X}(\theta_0))$, and the second equality follows from A5 after applying mean value expansions to the terms in E_1 and E_2

around θ_0 . By adding nulls and using properties of matrix norm, since $\widehat{\theta}(\widehat{\mathcal{W}}) = \theta_0 + o_p(1)$, we also have $\left\| \left(\frac{\partial \widehat{\mathcal{X}}(\theta)}{\partial \theta^\top} \Big|_{\theta=\widehat{\theta}(\widehat{\mathcal{W}})} \right)^\top \widehat{\mathcal{W}} - \nabla_{\mathcal{X}}^\top \mathcal{W} \right\| = o_p(1)$ under A2 and A5. Therefore $\widehat{\theta}(\widehat{\mathcal{W}})$ also satisfies

$$0 = \nabla_{\mathcal{X}}^\top \mathcal{W} \left(\widehat{\mathcal{U}} - \nabla_{\mathcal{X}} \left(\widehat{\theta}(\widehat{\mathcal{W}}) - \theta_0 \right) \right) + o_p \left(\frac{1}{\sqrt{N}} + \left\| \widehat{\theta}(\widehat{\mathcal{W}}) - \theta_0 \right\| \right),$$

with probability tending to 1. Then it follows that

$$\sqrt{N} \left(\widehat{\theta}(\widehat{\mathcal{W}}) - \theta_0 \right) = \left(\nabla_{\mathcal{X}}^\top \mathcal{W} \nabla_{\mathcal{X}} \right)^{-1} \nabla_{\mathcal{X}}^\top \mathcal{W} \widehat{\mathcal{U}} + o_p(1).$$

An application of Slutsky's theorem gives the result. ■

PROOF OF THEOREM 3. The proof for part (i) is standard (e.g. see Theorem 3.2 of Hansen (1982)). We claim the optimal weighting matrix converges in the limit to Σ^{-1} . Let $B = \mathcal{W} \nabla_{\mathcal{X}} \left(\nabla_{\mathcal{X}}^\top \mathcal{W} \nabla_{\mathcal{X}} \right)^{-1}$ and $C = \Sigma^{-1} \nabla_{\mathcal{X}} \left(\nabla_{\mathcal{X}}^\top \Sigma^{-1} \nabla_{\mathcal{X}} \right)^{-1}$, so we have $\Omega_{\mathcal{W}} = B^\top \Sigma B$ and $\Omega_{\Sigma^{-1}} = C^\top \Sigma C$. Using simple algebra, it can be shown that $B^\top \Sigma B - C^\top \Sigma C = (B - C)^\top \Sigma (B - C) \geq 0$. For part (ii), it follows from the proof of Theorem 2 that we did not use any specific information on $\widehat{\mathcal{W}}$ beyond the fact that it has a positive definite probability limit. ■

