cemmap

centre for microdata methods and practice

On the iterated estimation of dynamic discrete choice games

Federico A. Bugni
Jackson Bunting

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP13/18

E·S·R·C
ECONOMIC & SOCIAL RESEARCH COUNCIL

An ESRC Research Centre

On the iterated estimation of dynamic discrete choice games*

Federico A. Bugni
Department of Economics
Duke University
federico.bugni@duke.edu

Jackson Bunting
Department of Economics
Duke University
jackson.bunting@duke.edu

February 14, 2018

Abstract

We study the asymptotic properties of a class of estimators of the structural parameters in dynamic discrete choice games. We consider K-stage policy iteration (PI) estimators, where K denotes the number of policy iterations employed in the estimation. This class nests several estimators proposed in the literature. By considering a "maximum likelihood" criterion function, our estimator becomes the K-ML estimator in Aguirregabiria and Mira (2002, 2007). By considering a "minimum distance" criterion function, it defines a new K-MD estimator, which is an iterative version of the estimators in Pesendorfer and Schmidt-Dengler (2008) and Pakes et al. (2007).

First, we establish that the K-ML estimator is consistent and asymptotically normal for any K. This complements findings in Aguirregabiria and Mira (2007), who focus on $K = 1$ and K large enough to induce convergence of the estimator. Furthermore, we show that the asymptotic variance of the K-ML estimator can exhibit arbitrary patterns as a function K.

Second, we establish that the K-MD estimator is consistent and asymptotically normal for any K. For a specific weight matrix, the K-MD estimator has the same asymptotic distribution as the K-ML estimator. Our main result provides an optimal sequence of weight matrices for the K-MD estimator and shows that the optimally weighted K-MD estimator has an asymptotic distribution that is invariant to K. This new result is especially unexpected given the findings in Aguirregabiria and Mira (2007) for K-ML estimators. Our main result implies two new and important corollaries about the optimal 1-MD estimator (derived by Pesendorfer and Schmidt-Dengler (2008)). First, the optimal 1-MD estimator is optimal in the class of K-MD estimators for all K. In other words, additional policy iterations do not provide asymptotic efficiency gains relative to the optimal 1-MD estimator. Second, the optimal 1-MD estimator is more or equally asymptotically efficient than any K-ML estimator for all K.

Keywords: dynamic discrete choice problems, dynamic games, pseudo maximum likelihood estimator, minimum distance estimator, estimation, asymptotic efficiency.

JEL Classification Codes: C13, C61, C73

*We thank Peter Arcidiacono, Patrick Bayer, Joe Hotz, Matt Masten, Arnaud Maurel, Takuya Ura, and the participants at the Duke microeconometrics reading group for useful comments and suggestions. Of course, any and all errors are our own. The research of the first author was supported by NSF Grant SES-1729280.

1 Introduction

This paper investigates the asymptotic properties of a wide class of estimators of the structural parameters in a dynamic discrete choice game, i.e., a dynamic game with discrete actions. In particular, we consider the class of K-stage policy iteration (PI) estimators, where $K \in \mathbb{N}$ denotes the number of policy iterations employed in the estimation. By definition, the K-stage PI estimator is defined by:

$$\hat{\alpha}_K \equiv \arg\max_{\alpha \in \Theta} \hat{Q}(\alpha, \hat{P}_{K-1}), \tag{1.1}$$

where $\alpha^* \in \Theta_\alpha$ is the true value of the structural parameter of interest, $\hat{Q}$ is a criterion function, and $\hat{P}_k$ is the k-stage estimator of the conditional choice probabilities (CCPs), defined iteratively as follows. The preliminary or 0-stage estimator of the CCP is denoted by $\hat{P}_0$. Then, for any $k = 1, \ldots, K-1$,

$$\hat{P}_k \equiv \Psi(\hat{\alpha}_k, \hat{P}_{k-1}), \tag{1.2}$$

where Ψ is the best response CCP mapping of the structural game. Given any set of beliefs P, optimal or not, $\Psi(\alpha, P)$ indicates the corresponding optimal CCPs when the structural parameter is α. The idea of using iterations to estimate the CCPs in Eq. (1.2) was introduced to the literature in the seminal contributions of Aguirregabiria and Mira (2002, 2007).

Our K-stage PI estimator nests most of the estimators proposed in the dynamic discrete choice games literature. By appropriate choice of $\hat{Q}$ and K, our K-stage PI estimator coincides with the pseudo maximum likelihood (PML) estimator proposed by Aguirregabiria and Mira (2002, 2007), the asymptotic least squares estimators proposed by Pesendorfer and Schmidt-Dengler (2008), or the so-called simple estimators proposed by Pakes et al. (2007).

To implement the K-stage PI estimator, the researcher must determine the number of policy iterations $K \in \mathbb{N}$. This choice poses several related research questions. How should researchers choose K? Does it make a difference? If so, what is the "optimal" choice of K? The literature provides arguably incomplete answers to these questions. The main contribution of this paper is to answer these questions. Before describing our results, we review the main related findings in the literature.

Aguirregabiria and Mira (2002, 2007) propose K-stage PML estimators of the structural parameters in dynamic discrete choice problems. The earlier paper considers single-agent problems whereas the later one generalizes the analysis to multiple-agent problems, i.e., games. In both of these papers, the objective function is the pseudo log-likelihood criterion function $\hat{Q} = \hat{Q}_{ML}$, defined by:

$$\hat{Q}_{ML}(\alpha, P) \equiv \frac{1}{n} \sum_{i=1}^{n} \ln \left(\Psi(\alpha, P)(a_i | x_i) \right). \tag{1.3}$$

In this paper, we refer to the resulting estimator as the K-ML estimator. One of the main contributions of Aguirregabiria and Mira (2002, 2007) is to study the effect of the number of iterations K on the asymptotic distribution of the K-ML estimator.

In single-agent dynamic problems, Aguirregabiria and Mira (2002) show that the asymptotic distribution of the K-ML estimators is invariant to K. In other words, any additional round of (computationally costly) policy mapping iteration has no effect on the asymptotic distribution. This striking result is a consequence of the so-called "zero Jacobian property" that occurs when a single agent makes optimal decisions. This property does not hold in dynamic problems with multiple players, as each player makes optimal choices according to their preferences, which may not be aligned with their competitors' preferences. Thus, in dynamic discrete choice games, one should expect that the asymptotic distribution of K-ML estimators changes with K.

In multiple-agent dynamic games, Aguirregabiria and Mira (2007) show that the asymptotic distribution of the K-ML estimators is *not* invariant to K. Their contribution considers two specific choices of K. On one hand, they consider the 1-ML estimator, i.e., $K = 1$, which they refer to as the two-step pseudo maximum likelihood (PML) estimator. On the other hand, they consider the K-ML estimator that results from increasing K until the estimator converges (i.e., $\hat{\alpha}_{(K-1)-ML} = \hat{\alpha}_{K-ML}$), which they refer to as the nested pseudo likelihood (NPL) estimator. Kasahara and Shimotsu (2012) provide conditions that guarantee this convergence. For concreteness, we refer to this estimator as the ∞-ML estimator. Under some conditions, Aguirregabiria and Mira (2007) show that the 1-ML and ∞-ML estimators are consistent and asymptotically normal estimators of α, i.e.,

$$\begin{aligned} \sqrt{n}(\hat{\alpha}_{1-ML} - \alpha^*) &\overset{d}{\to} N(0, \Sigma_{1-ML}) \\ \sqrt{n}(\hat{\alpha}_{\infty-ML} - \alpha^*) &\overset{d}{\to} N(0, \Sigma_{\infty-ML}). \end{aligned} \tag{1.4}$$

More importantly, under additional conditions, Aguirregabiria and Mira (2007) show that $\Sigma_{1-ML} - \Sigma_{\infty-ML}$ is positive definite, that is, the ∞-ML estimator is asymptotically more efficient than the 1-ML estimator. So although iterating the K-ML estimator until convergence might be computationally costly, it improves asymptotic efficiency.

In later work, Pesendorfer and Schmidt-Dengler (2010) indicate that the ∞-ML estimator may be inconsistent in certain games with unstable equilibria. The intuition for this is as follows. To derive Eq. (1.4), Aguirregabiria and Mira (2007) consider an asymptotic framework with $n \to \infty$ and then $K \to \infty$ (in that order). In practice, however, researchers consider $K \to \infty$ for a large but fixed sample size n. Thus, the more relevant asymptotic framework for ∞-ML is $K \to \infty$ and then $n \to \infty$ (again, in that order). Pesendorfer and Schmidt-Dengler (2010) show that, in an unstable equilibrium, reversing the order of the limits can produce very different asymptotic results.[1] In any case, the main findings in Aguirregabiria and Mira (2007) are still applicable for dynamic games in which we observe data from a stable equilibrium.

Pesendorfer and Schmidt-Dengler (2008) consider estimation of dynamic discrete choice games using a class of minimum distance (MD) estimators. Specifically, their objective is to minimize the sample criterion function $\hat{Q} = \hat{Q}_{MD}$, given by:

$$\hat{Q}_{MD}(\alpha, \hat{P}_0) \equiv (\hat{P}_0 - \Psi(\alpha, \hat{P}_0))'\hat{W}(\hat{P}_0 - \Psi(\alpha, \hat{P}_0)),$$

where $\hat{W}$ is a weight matrix that converges in probability to a limiting weight matrix W.[2] This is a single-stage estimator and, consequently, we refer to it as the 1-MD estimator. Pesendorfer and Schmidt-Dengler (2008) show that the 1-MD estimator is a consistent and asymptotically normal estimator of α, i.e.,

$$\sqrt{n}(\hat{\alpha}_{1-MD} - \alpha^*) \overset{d}{\to} N(0, \Sigma_{1-MD}(W)).$$

Pesendorfer and Schmidt-Dengler (2008) show that an appropriate choice of $\hat{W}$ implies that the 1-MD is

[1] To avoid the problem raised by Pesendorfer and Schmidt-Dengler (2010), we consider the asymptotic analysis for K-stage PI estimators with fixed K and $n \to \infty$.

[2] Pesendorfer and Schmidt-Dengler (2008) also allow the preliminary estimators to include estimators other than $\hat{P}_0$ (e.g. estimator of the transition probability). We ignore this issue for simplicity for now, but allow it in later sections of the paper.

asymptotically equivalent to the 1-ML estimator in Aguirregabiria and Mira (2007) or the simple estimators proposed by Pakes et al. (2007). Furthermore, Pesendorfer and Schmidt-Dengler (2008) characterize the optimal choice of W (in terms of minimizing $\Sigma_{1-MD}(W)$), denoted by W^*_{1-MD}. In general, $\Sigma_{1-ML} - \Sigma_{1-MD}(W^*_{1-MD})$ is positive semidefinite, i.e., the optimal 1-MD estimator is more or equally asymptotically efficient than the 1-ML estimator.

In the light of the results in Aguirregabiria and Mira (2007) and Pesendorfer and Schmidt-Dengler (2008), it is natural to inquire whether an iterated version of the MD estimator could yield asymptotic efficiency gains relative to the 1-MD or the K-ML estimators. The consideration of iterated MD estimator opens several important research questions. How should we define the iterated version of the MD estimator? Does this strategy result in consistent and asymptotically normal estimators of α^*? If so, how should we choose the weight matrix W? What about the number of iterations K? Finally, does iterating the MD estimator produce asymptotic efficiency gains as Aguirregabiria and Mira (2007) find for the K-ML estimators? This paper provides answers to all of these questions.

We now summarize the main findings of our paper. We consider a standard dynamic discrete-choice game as in Aguirregabiria and Mira (2007) or Pesendorfer and Schmidt-Dengler (2008). In this context, we investigate the asymptotic properties of K-ML and K-MD estimators.

First, we establish that the K-ML estimator is consistent and asymptotically normal for any $K \in \mathbb{N}$. This complements findings in Aguirregabiria and Mira (2007), who focus on $K = 1$ and K large enough to induce convergence of the estimator. Furthermore, we show that the asymptotic variance of the K-ML estimator can exhibit arbitrary patterns as a function of K. In particular, depending on the parameters of the dynamic problem, the asymptotic variance can increase, decrease, or even oscillate with K.

Second, we also establish that the K-MD estimator is consistent and asymptotically normal for any $K \in \mathbb{N}$. This is a novel contribution relative to Pesendorfer and Schmidt-Dengler (2008) or Pakes et al. (2007), who focus on non-iterative 1-MD estimators. The asymptotic distribution of the K-MD estimator depends on the choice of the weight matrix. For a specific weight matrix, the K-MD has the same asymptotic distribution as the K-ML. We investigate the optimal choice of the weight matrix for the K-MD estimator.

Our main result, Theorem 4.3, shows that an optimal K-MD estimator has an asymptotic distribution that is invariant to K. This appears to be a novel result in the literature on PI estimation for games, and it is particularly surprising given the findings in Aguirregabiria and Mira (2007) for K-ML estimators. Our main result implies two important corollaries regarding the optimal 1-MD estimator (derived by Pesendorfer and Schmidt-Dengler (2008)):

1. The optimal 1-MD estimator is asymptotically efficient in the class of K-MD estimators for all $K \in \mathbb{N}$. In other words, additional policy iterations do not provide asymptotic efficiency gains relative to the optimal 1-MD estimator.