Representation Lemma

PROOF OF LEMMA R. First we introduce some additional notations that build on the terms defined in Section 2.2. Let $v_{i,\theta_i}^a = (v_{i,\theta_i}(a, x^1), \dots, v_{i,\theta_i}(a, x^J))$ for all a , and $\mathbf{v}_{i,\theta_i} = (v_{i,\theta_i}^0, \dots, v_{i,\theta_i}^K)^\top$, so that $\mathbf{v}_{i,\theta_i}$ is a $J(K+1)$ -vector. Let $\pi_{i,\theta_i}^{a_1 \dots a_I} = (\pi_{i,\theta_i}(a_1, \dots, a_I, x^1), \dots, \pi_{i,\theta_i}(a_1, \dots, a_I, x^J))$ for all $a_1, \dots, a_I$, and $\boldsymbol{\pi}_{i,\theta_i} = (\pi_{i,\theta_i}^{0 \dots 0}, \dots, \pi_{i,\theta_i}^{K \dots K})^\top$, so that $\boldsymbol{\pi}_{i,\theta_i}$ is a $J(K+1)^I$ -vector. For any k let: $\mathbf{I}_d$ denote an identity matrix of size d ; $\mathbf{H}_i$ denote a block-diagonal matrix $\mathbf{diag}(H_i^0, H_i^1, \dots, H_i^K)$, where H_i^a denotes a $J \times J$ matrix such that $(H_i^a)_{jj'} = \Pr[x_{t+1} = x^{j'} | x_t = x^j, a_{it} = a]$; $\mathbf{M} = \left(\mathbf{I}_{(K+1)^I} \otimes M \right)$ where $M = (\mathbf{I}_J - L)^{-1}$ and L denotes a $J \times J$ matrix such that $(L)_{jj'} = \beta \Pr[x_{t+1} = x^{j'} | x_t = x^j]$; $\mathbf{R} = \begin{bmatrix} P^{0 \dots 0} & \dots & P^{K \dots K} \\ \vdots & \dots & \vdots \\ P^{0 \dots 0} & \dots & P^{K \dots K} \end{bmatrix}$ be a $J(K+1)^I$ by $J(K+1)^I$ matrix, where $P^{a_1 \dots a_I} = \text{diag}(P(a_1, \dots, a_I | x^1), \dots, P(a_1, \dots, a_I | x^J))$ with $P(a_1, \dots, a_I | x) = \Pr[\alpha_{1,\theta_1}(s_{it}) = a_1, \dots, \alpha_{I,\theta_I}(s_{it}) = a_I | x_t = x] = \prod_{j=1}^I P_j(a_j | x)$, and let $\mathbf{R}_i = \begin{bmatrix} P_{i0}^{0 \dots 0} & \dots & P_{i0}^{K \dots K} \\ \vdots & \dots & \vdots \\ P_{iK}^{0 \dots 0} & \dots & P_{iK}^{K \dots K} \end{bmatrix}$ be a $J(K+1)$ by $J(K+1)^I$ matrix, where $P_{ik}^{a_1 \dots a_I} = \text{diag}(P_{ik}(a_1, \dots, a_I | x^1), \dots, P_{ik}(a_1, \dots, a_I | x^J))$ with $P_{ik}(a_1, \dots, a_I | x) = \Pr[\alpha_{1,\theta_1}(s_{it}) = a_1, \dots, \alpha_{i-1,\theta_{i-1}}(s_{i-1t}) = a_{i-1}, \alpha_{i,\theta_i}(s_{it}) = k, \alpha_{i+1,\theta_{i+1}}(s_{i+1t}) = a_{i+1}, \alpha_{I,\theta_I}(s_{It}) = a_I | x_t = x] =$

$P_i(k|x) \prod_{j \neq i}^I P_j(a_j|x)$. Define $\Delta v_{i,\theta_i}^a = (v_{i,\theta_i}(a, x^1) - v_{i,\theta_i}(0, x^1), \dots, v_{i,\theta_i}(a, x^J) - v_{i,\theta_i}(0, x^J))$ for all $a > 0$, and $\Delta \mathbf{v}_\theta = (\Delta v_{i,\theta_i}^1, \dots, \Delta v_{i,\theta_i}^K)^\top$. Let $\mathbf{D}$ denote the $JK^I \times J(K+1)^J$ matrix that performs the transformation $\mathbf{D}\mathbf{v}_\theta = \Delta \mathbf{v}_\theta$. Lastly, let $\underline{v}_i^a = (\underline{v}_i(a, x^1), \dots, \underline{v}_i(a, x^J))$ for all a , and define $\underline{\mathbf{v}}_i = (\underline{v}_i^0, \dots, \underline{v}_i^K)^\top$, so that $\Delta \underline{\mathbf{v}}_i = \mathbf{D}\underline{\mathbf{v}}_i$ is a $J(K+1)$ -vector. Then (6) immediately follows. ■

Asymptotic Equivalence of ALSEs

PROOF OF PROPOSITION E. In the proof of this proposition we shall assume standard regularity conditions hold throughout (i.e. we assume inverse of matrices exist, expected payoffs and functions are bounded and continuously differentiable etc.). As seen from the proof of Theorem 2, under standard regularity conditions $\widehat{\theta}(\widehat{\mathcal{W}})$ satisfies

$$\widehat{\theta}(\widehat{\mathcal{W}}) = \theta_0 + (\nabla_{\mathcal{X}}^\top \mathcal{W} \nabla_{\mathcal{X}})^{-1} \nabla_{\mathcal{X}}^\top \mathcal{W} \widehat{\mathcal{U}} + o_p\left(\frac{1}{\sqrt{N}}\right). \quad (15)$$