2. The optimal 1-MD estimator is more or equally asymptotically efficient than any K-ML estimator for all $K \in \mathbb{N}$.

The remainder of the paper is organized as follows. Section 2 describes the dynamic discrete choice game used in the paper, introduces the structure of the estimation problem and the main assumptions, and provides an illustrative example of the econometric model. Section 3 studies the asymptotic properties of the K-ML estimator. Section 4 introduces the K-MD estimation method, relates it to the K-ML method, and studies its asymptotic distribution. Section 5 presents results of Monte Carlo simulation and Section 6 concludes. The appendix of the paper collects all the proofs and intermediate results.

2 Setup

This section describes the econometric model, introduces the estimator and the assumptions, and provides an illustrative example.

2.1 Econometric model

We consider a standard dynamic discrete-choice game as described in Aguirregabiria and Mira (2007) or Pesendorfer and Schmidt-Dengler (2008).

The game has discrete time $t = 1, \ldots, T \equiv \infty$, and a finite set of players indexed by $j \in J \equiv \{1, \ldots, |J|\}$. In each period t, every player j observes a vector of state variables s_{jt} and chooses an action a_{jt} from a finite and common set of actions $A \equiv \{0, 1, \ldots, |A| - 1\}$ (with $|A| > 1$) with the objective of maximizing his expected discounted utility. The action denoted by 0 is referred to as the outside option, and we denote $\tilde{A} \equiv \{1, \ldots, |A| - 1\}$. All players choose their action simultaneously and independently upon observation of state variables.

The vector of state variables s_{jt} is composed of two subvectors x_t and ϵ_{jt}. The subvector $x_t \in X \equiv \{1, \ldots, |X|\}$ represents a state variable observed by all other players and the researcher, whereas the subvector $\epsilon_{jt} \in \mathbb{R}^{|A|}$ represents an action-specific state vector only observed by player j. We denote $\epsilon_t \equiv \{\epsilon_{jt} : j \in J\} \in \mathbb{R}^{|A| \times |J|}$ and $\vec{a}_t \equiv \{a_{jt} : j \in J\} \in A^{|J|}$.

We assume that ϵ_{jt} is randomly drawn from a strictly monotonic and continuous density function $dF_\epsilon(\cdot|x_t)$ with full support, that it is independent of ϵ_s for $s < t$ and $a_{\tilde{t}}$ for $\tilde{t} < t$, and also that $E[\epsilon_{jt}|\epsilon_{jt} \geq e]$ exists for all $e \in \mathbb{R}^{|A|}$. We also assume that the observed state variables behave according to a probability density function $dF_x(x_{t+1}|\vec{a}_t, x_t)$ that specifies the probability that the future observed state variable is x_{t+1} given that the actions are $\vec{a}_t$ and the current state is x_t. It then follows that $s_{t+1} = (x_{t+1}, \epsilon_{t+1})$ is a Markov process with a probability density $d\Pr(x_{t+1}, \epsilon_{t+1}|x_t, \epsilon_t, \vec{a}_t)$ that satisfies:

$$d\Pr(x_{t+1}, \epsilon_{t+1}|x_t, \epsilon_t, \vec{a}_t) \;=\; dF_\epsilon(\epsilon_{t+1}|x_{t+1}) \;\times\; dF_x(x_{t+1}|\vec{a}_t, x_t).$$

Every player j has a time-separable utility and discounts future payoffs by $\beta_j \in (0, 1)$. The period t payoff is received after every player made their choices and is given by:

$$\pi_j(\vec{a}_t, x_t) \;+\; \sum_{k \in A} \epsilon_{jt}(k)\; 1[a_{jt} = k].$$

Following the literature, we assume Markov perfect equilibrium (MPE) as the equilibrium concept for the game. By definition, a MPE is a collection of strategies and beliefs for each player such that each player has: (a) rational beliefs, (b) an optimal strategy given his beliefs and other players' choices, and (c) Markovian strategies. According to Pesendorfer and Schmidt-Dengler (2008, Theorem 1), this model has an MPE and, in fact, it could have multiple MPEs (e.g., see Pesendorfer and Schmidt-Dengler (2008, Sections 2 and 7)). We follow Aguirregabiria and Mira (2007) and assume that data come from one of the MPEs in which every player uses pure strategies.[3]

By definition, the MPE is a collection of equilibrium strategies and common beliefs. We denote the probability that player j will choose action $a \in A$ given observed state x by $P_j^*(a|x)$, and we denote $P^* \equiv \{P_j^*(a|x) : (j, a, x) \in J \times \tilde{A} \times X\}$. Note that beliefs only need to be specified in $\tilde{A}$ for every $(j, x) \in J \times X$,

[3]As explained in Aguirregabiria and Mira (2007, footnote 3), focusing on pure strategies can be rationalized by Harsanyi's Purification Theorem.

as $P_j^*(0|x) = 1 - \sum_{a\in\tilde{A}} P_j^*(a|x)$. We denote player j's equilibrium strategy by $\{a_j^*(e,x) : (e,x) \in \mathbb{R}^{|A|} \times X\}$, where $a_j^*(e,x)$ denotes player j's optimal choice when the current private shock is e and the observed state is x. Given that equilibrium strategies are time-invariant, we can abstract from calendar time for the remainder of the paper, and denote $\vec{a} = \vec{a}_t$, $\vec{a}' = \vec{a}_{t+1}$, $x = x_t$, and $x' = x_{t+1}$.

We use $\theta^* \in \Theta$ to denote the finite-dimensional parameter vector that collects the model elements $(\{\pi_j : j \in J\}, \{\beta_j : j \in J\}, dF_\epsilon, dF_x)$. Throughout this paper, we split the parameter vector as follows:

$$\theta^* \;\equiv\; (\alpha^*, g^*) \;\in\; \Theta \;\equiv\; \Theta_\alpha \times \Theta_g, \tag{2.1}$$

where $\alpha^* \in \Theta_\alpha$ denotes a parameter vector of interest that is estimated iteratively and $g^* \in \Theta_g$ denotes a nuisance parameter vector that is estimated directly from the data. In practice, structural parameters that determine the payoff functions $\{\pi_j : j \in J\}$ or the distribution dF_ε usually belong to α^*, while the transition probability density function dF_x is typically part of g^*.[4]

We now describe a fixed point mapping that characterizes equilibrium beliefs in any MPE. Let $P = \{P_j(a|x) : (j,a,x) \in J \times \tilde{A} \times X\}$ denote a set of beliefs, which need not be optimal. Given said beliefs, the ex-ante probability that player j chooses equilibrium action a given observed state x is:

$$\Psi_j(a,x;\alpha^*,g^*,P) \equiv \int_\epsilon \prod_{k\in A} 1[u_j(a,x,\alpha^*,g^*,P) + \epsilon_j(a) \geq u_j(k,x,\alpha^*,g^*,P) + \epsilon_j(k)] dF_\epsilon(\epsilon|x), \tag{2.2}$$

where $u_j(a,x,\alpha^*,g^*,P)$ denotes player j's continuation value net of the pay-off shocks under action a, state variable x, and with beliefs P. In turn,

$$u_j(a,x,\alpha^*,g^*,P) \equiv \sum_{\tilde{a}\in A^{|J|-1}} 1[\vec{a} = (a,\tilde{a})] \prod_{s\in J\setminus\{j\}} P_s(\tilde{a}_s|x)[\pi_j((a,\tilde{a}),x) + \beta \sum_{x'\in X} dF_x(x'|(a,\tilde{a}),x) V_j(x';P)],$$

where $\prod_{s\in J\setminus\{j\}} P_s(\tilde{a}_s|x)$ denotes the beliefs that the remaining players choose $\tilde{a} \equiv \{\tilde{a}_s : s \in J\setminus\{j\}\}$ conditional on x, and $V_j(x;P)$ is player j's ex-ante value function conditional on x.[5] By stacking up this mapping for all decisions and states $(a,x) \in \tilde{A} \times X$ and all players $j \in J$, we define the probability mapping $\Psi(\alpha^*,g^*,P) \equiv \{\Psi_j(a,x;\alpha^*,g^*,P) : (j,a,x) \in J \times \tilde{A} \times X\}$. Given any set of beliefs P (optimal or not), $\Psi(\alpha^*,g^*,P)$ indicates the corresponding optimal CCPs. Once again, note that $\Psi(\alpha^*,g^*,P)$ only needs to be specified in $\tilde{A}$ for every $(j,x) \in X \times J$, as $\Psi_j(0,x;\alpha^*,g^*,P) = 1 - \sum_{a\in\tilde{A}} \Psi(a,x;\alpha^*,g^*,P)$.

Aguirregabiria and Mira (2007, Representation Lemma) and Pesendorfer and Schmidt-Dengler (2008, Proposition 1) both show that the mapping Ψ fully characterizes equilibrium beliefs P^* in the MPE. That is, P^* is an equilibrium belief if and only if:

$$P^* \;=\; \Psi(\alpha^*,g^*,P^*), \tag{2.3}$$

The goal of the paper is to study the problem of inference of $\alpha^* \in \Theta_\alpha$ based on the fixed point equilibrium condition in Eq. (2.3).

[4]Note that the distinction between components of θ^* is without loss of generality, as one can choose to estimate all parameters iteratively by setting $\theta^* = \alpha^*$. The goal of estimating a nuisance parameter g^* directly from the data is to simplify the computation of the iterative procedure by reducing its dimensionality.

[5]The ex-ante value function is the discounted sum of future payoffs in the MPE given x and before players observe shocks and choose actions. It can be computed with the mapping valuation operator defined in Aguirregabiria and Mira (2007, Eqs. 10 and 14) or Pesendorfer and Schmidt-Dengler (2008, Eqs. 5 and 6).

2.2 Estimation procedure

The researcher estimates $\theta^* = (\alpha^*, g^*) \in \Theta \equiv \Theta_\alpha \times \Theta_g$ using a two-step and K-stage PI estimator. For any $K \in \mathbb{N}$, this estimator is defined as follows:

- **Step 1:** Estimate (g^*, P^*) with preliminary estimators $(\hat{g}, \hat{P}_0)$. We also refer to $\hat{P}_0$ as the 0-step estimator of the CCPs.
- **Step 2:** Estimate α^* with $\hat{\alpha}_K$, computed using the following algorithm. Initialize $k = 1$ and then:

  (a) Compute:

  $$\hat{\alpha}_k \quad \equiv \quad \arg\max_{\alpha \in \Theta_\alpha} \hat{Q}_k(\alpha, \hat{g}, \hat{P}_{k-1}), \tag{2.4}$$

  where $\hat{Q}_k : \Theta_\alpha \times \Theta_g \times \Theta_P \to \mathbb{R}$ is the k-th step sample objective function. If $k = K$, exit the algorithm. If $k < K$, go to (b).

  (b) Estimate P^* with the k-step estimator of the CCPs, given by:

  $$\hat{P}_k \quad \equiv \quad \Psi(\hat{\alpha}_k, \hat{g}, \hat{P}_{k-1}).$$

  Then, increase k by one unit and return to (a).

Throughout this paper, we consider α^* to be our main parameter of interest, while g^* is a nuisance parameter. For any $K \in \mathbb{N}$, the two-step and K-stage PI estimator of α^* is given by:

$$\hat{\alpha}_K \quad \equiv \quad \arg\max_{\alpha \in \Theta_\alpha} \hat{Q}_K(\alpha, \hat{g}, \hat{P}_{K-1}), \tag{2.5}$$

and the corresponding estimator of $\theta^* = (\alpha^*, g^*)$ is $\hat{\theta}_K = (\hat{\alpha}_K, \hat{g})$.

The algorithm does not specify the first-step estimators $(\hat{g}, \hat{P}_0)$ or the sequence of sample criterion functions $\{\hat{Q}_k : k \leq K\}$. Rather than determining these objects now, we restrict them by making assumptions in the next subsection. This allows us to use our framework to obtain results for multiple types of estimators, e.g., K-ML and K-MD for any $K \geq 1$, and several possible choices of preliminary estimators $(\hat{g}, \hat{P}_0)$.

We conclude this subsection by describing an alternative estimator that we refer to as the single-step K-stage PI estimator. By this, we refer to an estimation procedure in which the estimation of g^* is removed from the first step and incorporated into the parameter vector estimated in the second step. Note that this is the framework considered in Aguirregabiria and Mira (2007). Section A.3 in the Appendix describes this estimation procedure in detail and studies its properties. As we show there, every result obtained in this paper can be adapted to single-step estimation after suitable relabelling the parameters of the problem.

2.3 Assumptions

This section introduces the main assumptions used in our analysis. As explained in Section 2.1, the game has a MPE, but this need not be unique. To address this issue, we follow most papers in the literature and assume that the data come from a single MPE. The researcher observes an i.i.d. sample of the current state, the current actions, and the future state. To this end, we impose the following assumption.

Assumption A.1. (I.i.d. from one MPE) *The data $\{\{(\{a_{jt,i} : j \in J\}, x_{t,i}, x'_{t,i})\} : i \leq n\}$ are an i.i.d. sample from a single MPE. This MPE determines the data generating process (DGP) denoted by*

$\Pi^* \equiv \{\Pi_j^*(a,x,x') : (j,a,x,x') \in J \times A \times X \times X\}$, *where* $\Pi_j^*(a,x,x')$ *denotes the probability that player* j *chooses action* a *and the common state variable evolves from* x *to* x', *i.e.,*

$$\Pi_j^*(a,x,x') \;\equiv\; \Pr[\ (a_{jt}, x_t, x_{t+1}) = (a,x,x')\ ].$$

See Aguirregabiria and Mira (2007, Assumptions 5(A) and 5(D))) for a similar condition. The observations in the i.i.d. sample are indexed by $i = 1,\ldots,n$. Depending on the application, the index i can be used to denote different time periods (i.e., $i = \tau$ for some $\tau \in \{1,\ldots,\infty\}$, and so $(a_{jt,i}, x_{t,i}, x'_{t,i}) = (a_{j\tau}, x_\tau, x'_\tau)$ as in Pesendorfer and Schmidt-Dengler (2008)) or to denote different markets (as in Aguirregabiria and Mira (2007)). By Assumption A.1, the data identify the DGP, i.e., $\Pi_j^*(a,x,x')$ for every $(j,a,x,x') \in J \times A \times X \times X$. In turn, the DGP identifies the equilibrium CCPs, transition probabilities, and marginal distributions. To see why, note that for all $(j,\vec{a},x) \in J \times A^{|J|} \times X$,

$$\begin{aligned}
P_j^*(\vec{a}_j|x) &\equiv \frac{\sum_{x'\in X} \Pi_j^*(\vec{a}_j,x,x')}{\sum_{(\tilde{a},\tilde{x}')\in A\times X} \Pi_j^*(\tilde{a},x,\tilde{x}')}\\
\Lambda^*(x'|x,\vec{a}) &\equiv \frac{\prod_{j\in J} \Pi_j^*(\vec{a}_j,x,x')}{\sum_{(\breve{a},\breve{x})\in A^{|J|}\times X} \prod_{j\in J} \Pi_j^*(\breve{a}_j,\breve{x},x')}\\
m^*(x) &\equiv \sum_{(a,x')\in A\times X} \Pi_j^*(a,x,x'),
\end{aligned}$$

where $P_j^*(\vec{a}_j|x)$ denotes the probability that player j will choose action $\vec{a}_j$ given that the observed state is x, $\Lambda^*(x'|x,\vec{a})$ denotes the probability that the future state observed state is x' given that the current observed state is x and the action vector is $\vec{a}$, respectively, and $m^*(x)$ denotes the (unconditional) probability that the current observed state is x.[6]

Identification of the CCPs, however, is not sufficient for identification of the parameters of interest. To this end, we make the following assumption.

Assumption A.2. (Identification) $\Psi(\alpha, g^*, P^*) = P^*$ *if and only if* $\alpha = \alpha^*$.

This assumption is also imposed in the literature (e.g. Aguirregabiria and Mira (2007, Assumption 5(C)) and Pesendorfer and Schmidt-Dengler (2008, Assumption A4)). The problem of identification is studied in Pesendorfer and Schmidt-Dengler (2008, Section 5). In particular, Pesendorfer and Schmidt-Dengler (2008, Proposition 2) indicate the maximum number of parameters that could be identified from the model and Pesendorfer and Schmidt-Dengler (2008, Proposition 3) provides sufficient conditions for identification.

The K-stage PI estimator is an example of an extremum estimator. The following assumption imposes mild regularity conditions that are typically imposed for these estimators.

Assumption A.3. (Regularity conditions) *Assume the following conditions:*

(i) α^* *belongs to the interior of* Θ_α.

(ii) $\sup_{\alpha\in\Theta_\alpha} |\Psi(\alpha,\tilde{g},\tilde{P}) - \Psi(\alpha,g^*,P^*)| = o_p(1)$, *provided that* $(\tilde{g},\tilde{P}) = (g^*,P^*) + o_p(1)$.

(iii) $\inf_{\alpha\in\Theta_\alpha} \Psi_{jax}(\alpha,\tilde{g},\tilde{P}) > 0$ *for all* $(j,a,x) \in J \times A \times X$, *provided that* $(\tilde{g},\tilde{P}) = (g^*,P^*) + o_p(1)$.

(iv) $\Psi(\alpha,g,P)$ *is twice continuously differentiable in a neighborhood of* (α^*,g^*,P^*). *We use* $\Psi_\lambda \equiv \partial\Psi(\alpha^*,g^*,P^*)/\partial\lambda$ *for* $\lambda \in \{\alpha,g,P\}$.

[6]While the definition of $m^*(x)$ is function of the player's identity $j \in J$, all players should be in agreement in a MPE.

(v) $[\mathbf{I}_{d_P} - \Psi_P, -\Psi_g] \in \mathbb{R}^{d_P \times (d_P+d_g)}$ *and* $\Psi_\alpha \in \mathbb{R}^{d_P \times d_\alpha}$ *are full rank matrices.*

Assumption A.3(i) allows us to characterize $\hat{\alpha}_K$ using the first order condition produced by Eq. (2.5). It is imposed by Aguirregabiria and Mira (2007) and Pesendorfer and Schmidt-Dengler (2008, Assumption A2). Assumptions A.3(ii)-(iii) are related to the uniform consistency requirement used to establish the consistency of extremum estimators. Assumption A.3(iv) is a standard condition used to prove the asymptotic normality of extremum estimators and is related to Pesendorfer and Schmidt-Dengler (2008, Assumption A5). Finally, Assumption A.3(v) is a mild rank condition similar to Pesendorfer and Schmidt-Dengler (2008, Assumption A7).

We next discuss introduce assumptions on $(\hat{g}, \hat{P}_0)$, i.e., the preliminary estimators of (g^*, P^*). As a reference estimator for P^*, we consider sample frequency CCP estimator, given by:

$$\hat{P} \;\equiv\; \{\hat{P}_j(a|x) : (j,a,x) \in J \times \tilde{A} \times X\}, \tag{2.6}$$

with

$$\hat{P}_j(a|x) \equiv \frac{\sum_{i=1}^n \sum_{\tilde{x}' \in X} 1[(a_{jt,i}, x_{t,i}, x'_{t,i}) = (a, x, \tilde{x}')]/n}{\sum_{i=1}^n \sum_{(\check{a},\check{x}') \in A \times X} 1[(a_{jt,i}, x_{t,i}, x'_{t,i}) = (\check{a}, x, \check{x}')]/n}.$$

Under Assumption A.1, $\hat{P}$ is the maximum likelihood estimator (MLE) of P^*, and it satisfies:

$$\sqrt{n}(\hat{P} - P^*) \;\overset{d}{\to}\; N(0, \Omega_{PP}),$$

where $\Omega_{PP} \equiv diag\{\Sigma_{jx} : (j,x) \in J \times X\}$, and $\Sigma_{xj} \equiv (diag\{P^*_{jx}\} - P^*_{jx}P^{*\prime}_{jx})/m^*(x)$ and $P^*_{jx} \equiv \{P^*_{jax} : a \in \tilde{A}\}$ for all $(j,x) \in J \times X$. By standard arguments about the MLE, $\hat{P}$ is an efficient estimator of P^*; see Amemiya (1985, Section 4.2.4).

Rather than imposing specific estimators for $(\hat{g}, \hat{P}_0)$, we entertain two high-level assumptions.

Assumption A.4. (Baseline convergence) $(\hat{P}, \hat{P}_0, \hat{g})$ *satisfies the following condition:*

$$\sqrt{n}\begin{pmatrix} \hat{P} - P^* \\ \hat{P}_0 - P^* \\ \hat{g} - g^* \end{pmatrix} \;\overset{d}{\to}\; N\left(\begin{pmatrix} \mathbf{0}_{d_P} \\ \mathbf{0}_{d_P} \\ \mathbf{0}_{d_g} \end{pmatrix}, \begin{pmatrix} \Omega_{PP} & \Omega_{P0} & \Omega_{Pg} \\ \Omega'_{P0} & \Omega_{00} & \Omega_{0g} \\ \Omega'_{Pg} & \Omega'_{0g} & \Omega_{gg} \end{pmatrix}\right).$$

Assumption A.5. (Baseline convergence II) $(\hat{P}, \hat{P}_0, \hat{g})$ *satisfies the following conditions.*

(i) The asymptotic variance of $(\hat{P}, \hat{g})$ *is nonsingular.*

(ii) For any $M \in \mathbb{R}^{d_P \times d_P}$, $((\mathbf{I}_{d_P} - M)\hat{P} + M\hat{P}_0, \hat{g})$ *is not asymptotically more efficient than* $(\hat{P}, \hat{g})$.

Assumption A.4 imposes the consistency and the joint asymptotic normality of $(\hat{P}, \hat{P}_0, \hat{g})$. Assumption A.5(i) requires that the asymptotic variance of $(\hat{P}, \hat{g})$ is nonsingular. This condition is implicitly required by the definition of the optimal weight matrix in Pesendorfer and Schmidt-Dengler (2008, Proposition 5). To interpret Assumption A.5(ii), recall that $\hat{P}$ is an asymptotically efficient estimator of P^*, and so $(\mathbf{I}_{d_P} - M)\hat{P} + M\hat{P}_0$ cannot be asymptotically more efficient than $\hat{P}$. Assumption A.5(ii) requires that this conclusion applies also when these estimators are coupled with $\hat{g}$.

To illustrate Assumptions A.4 and A.5, we need to specify the parameter vector g^*. A common specification in the literature is the one in Pesendorfer and Schmidt-Dengler (2008), where g^* is the vector of state transition probabilities, i.e.,

$$g^* \;=\; \{\Pr(x'|a,x) : (a,x,x') \in A^{|J|} \times X \times \tilde{X}\},$$

and $\tilde{X} \equiv \{1,\ldots,|X|-1\}$ is the state space with the last action removed (to avoid redundancy). In this case, it is reasonable to propose sample frequencies estimators for $(\hat{P}_0, \hat{g})$, i.e., $\hat{P}_0 = \hat{P}$ and $\hat{g} \equiv \{\hat{g}(a,x,x') : (a,x,x') \in A^{|J|} \times X \times \tilde{X}\}$ with

$$\hat{g}(a,x,x') \;\equiv\; \frac{\sum_{i=1}^{n} 1[(\vec{a}_{t,i}, x_{t,i}, x'_{t,i}) = (a,x,\tilde{x}')]/n}{\sum_{i=1}^{n}\sum_{(\check{a},\check{x}') \in A^{|J|}\times X} 1[(a_{t,i}, x_{t,i}, x'_{t,i}) = (\check{a}, x, \check{x}')]/n}.$$

Then, $(\hat{P}_0, \hat{g})$ is the maximum likelihood estimator of (P^*, g^*), and standard arguments imply Assumptions A.4 and A.5. Note also that we would obtain the same conclusions if we replaced $(\hat{P}_0, \hat{g})$ with any asymptotically equivalent estimator, i.e., any estimator $(\tilde{P}_0, \tilde{g})$ such that $(\tilde{P}_0, \tilde{g}) = (\hat{P}_0, \hat{g}) + o(n^{-1/2})$. Examples of asymptotically equivalent estimators would be the ones resulting from the relaxation method in Kasahara and Shimotsu (2012) or the "undersmoothed" kernel estimator in Grund (1993, Theorem 5.3).

We conclude the section with a comment on single-step estimation. As mentioned earlier, single-step estimation case can be captured in our framework by eliminating the parameter g^* from the estimation problem. For the sake of completeness, Section A.3 of the Appendix shows how to adapt Assumptions A.1-A.5 to the single-step estimation case.

2.4 An illustrative example

We illustrate the framework with the two-player version dynamic entry game in Aguirregabiria and Mira (2007, Example 5). In each period $t = 1, \ldots, T \equiv \infty$, two firms indexed by $j \in J = \{1, 2\}$ simultaneously decide whether to enter or not into the market, upon observation of the state variables. Firm j's decision at time t is $a_{jt} \in A = \{0, 1\}$, which takes value one if firm j enters the market at time t, and zero otherwise. In each period t, the vector of state variables observed by firm j is $s_{jt} = (x_{jt}, \epsilon_{jt})$, where $\epsilon_{jt} = (\epsilon_{jt0}, \epsilon_{jt1}) \in \mathbb{R}^2$ represents the privately-observed vector of action-specific state variables and $x_{jt} = x_t \in X \equiv \{1, 2, 3, 4\}$ is a publicly-observed state variable that indicates the entry decisions in the previous period, i.e.,

$$x_{jt} \;=\; \left[\begin{array}{c} 1[(a_{1,t-1}, a_{2,t-1}) = (0,0)] + 2 \times 1[(a_{1,t-1}, a_{2,t-1}) = (0,1)] + \\ 3 \times 1[(a_{1,t-1}, a_{2,t-1}) = (1,0)] + 4 \times 1[(a_{1,t-1}, a_{2,t-1}) = (1,1)] \end{array}\right].$$

We specify the profit function as in Aguirregabiria and Mira (2007, Eq. (48)). If firm j enters the market in period t, its period profits are:

$$\pi_j((1, a_{-j,t}), x_t) \;=\; \lambda^*_{RS} - \lambda^*_{RN} \ln(1 + a_{-j,t}) - \lambda^*_{FC,j} - \lambda^*_{EC}(1 - a_{j,t-1}) + \epsilon_{jt1}, \tag{2.7}$$

where λ^*_{RS} represents fixed entry profits, λ^*_{RN} represents the effect of a competitor's entry, $\lambda^*_{FC,j}$ represents a firm-specific fixed cost, and λ^*_{EC} represents the entry. On the other hand, if firm j does not enter the market in period t, its period profits are:

$$\pi_j((0, a_{-j,t}), x_t) \;=\; \epsilon_{jt0}.$$

Firms discount future profits at a common discount factor $\beta^* \in (0,1)$.