Next we introduce ALSE_{PSD} . It shall be useful to bear in mind the illustrative discussion in Section 2.1. We first define some additional notations that build on the terms defined in Section 2.3. Let $\mathbf{P} = (\mathbf{P}_1^\top, \dots, \mathbf{P}_I^\top)^\top$ and $\mathbf{P}_\theta = (\mathbf{P}_{1,\theta_1}^\top, \dots, \mathbf{P}_{I,\theta_I}^\top)^\top$. Similarly, let $\Delta \mathbf{v} = (\Delta \mathbf{v}_1^\top, \dots, \Delta \mathbf{v}_I^\top)^\top$ and $\Delta \mathbf{v}_\theta = (\Delta \mathbf{v}_{1,\theta_1}^\top, \dots, \Delta \mathbf{v}_{I,\theta_I}^\top)^\top$. Then, by Hotz and Miller's inversion there exists an invertible and continuously differentiable map Γ such that $\mathbf{P} = \Gamma(\Delta \mathbf{v})$ and $\mathbf{P}_\theta = \Gamma(\Delta \mathbf{v}_\theta)$. In particular

$$\begin{aligned} \mathbf{P} &= \left(\Gamma_1(\Delta \mathbf{v}_1)^\top, \dots, \Gamma_I(\Delta \mathbf{v}_I)^\top \right)^\top, \text{ and} \\ \mathbf{P}_\theta &= \left(\Gamma_1(\Delta \mathbf{v}_{1,\theta_1})^\top, \dots, \Gamma_I(\Delta \mathbf{v}_{I,\theta_I})^\top \right)^\top, \end{aligned}$$

where Γ_i is the inverse of Φ_i , which is defined in the text. Therefore, in terms of $\mathcal{Y}$ and $\mathcal{X}(\theta)$,

$$\Delta \mathbf{v} - \Delta \mathbf{v}_\theta = \mathcal{Y} - \mathcal{X}(\theta).$$

Thus $\mathbf{P}$ and $\mathbf{P}_\theta$ are also deterministic functions of the preliminary estimators (that we denoted by γ_0). We denote the estimators of $\mathbf{P}$ and $\mathbf{P}_\theta$ by $\widetilde{\mathbf{P}}$ and $\widehat{\mathbf{P}}_\theta$ respectively, and these estimators are constructed based on the same $\widehat{\gamma}$ that define $\widehat{\mathcal{X}}$ and $\widehat{\mathcal{Y}}$. Note that, although $\mathbf{P} = \mathbf{P}_{\theta_0}$, $\widetilde{\mathbf{P}}$ and $\widehat{\mathbf{P}}_{\theta_0}$ are generally different. An ALSE_{PSD} , denoted by $\widehat{\theta}_{PSD}(\widehat{\mathcal{V}})$, is defined as the minimizer of

$$\min_{\theta \in \Theta} \left(\widetilde{\mathbf{P}} - \widehat{\mathbf{P}}_\theta \right)^\top \widehat{\mathcal{V}} \left(\widetilde{\mathbf{P}} - \widehat{\mathbf{P}}_\theta \right),$$

for some $\widehat{\mathcal{V}}$ that converges in probability to positive definite matrix $\mathcal{V}$ (cf. equation (21) on page 915 in Pesendorfer and Schmidt-Dengler (2008)). Under appropriate regularity conditions, it is straightforward to show, analogous to our Theorem 2, that

$$\sqrt{N} \left(\widehat{\theta}_{PSD}(\widehat{\mathcal{V}}) - \theta_0 \right) \xrightarrow{d} N(0, \Psi_{\mathcal{V}}).$$

For a first order asymptotic equivalence, it suffices to only consider the local asymptotic properties of ALSE_{PSD} around θ_0 . Let $\nabla_{\mathbf{P}}$ denote $\frac{\partial \mathbf{P}_\theta}{\partial \theta^\top} \big|_{\theta=\theta_0}$. An ALSE_{PSD} satisfies

$$0 = -\nabla_{\mathbf{P}}^\top \mathcal{V} \left(\tilde{\mathbf{P}} - \mathbf{P} - \left(\hat{\mathbf{P}}_{\hat{\theta}_{PSD}(\mathcal{V})} - \mathbf{P}_{\theta_0} \right) \right) + o_p \left(\frac{1}{\sqrt{N}} \right).$$