We assume that $\epsilon_{j,a,t}$ is i.i.d. drawn from an extreme value distribution with unit dispersion, i.e.,

$$dF_\epsilon(\epsilon_t = e_t|x_t) \;=\; \prod_{a=0}^{1}\prod_{j=1}^{2}\exp(-\exp(-e_{jta})),$$

Finally, since x_{t+1} is uniquely determined by $\vec{a}_t$,

$$dF_x(x_{t+1}|\vec{a}_t, x_t) \;=\; 1\left[x_{t+1} = \left[\begin{array}{c} 1[(a_{1,t-1}, a_{2,t-1}) = (0,0)] + 2\times 1[(a_{1,t-1}, a_{2,t-1}) = (0,1)] + \\ 3\times 1[(a_{1,t-1}, a_{2,t-1}) = (1,0)] + 4\times 1[(a_{1,t-1}, a_{2,t-1}) = (1,1)] \end{array}\right]\right].$$

This completes the specification of the econometric model up to $(\lambda^*_{RN}, \lambda^*_{EC}, \lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*)$. These parameters are known to the players but not necessarily known to the researcher.

We will use this econometric model to illustrate our theoretical results and for our Monte Carlo simulations. To ease the computational burden, we presume that the researcher knows that $(\lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*)$, and is interested in estimating $(\lambda^*_{RN}, \lambda^*_{EC})$. In addition, we assume that these parameters are estimated using a single-step PI estimator, i.e., $\theta^* = \alpha^* \equiv (\lambda^*_{RN}, \lambda^*_{EC})$.

3 Results for K-ML estimation

This section provides formal results for the K-ML estimator introduced in Aguirregabiria and Mira (2002, 2007) given an arbitrary number of iteration steps $K \in \mathbb{N}$. The K-ML estimator is defined by Eq. (2.5) with the pseudo log-likelihood criterion function, i.e., $\hat{Q}_K = \hat{Q}_{ML}$. That is,

- **Step 1:** Estimate (g^*, P^*) with preliminary step estimators $(\hat{g}, \hat{P}_0)$.
- **Step 2:** Estimate α^* with $\hat{\alpha}_{K-ML}$, computed using the following algorithm. Initialize $k = 1$ and then:

  (a) Compute:

  $$\hat{\alpha}_{k-ML} \;\equiv\; \arg\min_{\alpha\in\Theta_\alpha} \frac{1}{n}\sum_{i=1}^{n} \ln \Psi(\alpha, \hat{g}, \hat{P}_{k-1})(a_i|x_i).$$

  If $k = K$, exit the algorithm. If $k < K$, go to (b).

  (b) Estimate P^* with the k-step estimator of the CCPs, given by:

  $$\hat{P}_k \;\equiv\; \Psi(\hat{\alpha}_{k-ML}, \hat{g}, \hat{P}_{k-1}).$$

  Then, increase k by one unit and return to (a).

As explained in Section 1, The K-ML estimator is the K-stage PI estimator introduced by Aguirregabiria and Mira (2002) for dynamic single-agent problems and Aguirregabiria and Mira (2007) for dynamic games. Aguirregabiria and Mira (2007) study the asymptotic behavior of $\hat{\alpha}_{K-ML}$ for two extreme values of K: $K = 1$ and K large enough to induce the convergence of the estimator. Under some conditions, they show that iterating the K-ML estimator until convergence produces asymptotic efficiency gains.

The results in Aguirregabiria and Mira (2007) open certain avenues for further research. First, they focus on these two extreme values of K, without considering other possible values. Second, they restrict attention to single-step K-ML estimators for simplicity of the analysis. In Theorem 3.1, we complement the analysis in Aguirregabiria and Mira (2007) along these two dimensions. In particular, we derive the asymptotic distribution of the two-step K-ML estimator for any $K \geq 1$.

Theorem 3.1 (Two-step K-ML). *Fix $K \geq 1$ arbitrarily and assume Assumptions A.1-A.4. Then,*

$$\sqrt{n}(\hat{\alpha}_{K-ML} - \alpha^*) \quad \overset{d}{\to} \quad N(\mathbf{0}_{d_\alpha \times 1}, \Sigma_{K-ML}(\hat{P}_0, \hat{g})),$$

where

$$\Sigma_{K-ML}(\hat{P}_0) \equiv \left\{ \begin{array}{c} (\Psi_\alpha' \Omega_{PP}^{-1} \Psi_\alpha)^{-1} \Psi_\alpha' \Omega_{PP}^{-1} \times \\ \left[\begin{array}{c} (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P})' \\ -(\Psi_P \Phi_{K,0})' \\ -\Psi_g'(\mathbf{I}_{d_P} + \Psi_P \Phi_{K,g})' \end{array} \right]' \left(\begin{array}{ccc} \Omega_{PP} & \Omega_{P0} & \Omega_{Pg} \\ \Omega_{P0}' & \Omega_{00} & \Omega_{0g} \\ \Omega_{Pg}' & \Omega_{0g}' & \Omega_{gg} \end{array} \right) \left[\begin{array}{c} (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P})' \\ -(\Psi_P \Phi_{K,0})' \\ -\Psi_g'(\mathbf{I}_{d_P} + \Psi_P \Phi_{K,g})' \end{array} \right] \\ \times \Omega_{PP}^{-1} \Psi_\alpha (\Psi_\alpha' \Omega_{PP}^{-1} \Psi_\alpha)^{-1} \end{array} \right\},$$

and $\{\Phi_{k,P} : k \leq K\}$, $\{\Phi_{k,0} : k \leq K\}$, and $\{\Phi_{k,g} : k \leq K\}$ are defined as follows. Set $\Phi_{1,P} \equiv \mathbf{0}_{d_P \times d_P}$, $\Phi_{1,0} \equiv \mathbf{I}_{d_P}$, $\Phi_{1,g} \equiv \mathbf{0}_{d_P \times d_P}$ and, for any $k \leq K-1$,

$$\begin{aligned} \Phi_{k+1,P} &\equiv (\mathbf{I}_{d_P} - \Psi_\alpha (\Psi_\alpha' \Omega_{PP}^{-1} \Psi_\alpha)^{-1} \Psi_\alpha' \Omega_{PP}^{-1}) \Psi_P \Phi_{k,P} + \Psi_\alpha (\Psi_\alpha' \Omega_{PP}^{-1} \Psi_\alpha)^{-1} \Psi_\alpha' \Omega_{PP}^{-1}, \\ \Phi_{k+1,0} &\equiv (\mathbf{I}_{d_P} - \Psi_\alpha (\Psi_\alpha' \Omega_{PP}^{-1} \Psi_\alpha)^{-1} \Psi_\alpha' \Omega_{PP}^{-1}) \Psi_P \Phi_{k,0}, \\ \Phi_{k+1,g} &\equiv (\mathbf{I}_{d_P} - \Psi_\alpha (\Psi_\alpha' \Omega_{PP}^{-1} \Psi_\alpha)^{-1} \Psi_\alpha' \Omega_{PP}^{-1}) \Psi_P \Phi_{k,g} + (\mathbf{I}_{d_P} - \Psi_\alpha (\Psi_\alpha' \Omega_{PP}^{-1} \Psi_\alpha)^{-1} \Psi_\alpha' \Omega_{PP}^{-1}). \end{aligned} \tag{3.1}$$

There are several important comments regarding this result. First, Theorem 3.1 considers $K \geq 1$ but *fixed* as $n \to \infty$. Because of this, our asymptotic framework is not subject to the criticism raised by Pesendorfer and Schmidt-Dengler (2010). Second, we can deduce an analogous result for the single-step K-ML estimator by eliminating g^* from the estimation problem. See Theorem A.2 in the appendix for the corresponding result. Finally, note that we can consistently estimate $\Sigma_{K-ML}(\hat{P}_0, \hat{g})$ for any $K \geq 1$ based on consistent estimators of the asymptotic variance in Assumption A.4 and the parameter vector (α^*, g^*) (e.g., $(\hat{\alpha}_{1-ML}, \hat{g})$).

Theorem 3.1 reveals that the K-ML estimator of α^* is consistent and asymptotically normally distributed for all $K \geq 1$. Thus, the asymptotic mean squared error of the K-ML estimator is equal to its asymptotic variance, $\Sigma_{K-ML}(\hat{P}_0)$. The goal for the rest of the section is to investigate how this asymptotic variance changes with the number of iterations K.

In single-agent dynamic problems, Aguirregabiria and Mira (2002) show that the so-called zero Jacobian property holds, i.e., $\Psi_P = \mathbf{0}_{dP \times dP}$. If we plug in this information into Theorem 3.1, we conclude the asymptotic variance of the K-ML estimator is given by:

$$\Sigma_{K-ML}(\hat{P}_0) = \; (\Psi_\alpha' \Omega_{PP}^{-1} \Psi_\alpha)^{-1} \Psi_\alpha' \Omega_{PP}^{-1} (\Omega_{PP} + \Psi_g \Omega_{gg} \Psi_g' - \Psi_g \Omega_{Pg}' - \Omega_{Pg} \Psi_g') \Omega_{PP}^{-1} \Psi_\alpha (\Psi_\alpha' \Omega_{PP}^{-1} \Psi_\alpha)^{-1} \; .$$

This expression is invariant to K, corresponding to the main finding in Aguirregabiria and Mira (2002).

In multiple-agent dynamic problems, however, the zero Jacobian property no longer holds. Theorem 3.1 reveals that the asymptotic variance of the K-ML estimator is a complicated function of the number of iteration steps K. We illustrate this complexity using the example of Section 2.4. In this example, the researcher is interested in estimating $(\lambda_{RN}^*, \lambda_{EC}^*)$ and uses single-step estimator. For simplicity, we set

$\hat{P}_0 = \hat{P}$. In this context, the asymptotic variance of $\hat{\alpha}_{K-ML} = (\hat{\lambda}_{RN,K-ML}, \hat{\lambda}_{EC,K-ML})$ is given by:

$$\Sigma_{K-ML}(\hat{P}_0) = (\Psi_\alpha' \Omega_{PP}^{-1} \Psi_\alpha)^{-1} \Psi_\alpha' \Omega_{PP}^{-1} (\mathbf{I}_{d_P} - \Psi_P \Phi_{k,P0}) \Omega_{PP} (\mathbf{I}_{d_P} - \Psi_P \Phi_{k,P})' \Omega_{PP}^{-1} \Psi_\alpha (\Psi_\alpha' \Omega_{PP}^{-1} \Psi_\alpha)^{-1}, \quad (3.2)$$

where $\{\Phi_{k,P0} : k \leq K\}$ is defined by $\Phi_{k,P0} \equiv \Phi_{k,P} + \Phi_{k,0}$, with $\{\Phi_{k,P} : k \leq K\}$ and $\{\Phi_{k,0} : k \leq K\}$ as in Eq. (3.1). For any true parameter vector and any $K \in \mathbb{N}$, we can numerically compute Eq. (3.2). For exposition, we focus on the [1,1]-element of $\Sigma_{K-ML}(\hat{P}_0)$, which corresponds to the asymptotic variance of $\hat{\lambda}_{RN,K-ML}$. Figures 1, 2, and 3 show the asymptotic variance of $\hat{\lambda}_{RN,K-ML}$ as a function of K for $(\lambda_{RN}^*, \lambda_{EC}^*, \lambda_{RS}^*, \lambda_{FC,1}^*, \lambda_{FC,2}^*, \beta^*) = (2.8, 0.8, 0.7, 0.6, 0.4, 0.95)$, $(2, 1.8, 0.2, 0.1, 0.3, 0.95)$, and $(2.2, 1.45, 0.45, 0.22, 0.29, 0.95)$, respectively. These figures confirm that, in general, the asymptotic variance of the K-ML estimator can decrease, increase, or even wiggle with the number of iterations K. Note that these widely different patterns occur within the same econometric model.

We view the fact that $\Sigma_{K-ML}(\hat{P}_0)$ can change so much with the number of iterations K as a negative feature of the K-ML estimator. A researcher who uses the K-ML estimator and desires asymptotic efficiency faces difficulties when choosing K. Prior to estimation, the researcher cannot be certain regarding the effect of K on the asymptotic efficiency of the K-ML estimator. Additional iterations could help asymptotic efficiency (as in Figure 1) or hurt asymptotic efficiency (as in Figure 2). In principle, the researcher can consistently estimate the asymptotic variance of K-ML for each K by plugging in any consistent estimator of the structural parameters (e.g., $\hat{\alpha}_{1-ML}$ and $\hat{g}$). This idea has two important drawbacks. First, the conclusions are contaminated by sampling error. Second, the entire procedure represents an important computational burden on the researcher.

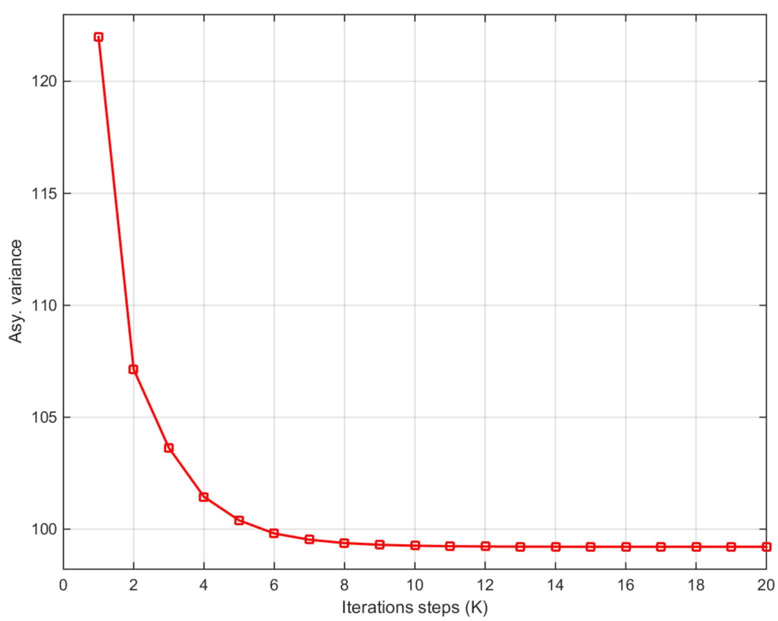

Figure 1: Asymptotic variance of the K-ML estimator of λ_{RN}^* as a function of the number of iterations K when $(\lambda_{RN}^*, \lambda_{EC}^*, \lambda_{RS}^*, \lambda_{FC,1}^*, \lambda_{FC,2}^*, \beta^*) = (2.8, 0.8, 0.7, 0.6, 0.4, 0.95)$.

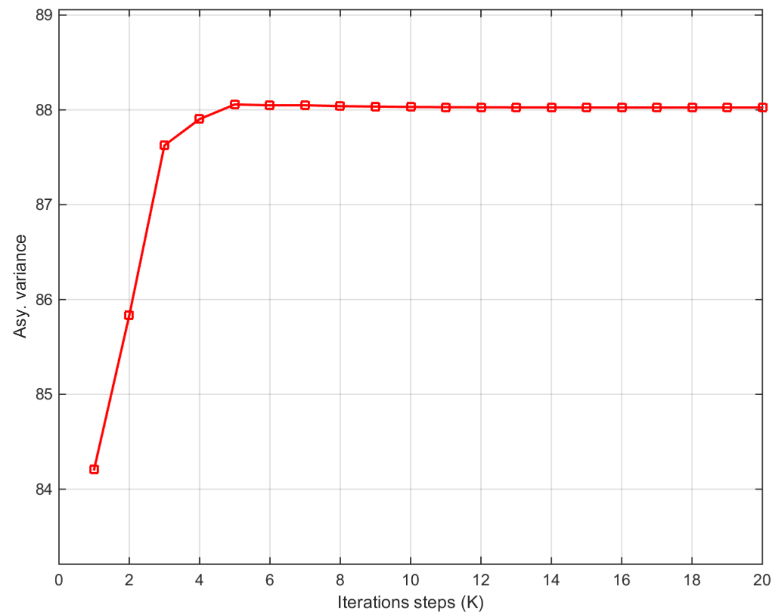

Figure 2: Asymptotic variance of the K-ML estimator of λ^*_{RN} as a function of the number of iterations K when $(\lambda^*_{RN}, \lambda^*_{EC}, \lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*) = (2, 1.8, 0.2, 0.1, 0.3, 0.95)$.

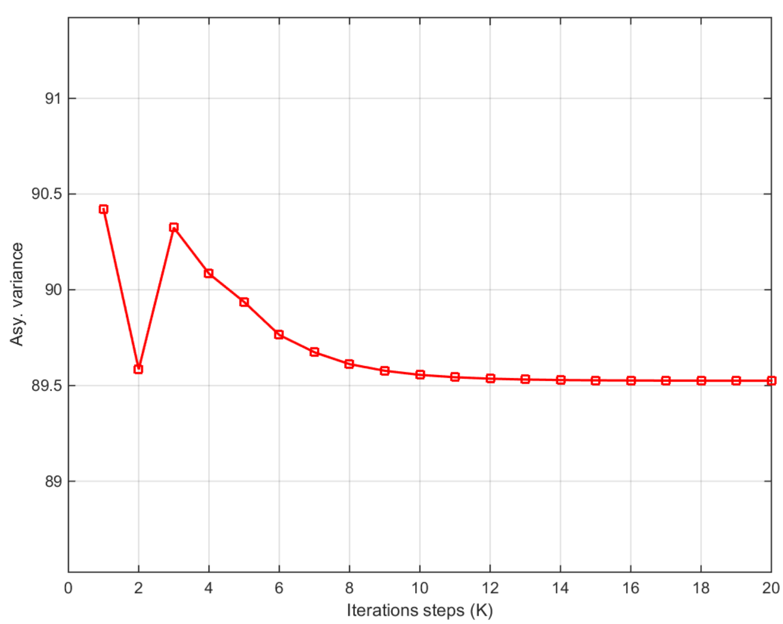

Figure 3: Asymptotic variance of the K-ML estimator of λ^*_{RN} as a function of the number of iterations K when $(\lambda^*_{RN}, \lambda^*_{EC}, \lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*) = (2.2, 1.45, 0.45, 0.22, 0.29, 0.95)$.

4 Results for K-MD estimation

In this section, we introduce a novel K-stage PI estimator, referred to as the K-MD estimator. We demonstrate that this has several advantages over the K-ML estimator. In particular, we show that an optimal K-MD estimator is easy to compute, it dominates the K-ML estimator in asymptotic efficiency, and its asymptotic variance does not change with K.

The K-MD estimator is defined by Eq. (2.5) with (negative) minimum distance criterion function:

$$\hat{Q}_{K-MD}(\alpha, g, P) \;\equiv\; -(\hat{P} - \Psi(\alpha, g, P))'\hat{W}_K(\hat{P} - \Psi(\alpha, g, P)),$$

where $\{\hat{W}_k : k \leq K\}$ is a sequence of positive semidefinite weight matrices. That is,

- **Step 1:** Estimate (g^*, P^*) with preliminary step estimators $(\hat{g}, \hat{P}_0)$.
- **Step 2:** Estimate α^* with $\hat{\alpha}_{K-MD}$, computed using the following algorithm. Initialize $k = 1$ and then:

  (a) Compute:

  $$\hat{\alpha}_{k-MD} \;\equiv\; \underset{\alpha \in \Theta_\alpha}{\arg\min}\; (\hat{P} - \Psi(\alpha, \hat{g}, \hat{P}_{k-1}))'\hat{W}_k(\hat{P} - \Psi(\alpha, \hat{g}, \hat{P}_{k-1})).$$

  If $k = K$, exit the algorithm. If $k < K$, go to (b).

  (b) Estimate P^* with the k-step estimator of the CCPs, given by:

  $$\hat{P}_k \;\equiv\; \Psi(\hat{\alpha}_{k-MD}, \hat{g}, \hat{P}_{k-1}).$$

  Then, increase k by one unit and return to (a).

The implementation of the K-MD estimator requires several choices: the number of iteration steps K and the associated weight matrices $\{\hat{W}_k : k \leq K\}$. For instance, the least squares estimator in Pesendorfer and Schmidt-Dengler (2008) is a particular case of our 1-MD estimator with $\hat{P}_0 = \hat{P}$. In this sense, our K-MD estimator can be considered as an iterative version of their least squares estimator. The primary goal of this section is to study how to make optimal choices of K and $\{\hat{W}_k : k \leq K\}$.

To establish the asymptotic properties of the K-MD estimator, we add the following assumption.

Assumption A.6. (Weight matrices) *$\hat{W}_k \xrightarrow{p} W_k$ for all $k \leq K$, where $W_k \in \mathbb{R}^{d_P \times d_P}$ is positive definite and symmetric for all $k \leq K$.*

The next result derives the asymptotic distribution of the two-step K-MD estimator for any $K \geq 1$.

Theorem 4.1 (Two-step K-MD)**.** *Fix $K \geq 1$ arbitrarily and assume Assumptions A.1-A.4 and A.6. Then,*

$$\sqrt{n}(\hat{\alpha}_{K-MD} - \alpha^*) \;\xrightarrow{d}\; N(\mathbf{0}_{d_\alpha \times 1}, \Sigma_{K-MD}(\hat{P}_0, \{W_k : k \leq K\})),$$

where

$$\Sigma_{K-MD}(\hat{P}_0, \{W_k : k \leq K\}) \;\equiv\; \left\{ \begin{array}{c} (\Psi_\alpha' W_K \Psi_\alpha)^{-1}\Psi_\alpha' W_K \times \\ \left[\begin{array}{c} (\mathbf{I}_{d_P} - \Psi_P\Phi_{K,P})' \\ -(\Psi_P\Phi_{K,0})' \\ -\Psi_g'(\mathbf{I}_{d_P} + \Psi_P\Phi_{K,g})' \end{array}\right]' \left(\begin{array}{ccc} \Omega_{PP} & \Omega_{P0} & \Omega_{Pg} \\ \Omega_{P0}' & \Omega_{00} & \Omega_{0g} \\ \Omega_{Pg}' & \Omega_{0g}' & \Omega_{gg} \end{array}\right) \left[\begin{array}{c} (\mathbf{I}_{d_P} - \Psi_P\Phi_{K,P})' \\ -(\Psi_P\Phi_{K,0})' \\ -\Psi_g'(\mathbf{I}_{d_P} + \Psi_P\Phi_{K,g})' \end{array}\right] \\ \times W_K'\Psi_\alpha(\Psi_\alpha' W_K'\Psi_\alpha)^{-1} \end{array} \right\},$$

and $\{\Phi_{k,0} : k \leq K\}$, $\{\Phi_{k,P} : k \leq K\}$, *and* $\{\Phi_{k,g} : k \leq K\}$ *defined as follows. Set* $\Phi_{1,P} \equiv \mathbf{0}_{d_P \times d_P}$, $\Phi_{1,0} \equiv \mathbf{I}_{d_P}$, $\Phi_{1,g} \equiv \mathbf{0}_{d_P \times d_P}$ *and, for any* $k \leq K-1$,

$$\begin{aligned}\Phi_{k+1,P} &\equiv (\mathbf{I}_{d_P} - \Psi_\alpha(\Psi'_\alpha W_k \Psi_\alpha)^{-1}\Psi'_\alpha W_k)\Psi_P \Phi_{k,P} + \Psi_\alpha(\Psi'_\alpha W_k \Psi_\alpha)^{-1}\Psi'_\alpha W_k,\\ \Phi_{k+1,0} &\equiv (\mathbf{I}_{d_P} - \Psi_\alpha(\Psi'_\alpha W_k \Psi_\alpha)^{-1}\Psi'_\alpha W_k)\Psi_P \Phi_{k,0},\\ \Phi_{k+1,g} &\equiv (\mathbf{I}_{d_P} - \Psi_\alpha(\Psi'_\alpha W_k \Psi_\alpha)^{-1}\Psi'_\alpha W_k)(\mathbf{I}_{d_P} + \Psi_P \Phi_{k,g}).\end{aligned} \tag{4.1}$$

Several important comments are in order. First, as in Theorem 3.1, Theorem 4.1 considers $K \geq 1$ but *fixed* as $n \to \infty$, and it thus is free from the criticism raised by Pesendorfer and Schmidt-Dengler (2010). Second, we can deduce an analogous result for the single-step K-MD estimator by eliminating g^* from the estimation problem, see Theorem A.3 in the appendix. Third, we can consistently estimate $\Sigma_{K-MD}(\hat{P}_0, \hat{g})$ for any $K \geq 1$ based on consistent estimators of the asymptotic variance in Assumption A.4 and the parameter vector (α^*, g^*) (e.g., $(\hat{\alpha}_{1-MD}, \hat{g})$). Finally, note that the asymptotic distribution of the K-ML estimator coincides with that of the K-MD estimator when $W_k \equiv \Omega_{PP}^{-1}$ for all $k \leq K$. We record this result in the following corollary.