As the problem is smooth, it can be shown generally that the condition above simplifies further to

$$0 = -\nabla_{\mathbf{P}}^\top \mathcal{V} \left(\tilde{\mathbf{P}} - \mathbf{P} - \left(\hat{\mathbf{P}}_{\theta_0} - \mathbf{P}_{\theta_0} + \mathbf{P}_{\hat{\theta}_{PSD}(\mathcal{V})} - \mathbf{P}_{\theta_0} \right) \right) + o_p \left(\frac{1}{\sqrt{N}} \right).$$

So that we have

$$\hat{\theta}_{PSD}(\mathcal{V}) = \theta_0 + (\nabla_{\mathbf{P}}^\top \mathcal{V} \nabla_{\mathbf{P}})^{-1} \nabla_{\mathbf{P}}^\top \mathcal{V} \left(\tilde{\mathbf{P}} - \mathbf{P} - (\hat{\mathbf{P}}_{\theta_0} - \mathbf{P}_{\theta_0}) \right) + o_p \left(\frac{1}{\sqrt{N}} \right).$$

By chain rule $\nabla_{\mathbf{P}}$ equals $\nabla_{\Gamma} \nabla_{\mathcal{X}}$, where ∇_{Γ} denotes the Jacobian of Γ evaluated at $\Delta \mathbf{v}$, and $\frac{\partial \Delta \mathbf{v}_\theta}{\partial \theta^\top} \big|_{\theta=\theta_0}$ equals $\nabla_{\mathcal{X}}$. Thus, we can write

$$\begin{aligned} \hat{\theta}_{PSD}(\mathcal{V}) &= \theta_0 + (\nabla_{\mathcal{X}}^\top \nabla_{\Gamma}^\top \mathcal{V} \nabla_{\Gamma} \nabla_{\mathcal{X}})^{-1} \nabla_{\mathcal{X}}^\top \nabla_{\Gamma}^\top \mathcal{V} \left(\tilde{\mathbf{P}} - \mathbf{P} - (\hat{\mathbf{P}}_{\theta_0} - \mathbf{P}_{\theta_0}) \right) + o_p \left(\frac{1}{\sqrt{N}} \right) \\ &= \theta_0 + (\nabla_{\mathcal{X}}^\top \nabla_{\Gamma}^\top \mathcal{V} \nabla_{\Gamma} \nabla_{\mathcal{X}})^{-1} \nabla_{\mathcal{X}}^\top \nabla_{\Gamma}^\top \mathcal{V} \nabla_{\Gamma} \hat{\mathcal{U}} + o_p \left(\frac{1}{\sqrt{N}} \right), \end{aligned}$$

where the last equality follows from linearizing $\tilde{\mathbf{P}} - \mathbf{P} - (\hat{\mathbf{P}}_{\theta_0} - \mathbf{P}_{\theta_0})$ in terms of $\hat{\mathcal{Y}} - \hat{\mathcal{X}}(\theta_0)$. By defining $\mathcal{W}_{\mathcal{V}} = \nabla_{\Gamma}^\top \mathcal{V} \nabla_{\Gamma}$, we have

$$\hat{\theta}_{PSD}(\mathcal{V}) = \theta_0 + (\nabla_{\mathcal{X}}^\top \mathcal{W}_{\mathcal{V}} \nabla_{\mathcal{X}})^{-1} \nabla_{\mathcal{X}}^\top \mathcal{W}_{\mathcal{V}} \hat{\mathcal{U}} + o_p \left(\frac{1}{\sqrt{N}} \right). \quad (16)$$