Corollary 4.1 (K-ML is a special case of K-MD)**.** *Fix* $K \geq 1$ *arbitrarily and assume Assumptions A.1-A.4. The asymptotic distribution of* K*-ML is a special case of* K*-MD with* $W_k \equiv \Omega_{PP}^{-1}$ *for all* $k \leq K$.

Theorem 4.1 reveals that the asymptotic variance of the K-MD estimator is a complicated function of the number of iteration steps K and sequence of limiting weighting matrices $\{W_k : k \leq K\}$. A natural question to ask is the following: Is there an optimal way of choosing these parameters? In particular, what is the optimal choice of K and $\{W_k : k \leq K\}$ that minimizes the asymptotic variance of the K-MD estimator? We devote the rest of this section to this question.

As a first approach to this problem, we consider the non-iterative 1-MD estimator. As shown in Pesendorfer and Schmidt-Dengler (2008), the asymptotic distribution of this estimator is analogous to that of a GMM estimator so we can leverage well-known optimality results. The next result provides a concrete answer regarding the optimal choices of $\hat{P}_0$ and W_1.

Theorem 4.2 (Optimality with $K = 1$)**.** *Assume Assumptions A.1-A.6. Let* $\hat{\alpha}^*_{1-MD}$ *denote the* 1*-MD estimator with* $\hat{P}_0 = \hat{P}$ *and* $W_1 = W_1^*$ *with:*

$$W_1^* \equiv [(\mathbf{I}_{d_P} - \Psi_P)\Omega_{PP}(\mathbf{I}_{d_P} - \Psi'_P) + \Psi_g \Omega_{gg} \Psi'_g - \Psi_g \Omega'_{Pg}(\mathbf{I}_{d_P} - \Psi'_P) - (\mathbf{I}_{d_P} - \Psi_P)\Omega_{Pg}\Psi'_g]^{-1}. \tag{4.2}$$

Then,

$$\sqrt{n}(\hat{\alpha}^*_{1-MD} - \alpha^*) \overset{d}{\to} N(\mathbf{0}_{d_\alpha \times 1}, \Sigma^*),$$

with

$$\Sigma^* \equiv (\Psi'_\alpha[(\mathbf{I}_{d_P} - \Psi_P)\Omega_{PP}(\mathbf{I}_{d_P} - \Psi'_P) + \Psi_g \Omega_{gg} \Psi'_g - \Psi_g \Omega'_{Pg}(\mathbf{I}_{d_P} - \Psi'_P) - (\mathbf{I}_{d_P} - \Psi_P)\Omega_{Pg}\Psi'_g]^{-1}\Psi_\alpha)^{-1}. \tag{4.3}$$

Furthermore, $\Sigma_{1-MD}(\hat{P}_0, W_1) - \Sigma^*$ *is positive semidefinite for all* $(\hat{P}_0, W_1)$*, i.e.,* $\hat{\alpha}^*_{1-MD}$ *is optimal among all* 1*-MD estimators that satisfy our assumptions.*

Theorem 4.2 indicates that $\hat{P}_0 = \hat{P}$ and $W_1 = W_1^*$ produce an optimal 1-MD estimator. On the one hand, $\hat{P}_0 = \hat{P}$ is a reasonable choice for an optimal estimator of the CCPs due to the asymptotic efficiency of the MLE. Given this choice, the asymptotic distribution of the 1-MD estimator is analogous to that of a standard GMM problem. From this observation, $W_1 = W_1^*$ follows. As one would expect, W_1^* coincides with the optimal weight matrix in the non-iterative analysis in Pesendorfer and Schmidt-Dengler (2008, Proposition 5). Finally, note that $W_1^* \neq \Omega_{PP}^{-1}$, i.e., the optimal weight matrix does not coincide the one that produces the 1-ML estimator. In fact, the 1-ML estimator need not be an asymptotical optimal 1-MD estimator, i.e., $\Sigma_{1-MD}(\hat{P}_0, \Omega_{PP}^{-1}) \neq \Sigma^*$.

We now move on to the general case with $K \geq 1$. According to Theorem 4.1, the asymptotic variance of the K-MD estimator depends on the number of iteration steps K, the asymptotic distribution of $\hat{P}_0$, and the entire sequence of limiting weight matrices $\{W_k : k \leq K\}$. In this sense, determining an optimal K-MD estimator appears to be a complicated task. The next result provides a concrete answer to this problem.

Theorem 4.3 (Invariance and optimality)**.** *Fix $K \geq 1$ arbitrarily and assume Assumptions A.1-A.6. In addition, assume that the sequence of weight matrices $\{W_k : k \leq K-1\}$ is such that the matrix*

$$\Lambda_K \equiv \mathbf{I}_{d_P} + 1[K>1] \sum_{b=1}^{K-1} \prod_{c=1}^{b} \Psi_P(\mathbf{I}_{d_P} - \Psi_\alpha(\Psi_\alpha' W_{K-c} \Psi_\alpha)^{-1} \Psi_\alpha' W_{K-c}) \in \mathbb{R}^{d_P \times d_P} \quad (4.4)$$

is non-singular. Then, we have the following two results.

1. *Invariance. Let $\hat{\alpha}_{K-MD}^*$ denote the K-MD estimator with $\hat{P}_0 = \hat{P}$, weight matrices $\{W_k : k \leq K-1\}$ for steps $1, \ldots, K-1$ (if $K > 1$), and the corresponding optimal weight matrix in step K. Then,*

$$\sqrt{n}(\hat{\alpha}_{K-MD}^* - \alpha^*) \overset{d}{\to} N(\mathbf{0}_{d_\alpha \times 1}, \Sigma^*),$$

   where Σ^* is as in Eq. (4.3).

2. *Optimality. Let $\hat{\alpha}_{K-MD}$ denote the K-MD estimator with $\hat{P}_0$ and weight matrices $\{W_k : k \leq K\}$. Then,*

$$\sqrt{n}(\hat{\alpha}_{K-MD} - \alpha^*) \overset{d}{\to} N(\mathbf{0}_{d_\alpha \times 1}, \Sigma_{K-MD}(\hat{P}_0, \{W_k : k \leq K\})).$$

   Furthermore, $\Sigma_{K-MD}(\hat{P}_0, \{W_k : k \leq K\}) - \Sigma^*$ is positive semidefinite, i.e., $\hat{\alpha}_{K-MD}^*$ is optimal among all K-MD estimators that satisfy our assumptions.

Besides Assumptions A.1-A.6, Theorem 4.3 also requires that the matrix Λ_K defined in Eq. (4.4) is non-singular. This additional requirement appears to be very mild, as it was satisfied in the vast majority of our Monte Carlo simulations. We now provide some theoretical interpretation of this condition. In single-agent problems, the zero Jacobian property (i.e. $\Psi_P = \mathbf{0}_{dP \times dP}$) implies that $\Lambda_K = \mathbf{I}_{d_P}$, which is always invertible. In multiple-agent problems, the zero Jacobian property no longer holds and, thus, a singular Λ_K is plausible. However, note that Λ_K will be non-singular as long as the amount of strategic interaction is sufficiently small, i.e., if Ψ_P is sufficiently close to zero. We can then interpret the invertibility of Λ_K as a high-level condition that restricts the amount of strategic interaction between the players in the game.

Theorem 4.3 is the main finding of this paper, and it establishes two central results regarding the asymptotic optimality of the K-MD estimator. We begin by discussing the first one, referred to as "invariance". This result focuses on a particular CCP estimator: $\hat{P}_0 = \hat{P}$. This is a natural choice to consider, as $\hat{P}$ is asymptotically efficient under our assumptions. Given this choice, the asymptotic variance of the K-MD estimator depends on the entire sequence of weight matrices $\{W_k : k \leq K\}$. While the dependence on the first $K-1$ weight matrices is fairly complicated, the dependence on the last weight matrix (i.e., W_K)

resembles that of the weight matrix in a standard GMM problem. If the matrix in Eq. (4.4) is invertible, we can use standard GMM results to define an optimal choice for W_K, *given the sequence of first* $K-1$ *weight matrices*. In principle, one might expect that the resulting asymptotic variance depends on the first $K-1$ weight matrices. The "invariance" result reveals that this is not the case. In other words, for $\hat{P}_0 = \hat{P}$ and an optimal choice of W_K, the asymptotic distribution of the K-MD estimator is invariant to the first $K-1$ weight matrices, or even K. Furthermore, the resulting asymptotic distribution coincides with that of the optimal 1-MD estimator obtained in Theorem 4.2.

The "invariance" result is the key to the second result in Theorem 4.3, referred to as "optimality". This second result characterizes the optimal choice of $\hat{P}_0$ and $\{W_k : k \leq K\}$ for K-MD estimators. The intuition of the result is as follows. First, given that $\hat{P}$ is the asymptotically efficient estimator of the CCPs, it is intuitive that $\hat{P}_0 = \hat{P}$ is optimal. Second, it is also intuitive that asymptotic efficiency requires setting W_K to be optimal, *given the sequence of first* $K-1$ *weight matrices*. At this point, our "invariance" result indicates that the asymptotic distribution does not depend on K or the first $K-1$ weight matrices. From this, we can then conclude that the K-MD estimator with $\hat{P}_0 = \hat{P}$ and an optimal last weight matrix W_K (given any first $K-1$ weight matrices) is asymptotically efficient among all K-MD estimators. Also, this optimal weight matrix can be estimated by sample analogues. By the usual asymptotic arguments based on Slutsky's theorem, the feasible K-MD estimator that uses the sample analogue of the optimal weight matrix is also asymptotically efficient.

Theorem 4.3 implies two important corollaries regarding the optimal 1-MD estimator discussed in Theorem 4.2. The first corollary is that the optimal 1-MD estimator is asymptotically efficient in the class of all K-MD estimators that satisfy the assumptions of Theorem 4.3. In other words, additional policy iterations do not provide asymptotic efficiency gains relative to the optimal 1-MD estimator. From a computational point of view, this result suggests that researchers should restrict attention to the non-iterative 1-MD estimator with $\hat{P}_0 = \hat{P}$ and $W_1 = W_1^*$, studied in Theorem 4.2. The intuition behind this result is that the multiple iteration steps of the K-MD estimator are merely reprocessing the sample information, i.e., no new information is added with each iteration step. Provided that the criterion function is optimally weighted, the non-iterative 1-MD estimator is capable of processing the sample information in an asymptotically efficient manner. Thus, additional iteration steps do not provide any additional asymptotic efficiency.

The second corollary of Theorem 4.3 is that the K-ML estimator is *usually* not asymptotically efficient, and can be feasibly improved upon. In particular, provided that the assumptions of Theorem 4.3 are satisfied when $W_k = \Omega_{PP}^{-1}$ for all $k \leq K$, the non-iterative 1-MD estimator with $\hat{P}_0 = \hat{P}$ and $W_1 = W_1^*$ is more or equally asymptotically efficient than the K-ML estimator.

We illustrate the results of this section by revising the example of Section 2.4 with the parameter values considered in Section 3. The asymptotic variance of $\hat{\alpha}_{K-MD} = (\hat{\lambda}_{RN,K-MD}, \hat{\lambda}_{EC,K-MD})$ is given by:

$$\Sigma_{K-MD}(\hat{P}_0, \{W_k : k \leq K\}) = (\Psi_\alpha' W_K \Psi_\alpha)^{-1} \Psi_\alpha' W_K (\mathbf{I}_{d_P} - \Psi_P \Phi_{k,P0}) \Omega_{PP} (\mathbf{I}_{d_P} - \Psi_P \Phi_{k,P})' W_K \Psi_\alpha (\Psi_\alpha' W_K \Psi_\alpha)^{-1}, \tag{4.5}$$

where $\{\Phi_{k,P0} : k \leq K\}$ is defined by $\Phi_{k,P0} \equiv \Phi_{k,P} + \Phi_{k,0}$, with $\{\Phi_{k,P} : k \leq K\}$ and $\{\Phi_{k,0} : k \leq K\}$ as in Eq. (4.1). For any true parameter vector and any $K \in \mathbb{N}$, we can numerically compute Eq. (4.5).

In Section 3, we considered three specific parameter values of $(\lambda_{RN}^*, \lambda_{EC}^*)$ that produced an asymptotic variance of the K-ML estimator of λ_{RN}^* decreased, increased, and wiggled with K. We now compute the optimal K-MD estimator of λ_{RN}^* for the same parameter values. The results are presented in Figures 4, 5, and 6, respectively.[7] These graphs illustrate the findings in Theorem 4.3. In accordance to the "invariance" result, the asymptotic variance of the optimal K-MD estimator does not vary with the number of iterations K. Also in accordance to the "optimality" result, the asymptotic variance of the optimal K-MD estimator is lower than any other K-MD estimator. In turn, since the asymptotic distribution of the K-ML estimator is a special case of the K-MD estimator, the asymptotic variance of the optimal K-MD estimator is lower than the K-ML estimator for all $K \in \mathbb{N}$. Combining both results, the optimal non-iterative 1-MD estimator is both computationally convenient and asymptotically efficient among the estimators under consideration.

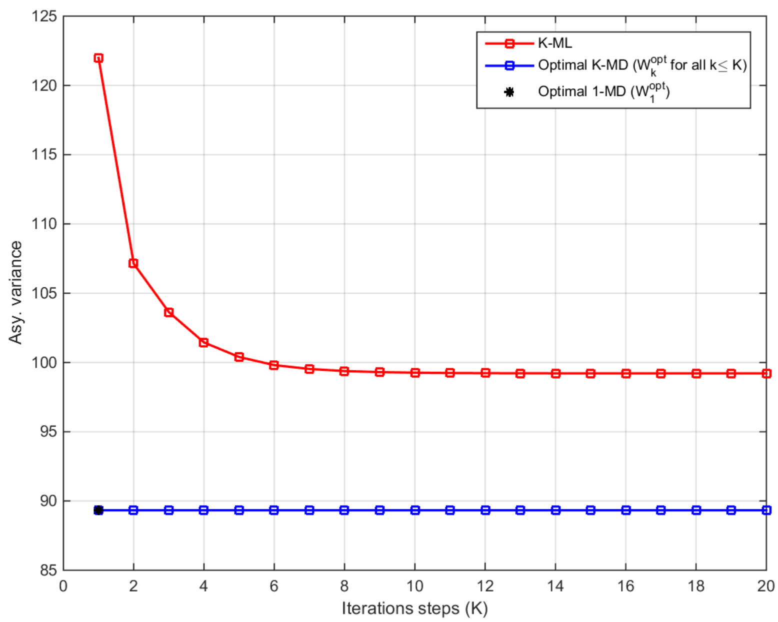

Figure 4: Asymptotic variance of the K-ML and optimal K-MD estimator of λ^*_{RN} as a function of the number of iterations K when $(\lambda^*_{RN}, \lambda^*_{EC}, \lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*) = (2.8, 0.8, 0.7, 0.6, 0.4, 0.95)$. The optimal K-MD estimator is computed using the optimal weighting matrix in every iteration step.

[7] According to the "invariance" result in Theorem 4.3, there are multiple asymptotically equivalent ways of implementing the optimal K-MD estimator. For concreteness, we set the weight matrix optimally in each iteration step.

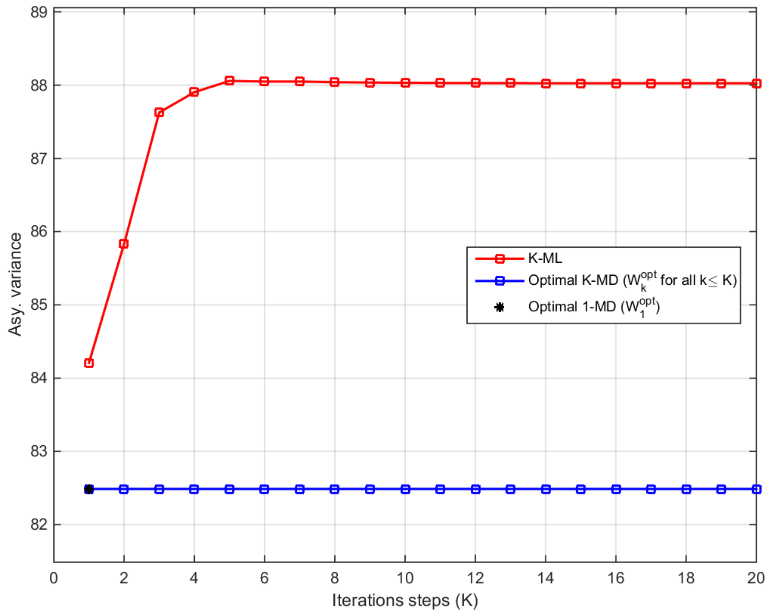

Figure 5: Asymptotic variance of the K-ML and optimal K-MD estimator of λ^*_{RN} as a function of the number of iterations K when $(\lambda^*_{RN}, \lambda^*_{EC}, \lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*) = (2, 1.8, 0.2, 0.1, 0.3, 0.95)$. The optimal K-MD estimator is computed using the optimal weighting matrix in every iteration step.

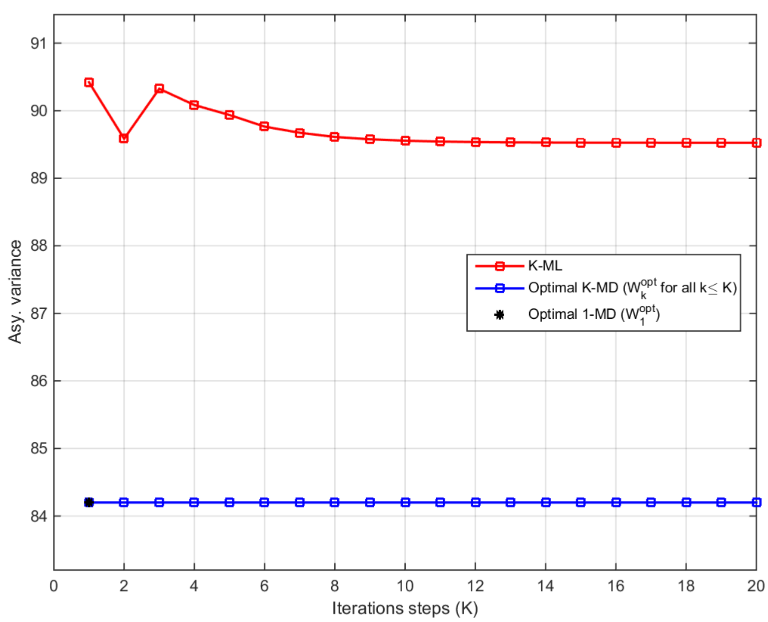

Figure 6: Asymptotic variance of the K-ML and optimal K-MD estimator of λ^*_{RN} as a function of the number of iterations K when $(\lambda^*_{RN}, \lambda^*_{EC}, \lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*) = (2.2, 1.45, 0.45, 0.22, 0.29, 0.95)$. The optimal K-MD estimator is computed using the optimal weighting matrix in every iteration step.

5 Monte Carlo simulations

This section investigates the finite sample performance of K-ML and K-MD estimators considered in previous sections. We simulate data using the two-player dynamic entry game described in Section 2.4. Recall that this model is specified up to the parameters $(\lambda^*_{RS}, \lambda^*_{RN}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \lambda^*_{EC}, \beta^*)$. For simplicity, we assume that the researcher knows $(\lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*)$ and wants to estimate $\alpha^* \equiv (\lambda^*_{RN}, \lambda^*_{EC})$. We consider the three specific parameter values that were used to illustrate the theoretical results in Sections 3 and 4.

Our simulation results are the average of $S = 10,000$ independent datasets $\{\{(\{a_{jt,i} : j \in J\}, x_{t,i}, x'_{t,i})\} : i \leq n\}$ that are i.i.d. distributed according to the econometric model. We show results in tables for sample sizes $n \in \{500,\ 1,000,\ 2,000\}$. For brevity, we only show simulation results for λ^*_{RN}, as this was the focus of the discussion in Sections 3 and 4.

Table 1 provides results for the first parameter value, i.e., $(\lambda^*_{RN}, \lambda^*_{EC}, \lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*) = (2.8, 0.8, 0.7, 0.6, 0.4, 0.95)$. Recall from previous sections that this parameter value produces an asymptotic variance of the K-ML estimator of λ^*_{RN} that decreases with K (see Figures 1 and 4, repeated at the bottom of the Table 1).

Let us first focus on the results for the K-ML estimator. The simulation results closely resemble the predictions from the asymptotic approximation. First, the empirical variance and mean squared error are extremely close, indicating that the asymptotic bias is almost negligible. Second, the empirical variance is decreasing with K and is close to the one predicted by our asymptotic analysis.

Next, we turn attention to the optimal K-MD estimator. Recall that the "invariance" result in Theorem 4.3 indicates that there are multiple asymptotically equivalent ways of implementing the optimal K-MD estimator. Throughout this section, the optimal K-MD estimator is a *feasible* estimator of the optimal K-MD estimator derived in Theorem 4.3 where, in each iteration step, we estimate the optimal weight matrix. According to our theoretical results, this feasible optimal K-MD estimator is asymptotically optimal among K-MD estimators, has zero asymptotic bias, and has an asymptotic variance that does not change with K. For the most part, these predictions are satisfied in our simulations. First, the empirical variance and mean squared error are again extremely close, and so the finite-sample bias is almost negligible. Second, the empirical variance is close to the one predicted by our asymptotic analysis. As the predicted by the "optimality" result in Theorem 4.3, the feasible optimal K-MD estimator is more efficient than the K-ML estimator. For most values of K under consideration, the empirical variance of the K-MD estimator appears to be invariant to K, especially for the larger sample sizes. However, we find that the empirical variance slightly decreases between $K = 1$ and $K = 2$. Our asymptotic analysis cannot explain this last empirical finding. This anomalous behavior for low values of K is analogous to the one found by Aguirregabiria and Mira (2002) and could be related to higher-order analysis in Kasahara and Shimotsu (2008). A high-order analysis of these estimators is out of the scope of this paper and is left for future work.

Estimator	Statistic	$K=1$	$K=2$	$K=3$	$K=4$	$K=5$	$K=10$	$K=15$	$K=20$
					$n=500$				
K-ML	Var	123.11	107.36	104.81	100.87	99.16	95.99	95.77	95.75
	MSE	123.13	107.68	105.34	101.09	99.32	96.05	95.83	95.81
Opt. K-MD	Var	94.77	88.34	87.46	86.49	86.02	85.10	85.03	85.03
	MSE	96.16	88.68	87.46	86.50	86.09	85.24	85.18	85.17
					$n=1,000$				
K-ML	Var	126.01	109.00	105.87	102.65	101.24	99.01	98.89	98.88
	MSE	126.07	109.34	106.34	102.91	101.44	99.12	99.00	98.99
Opt. K-MD	Var	93.39	89.61	88.91	88.43	88.14	87.70	87.67	87.67
	MSE	94.30	89.89	88.92	88.43	88.15	87.73	87.70	87.70
					$n=2,000$				
K-ML	Var	122.87	107.33	103.80	101.12	99.85	98.18	98.10	98.10
	MSE	122.87	107.42	103.92	101.17	99.89	98.20	98.12	98.12
Opt. K-MD	Var	90.11	88.18	87.80	87.59	87.44	87.25	87.22	87.23
	MSE	90.40	88.27	87.80	87.60	87.46	87.29	87.27	87.28
					Asymptotic results				
K-ML	Asy. Var	121.98	107.13	103.63	101.44	100.39	99.26	99.21	99.21
Opt. K-MD	Asy. Var	89.33	89.33	89.33	89.33	89.33	89.33	89.33	89.33

Table 1: Simulation results for estimation of λ^*_{RN} when $(\lambda^*_{RN}, \lambda^*_{EC}, \lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*) = (2.8, 0.8, 0.7, 0.6, 0.4, 0.95)$. "$K$-ML" denotes the K-ML estimator and "Opt. K-MD" denotes the feasible optimal K-MD estimator computed with an estimated optimal weight matrix in every iteration step. "Var" denotes the average empirical variance scaled by n and "MSE" denotes the average scaled mean squared error scaled by n and, for both statistics, the average is computed over $S = 10,000$ simulations. Finally, "Asy. Var" denotes the asymptotic variance according to Theorems 3.1, 4.1, and 4.3.

Table 2 provides results for the second parameter value, i.e., $(\lambda^*_{RN}, \lambda^*_{EC}, \lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*) = (2, 1.8, 0.2, 0.1, 0.3, 0.95)$. Recall that this parameter value produced an asymptotic variance of the K-ML estimator of λ^*_{RN} that increases with K (see Figures 2 and 5, repeated at the bottom of the Table 2). In turn, Table 3 provides results for the third parameter value, i.e., $(\lambda^*_{RN}, \lambda^*_{EC}, \lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*) = (2.2, 1.45, 0.45, 0.22, 0.29, 0.95)$. This parameter value produced an asymptotic variance of the K-ML estimator of λ^*_{RN} that wiggles with K (see Figures 3 and 6, repeated at the bottom of the Table 3).

The simulation results for these two parameter values are qualitatively similar to the ones obtained for the first parameter value and, for the most part, support our theoretical conclusions. First, both estimators have very little empirical bias. Second, all the estimators have an empirical variance that is very close to the one predicted by the asymptotic analysis. In particular, the empirical variance of the K-ML estimator is increasing in K for the second parameter value and wiggles for the third parameter value. Third, in most cases, the empirical variance of the optimal K-MD estimator is lower than that of the K-ML estimator. Finally, the empirical variance of the optimal K-MD estimator is invariant to K except for small values of K for which it is decreasing. One notable difference relative to the simulation is that the range of iterations over which the empirical variance decreases now extends between $K = 1$ and $K = 5$. Once again, we conjecture that this phenomenon is related to high-order effect, which is out of the scope of the analysis of this paper.