Therefore, by comparing (15) and (16), $\hat{\theta}_{PSD}(\mathcal{V})$ has the same asymptotic distribution as $\hat{\theta}(\mathcal{W}_{\mathcal{V}})$. In particular, let $\mathcal{V}^*$ denote the efficient weighting matrix for ALSE_{PSD} so that $\Psi_{\mathcal{V}^*} \leq \Psi_{\mathcal{V}}$ for any positive definite matrix $\mathcal{V}$. Therefore the efficient ALSE_{PSD} , denoted by $\hat{\theta}_{PSD}^*$, has the same asymptotic distribution as $\hat{\theta}(\mathcal{W}_{\mathcal{V}^*})$ with $\mathcal{W}_{\mathcal{V}^*} = \nabla_{\Gamma}^\top \mathcal{V}^* \nabla_{\Gamma}$. Then it must hold, by Theorem 3(i), that $\Omega_{\Sigma^{-1}} \leq \Psi_{\mathcal{V}^*}$ since $\Omega_{\Sigma^{-1}}$ is the lower variance bound. To complete the proof, an identical argument can be made in the reverse direction. It is easy to show that any $\hat{\theta}(\mathcal{W})$ that satisfies (15) also has the same asymptotic distribution as $\hat{\theta}_{PSD}(\mathcal{V}_{\mathcal{W}})$, where $\mathcal{V}_{\mathcal{W}} = \nabla_{\Gamma^{-1}}^\top \mathcal{W} \nabla_{\Gamma^{-1}}$ (cf. $\mathcal{W}_{\mathcal{V}}$), and $\nabla_{\Gamma^{-1}}$ denotes the Jacobian of Γ^{-1} evaluated at $\mathbf{P}$ (that equals $(\nabla_{\Gamma})^{-1}$ by the inverse function theorem). We omit further details to avoid repetition. Thus, it follows that $\Psi_{\mathcal{V}^*} \leq \Omega_{\Sigma^{-1}}$, hence we can also conclude that $\Psi_{\mathcal{V}^*} = \Omega_{\Sigma^{-1}}$.

In summary:

$$\begin{aligned} \sqrt{N} \left(\hat{\theta}(\mathcal{W}) - \theta_0 \right) &= \sqrt{N} \left(\hat{\theta}_{PSD}(\mathcal{V}_{\mathcal{W}}) - \theta_0 \right) + o_p \left(\frac{1}{\sqrt{N}} \right) \text{ with } \mathcal{V}_{\mathcal{W}} = \nabla_{\Gamma^{-1}}^\top \mathcal{W} \nabla_{\Gamma^{-1}}, \quad (17) \\ \sqrt{N} \left(\hat{\theta}_{PSD}(\mathcal{V}) - \theta_0 \right) &= \sqrt{N} \left(\hat{\theta}(\mathcal{W}_{\mathcal{V}}) - \theta_0 \right) + o_p \left(\frac{1}{\sqrt{N}} \right) \text{ with } \mathcal{W}_{\mathcal{V}} = \nabla_{\Gamma}^\top \mathcal{V} \nabla_{\Gamma}, \end{aligned}$$

and $(\mathcal{V}, \mathcal{W})$ can be replaced by any consistent estimators $(\widehat{\mathcal{V}}, \widehat{\mathcal{W}})$. Therefore our estimator and $ALSE_{PSD}$ can always be constructed to have the same asymptotic distribution and achieve the same lower variance bound. ■

References

- [1] Akerberg, D., L. Benkard, S. Berry, and A. Pakes (2005), “Econometric Tools for Analyzing Market Outcome,” *Handbook of Econometrics*, vol. 6, eds. J. Heckman and E. Leamer. North-Holland.
- [2] Aguirregabiria, V., and P. Mira (2002): “Swapping Nested Fixed Point Algorithm: a Class of Estimators for Discrete Markov Decision Models,” *Econometrica*, **70**, 1519-1543.
- [3] Aguirregabiria, V. and P. Mira (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, **75**, 1-53.
- [4] Aguirregabiria, V., and P. Mira (2010): “Dynamic Discrete Choice Structural Models: A Survey,” *Journal of Econometrics*, **156**, 38-67
- [5] Aguirregabiria, V. and J. Suzuki (2013): “Identification and Counterfactuals in Dynamic Models of Market Entry and Exit,” *Working Paper, University of Toronto*.
- [6] Amemiya, T. (1985): *Advanced Econometrics*, Harvard University Press.
- [7] Bajari, P. and H. Hong (2006): “Semiparametric Estimation of a Dynamic Game of Incomplete Information,” *NBER Technical Working Paper 320*.
- [8] Bajari, P., C.L. Benkard, and J. Levin (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, **75**, 1331-1370.
- [9] Bajari, P., V. Chernozhukov, H. Hong and D. Nekipelov (2009): “Identification and Efficient Estimation of a Dynamic Discrete Game,” *Working paper*, University of Minnesota.
- [10] Chamberlain, G. (1982): “Multivariate Regression Models for Panel Data,” *Journal of Econometrics*, **18**, 5-46.
- [11] Collard-Wexler, A. (2013): “Demand Fluctuations and Plant Turnover in the Ready-Mix Concrete Industry,” *forthcoming in Econometrica*.
- [12] Davidson, R. and J.G. MacKinnon (1993): *Estimation and Inference in Econometrics*, Oxford University Press.

- [13] Egedal, M., Z. Lai and C. Su (2013): “Estimating Dynamic Discrete-Choice Games of Incomplete Information,” *Working Paper, University of Chicago Booth School of Business*.
- [14] Frisch, R. and F.V. Waugh (1933): “Partial Time Regressions as Compared with Individual Trends,” *Econometrica*, **1**, 387-401.
- [15] Gourieroux, C. and A. Monfort (1995): *Statistics and Econometric Models: Volume 1, General Concepts, Estimation, Prediction and Algorithms*, Themes in Modern Econometrics, Cambridge University Press.
- [16] Hansen, L.P. (1982): “Large Sample Properties of Generalized Method of Moments Estimators,” *Econometrica*, **50**, 1029 -1054.
- [17] Hotz, V., and R.A. Miller (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, **60**, 497-531.
- [18] Hotz, V., R.A. Miller, S. Smith and J. Smith (1994): “A Simulation Estimator for Dynamic Models of Discrete Choice,” *Review of Economic Studies*, **61**, 265-289.
- [19] Hu, Y. and M. Shum (2012): “Nonparametric Identification of Dynamic Models with Unobserved State Variables,” *Journal of Econometrics*, **171**, 32-44.
- [20] Kasahara, H. and K. Shimotsu (2008): “Pseudo-Likelihood Estimation and Bootstrap Inference for Structural Discrete Markov Decision Models,” *Journal of Econometrics*, **146**, 92-106.
- [21] Kasahara, H. and K. Shimotsu (2009), “Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices,” *Econometrica*, **77**, 135-175.
- [22] Kasahara, H. and K. Shimotsu (2012), “Sequential Estimation of Structural Models with a Fixed Point Constraint,” *Econometrica*, **80**, 2303-2319
- [23] Li, Q. and J.S. Racine (2006): *Nonparametric Econometrics*, Princeton University Press.
- [24] Magnac, M. and D. Thesmar (2002): “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, **70**, 801-816.
- [25] Matzkin, R.L. (1991): “Semiparametric Estimation of Monotone and Concave Utility Functions for Polychotomous Choice Models,” *Econometrica*, **59**, 1315–1327.
- [26] Newey, W.K. and D.L. McFadden (1994): “Large Sample Estimation and Hypothesis Testing,” *Handbook of Econometrics*, vol. 4, eds. R.F. Engle and D. McFadden. North-Holland.

- [27] Norets, A. and S. Takahashi (2013): “On the Surjectivity of the Mapping Between Utilities and Choice Probabilities,” *Quantitative Economics*, **4**, 149–155.
- [28] 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.
- [29] Pesendorfer, M., and P. Schmidt-Dengler (2008): “Asymptotic Least Squares Estimator for Dynamic Games,” *Review of Economics Studies*, **75**, 901–928.
- [30] Ryan, S. P. (2012): “The Costs of Environmental Regulation in a Concentrated Industry,” *Econometrica*, **80**, 1019–1061.
- [31] Rust, J. (1994): “Structural Estimation of Markov Decision Processes,” *Handbook of Econometrics*, vol. 4, eds. R.F. Engle and D. McFadden. North-Holland.
- [32] Rust, J. (1996): “Numerical Dynamic Programming in Economics,” *Handbook of Computational Economics*, vol. 1, eds. H.M. Aumann, D.A. Kendrick and J. Rust. Elsevier.
- [33] Srisuma, S. (2013a): “Identification in Markov Decision Models,” *Working Paper, University of Cambridge*.
- [34] Srisuma, S. (2013b): “Minimum Distance Estimators for Dynamic Games,” *Quantitative Economics*, **4**, 549–583.
- [35] Srisuma, S. and O.B. Linton (2012): “Semiparametric Estimation of Markov Decision Processes with Continuous State Space,” *Journal of Econometrics*, **166**, 320–341.
- [36] Tamer, E. (2003), “Incomplete Simultaneous Discrete Response Model with Multiple Equilibria,” *Review of Economic Studies*, **70**, 147–165.