Estimator	Statistic	$K=1$	$K=2$	$K=3$	$K=4$	$K=5$	$K=10$	$K=15$	$K=20$
					$n=500$				
K-ML	Var	85.84	87.48	91.57	91.55	92.18	91.63	91.18	90.90
	MSE	85.84	87.82	92.39	92.21	92.83	92.18	91.70	91.42
Opt. K-MD	Var	111.25	96.26	89.68	88.36	87.41	84.98	84.81	84.46
	MSE	116.56	98.19	90.09	88.57	87.50	84.99	84.82	84.47
					$n=1,000$				
K-ML	Var	86.09	87.45	90.75	90.84	91.28	90.87	90.62	90.52
	MSE	86.12	87.73	91.32	91.32	91.76	91.27	91.01	90.91
Opt. K-MD	Var	99.33	89.30	85.78	85.01	84.20	83.13	82.85	82.77
	MSE	102.77	90.42	86.05	85.15	84.26	83.15	82.86	82.78
					$n=2,000$				
K-ML	Var	82.42	83.96	86.38	86.59	86.84	86.62	86.54	86.52
	MSE	82.42	84.05	86.56	86.73	86.98	86.74	86.65	86.63
Opt. K-MD	Var	87.91	83.43	81.53	81.26	80.86	80.39	80.28	80.27
	MSE	89.53	83.92	81.63	81.31	80.88	80.39	80.29	80.27
					Asymptotic results				
K-ML	Asy. Var	84.21	85.83	87.63	87.90	88.06	88.03	88.03	88.03
Opt. K-MD	Asy. Var	82.49	82.49	82.49	82.49	82.49	82.49	82.49	82.49

Table 2: Simulation results for estimation of λ^*_{RN} when $(\lambda^*_{RN}, \lambda^*_{EC}, \lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*) = (2, 1.8, 0.2, 0.1, 0.3, 0.95)$. "$K$-ML" denotes the K-ML estimator and "Opt. K-MD" denotes the feasible optimal K-MD estimator computed with an estimated optimal weight matrix in every iteration step. "Var" denotes the average empirical variance scaled by n and "MSE" denotes the average scaled mean squared error scaled by n and, for both statistics, the average is computed over $S = 10,000$ simulations. Finally, "Asy. Var" denotes the asymptotic variance according to Theorems 3.1, 4.1, and 4.3.

Estimator	Statistic	$K=1$	$K=2$	$K=3$	$K=4$	$K=5$	$K=10$	$K=15$	$K=20$
					$n=500$				
K-ML	Var	92.71	91.47	94.02	92.98	92.65	89.98	89.31	89.17
	MSE	92.72	91.88	94.83	93.58	93.20	90.36	89.66	89.52
Opt. K-MD	Var	101.56	88.31	84.28	83.60	82.52	81.10	80.76	80.70
	MSE	105.39	89.50	84.45	83.68	82.54	81.10	80.77	80.71
					$n=1,000$				
K-ML	Var	92.01	91.95	94.34	93.71	93.54	92.07	91.78	91.73
	MSE	92.02	92.19	94.83	94.08	93.88	92.32	92.02	91.97
Opt. K-MD	Var	95.01	86.93	84.72	84.47	83.92	83.30	83.17	83.16
	MSE	97.29	87.57	84.82	84.51	83.93	83.30	83.17	83.16
					$n=2,000$				
K-ML	Var	89.51	89.15	90.60	90.18	90.00	89.09	88.95	88.94
	MSE	89.52	89.26	90.80	90.33	90.14	89.18	89.04	89.03
Opt. K-MD	Var	86.67	83.69	82.33	82.26	81.94	81.61	81.56	81.55
	MSE	87.80	84.03	82.38	82.28	81.95	81.61	81.56	81.55
					Asymptotic results				
K-ML	Asy. Var	90.42	89.58	90.32	90.08	89.94	89.56	89.53	89.52
Opt. K-MD	Asy. Var	84.20	84.20	84.20	84.20	84.20	84.20	84.20	84.20

Table 3: Simulation results for estimation of λ^*_{RN} when $(\lambda^*_{RN}, \lambda^*_{EC}, \lambda^*_{RS}, \lambda^*_{FC,1}, \lambda^*_{FC,2}, \beta^*) = (2.2, 1.45, 0.45, 0.22, 0.29, 0.95)$. "$K$-ML" denotes the K-ML estimator and "Opt. K-MD" denotes the feasible optimal K-MD estimator computed with an estimated optimal weight matrix in every iteration step. "Var" denotes the average empirical variance scaled by n and "MSE" denotes the average scaled mean squared error scaled by n and, for both statistics, the average is computed over $S = 10,000$ simulations. Finally, "Asy. Var" denotes the asymptotic variance according to Theorems 3.1, 4.1, and 4.3.

6 Conclusions

This paper investigates the asymptotic properties of a class of estimators of the structural parameters in dynamic discrete choice games. We consider K-stage policy iteration (PI) estimators, where $K \in \mathbb{N}$ denotes the number of policy iterations employed in the estimation. This class nests several estimators proposed in the literature. By considering a "maximum likelihood" criterion function, the K-stage PI estimator becomes the K-ML estimator in Aguirregabiria and Mira (2002, 2007). By considering a "minimum distance" criterion function, K-stage PI estimator defines a novel K-MD estimator, which is an iterative version of the estimators in Pesendorfer and Schmidt-Dengler (2008) and Pakes et al. (2007). Since we consider an asymptotic framework with fixed K and $n \to \infty$, our analysis is not affected by the problems described in Pesendorfer and Schmidt-Dengler (2010).

First, we establish that the K-ML estimator is consistent and asymptotically normal for any $K \in \mathbb{N}$. This complements findings in Aguirregabiria and Mira (2007), who focus on $K = 1$ and K large enough to induce convergence of the estimator. Furthermore, we show that the asymptotic variance of the K-ML estimator can exhibit arbitrary patterns as a function K. In particular, we show that by changing the parameter values in a typical dynamic discrete choice game, the asymptotic variance of the K-ML estimator can increase, decrease, or even oscillate with K.

Second, we also establish that the K-MD estimator is consistent and asymptotically normal for any $K \in \mathbb{N}$. Its asymptotic distribution depends on the choice of the weight matrix. For a specific weight matrix, the K-MD has the same asymptotic distribution as the K-ML. We investigate the optimal choice of the weight matrix for the K-MD estimator. Our main result shows that an optimally weighted K-MD estimator has an asymptotic distribution that is invariant to K. This appears to be a novel result in the literature on PI estimation for games, and it is particularly surprising given the findings in Aguirregabiria and Mira (2007) for K-ML estimators.

The main result in our paper implies two important corollaries regarding the optimal 1-MD estimator (derived by Pesendorfer and Schmidt-Dengler (2008)). First, the optimal 1-MD estimator is optimal in the class of K-MD estimators for all $K \in \mathbb{N}$. In other words, additional policy iterations do not provide asymptotic efficiency gains relative to the optimal 1-MD estimator. Second, the optimal 1-MD estimator is more or equally asymptotically efficient than any K-ML estimator for all $K \in \mathbb{N}$.

We explored our theoretical findings in Monte Carlo simulations. For the most part, our finite-sample simulation evidence supports our asymptotic conclusions. The K-ML and the optimal K-MD estimators have negligible empirical bias and have an empirical variance that is very close to the one predicted by the asymptotic analysis. In most cases, the empirical variance of the optimal K-MD estimator is lower than that of the K-ML estimator. Also, it appears to be invariant to K except for very small values of K for which it is decreasing in K. The behavior for low values of K is analogous to the one found by Aguirregabiria and Mira (2002) and could be related to higher-order analysis in Kasahara and Shimotsu (2008).

A Appendix

Throughout this appendix, "s.t." abbreviates "such that", "RHS" abbreviates, and "PSD" abbreviates "positive semidefinite".

Several results in the appendix make use of the following high-level assumption. Note that whenever this assumption is used to prove results in the main text, we first verify that it holds under the lower-level conditions.

Assumption A.7. (High-level assumptions for iterative estimators). *There is a sequence of limiting criterion* *functions* $\{Q_k : k \leq K\}$ *with* $Q_k : \Theta_\alpha \times \Theta_g \times \Theta_P \to \mathbb{R}$ *such that:*

1. $\sup_{\alpha\in\Theta_\alpha} |\hat{Q}_k(\alpha, \tilde{g}, \tilde{P}) - Q_k(\alpha, g^*, P^*)| = o_p(1)$*, provided that* $(\tilde{g}, \tilde{P}) = (g^*, P^*) + o_p(1)$.
2. $Q_k(\alpha, g^*, P^*)$ *is uniquely maximized at* α^*.
3. $\sqrt{n}\partial\hat{Q}_k(\alpha^*, g^*, P^*)/\partial\alpha = \Xi_k\sqrt{n}(\hat{P} - P^*) + o_p(1)$*, for some matrix* Ξ_k.
4. *For any* $\lambda \in \{\alpha, g, P\}$, $\partial^2\hat{Q}_k(\tilde{\alpha}, \tilde{g}, \tilde{P})/\partial\alpha\partial\lambda' = \partial^2 Q_k(\alpha^*, g^*, P^*)/\partial\alpha\partial\lambda' + o_p(1)$*, provided that* $(\tilde{\alpha}, \tilde{g}, \tilde{P}) = (\alpha^*, g^*, P^*) + o_p(1)$.
5. $\partial^2 Q_k(\alpha^*, g^*, P^*)/\partial\alpha\partial\alpha'$ *is non-singular.*

A.1 Proofs of results in the main text

Proof of Theorem 3.1. This proof will require the following notation. For any $(\alpha, g, P) \in \Theta_\alpha \times \Theta_g \times \Theta_P$ and $(j, a, x) \in J \times A \times X$, let $n_{jax} \equiv \sum_{i=1}^n 1[(a_{jt,i}, x_{t,i}) = (a, x)]$, $n_x \equiv \sum_{i=1}^n 1[x_{t,i} = x]$, $\Psi_{jax}(\alpha, g, P) \equiv \prod_{j\in J}\Psi_j(\alpha, g, P)(a|x)$, $\hat{P}_{jax} \equiv n_{jax}/n_x$, $P^*_{jax} \equiv P^*_j(a|x) = \sum_{x'\in X}\Pi^*_j(a, x, x')/\sum_{(\tilde{a},\tilde{x}')\in A\times X}\Pi^*_j(\tilde{a}, x, \tilde{x}')$, and $m^*(x) \equiv \sum_{(a,x')\in A\times X}\Pi^*_j(a, x, x')$. Note that RHS of the last equation does not change with $j \in J$ by the equilibrium assumption in Assumption A.1. Also, Assumptions A.2 and A.3 imply $P^*_{jax} > 0$ and $m^*(x) > 0$ for every $(j, a, x) \in J \times A \times X$.

As we now show, Theorem 3.1 is a consequence of applying Theorem A.1 with $\hat{Q}_k \equiv \hat{Q}_{ML}$ and $Q_k \equiv Q_{ML}$ where, for any $(\alpha, g, P) \in \Theta_\alpha \times \Theta_g \times \Theta_P$,

$$\begin{aligned}\hat{Q}_{ML}(\alpha, g, P) &\equiv \frac{1}{n}\sum_{i=1}^n \ln\Psi(\alpha, g, P)(a_i|x_i) = \sum_{(j,a,x)\in J\times A\times X}\frac{n_{jax}}{n}\ln\Psi_{jax}(\alpha, g, P)\\ &= \sum_{(j,x)\in J\times X}\frac{n_x}{n}\left[\sum_{a\in\tilde{A}}\hat{P}_{jax}\ln\Psi_{jax}(\alpha, g, P) + \hat{P}_{j0x}\ln(1 - \sum_{a\in\tilde{A}}\Psi_{jax}(\alpha, g, P))\right],\end{aligned}$$

and

$$\begin{aligned}Q_{ML}(\alpha, g, P) &\equiv \sum_{(j,a,x)\in J\times A\times X} m^*(x)P^*_{jax}\ln\Psi_{jax}(\alpha, g, P)\\ &= \sum_{(j,x)\in J\times X} m^*(x)\left[\sum_{a\in\tilde{A}}P^*_{jax}\ln\Psi_{jax}(\alpha, g, P) + P^*_{j0x}\ln(1 - \sum_{a\in\tilde{A}}\Psi_{jax}(\alpha, g, P))\right],\end{aligned}$$

and where we have used that $\Psi_{j0x}(\alpha, g, P) = 1 - \sum_{a\in\tilde{A}}\Psi_{jax}(\alpha, g, P)$.

For any $\lambda \in \{\alpha, g, P\}$, note that:

$$\begin{aligned}\frac{\partial^2 Q_{ML}(\alpha^*, g^*, P^*)}{\partial\alpha\partial\lambda'} &= -\sum_{(j,x)\in J\times X} m^*(x)\left[\begin{array}{c}\frac{1}{\Psi_{jax}(\alpha^*,g^*,P^*)}\sum_{a\in\tilde{A}}\frac{\partial\Psi_{jax}(\alpha^*,g^*,P^*)}{\partial\alpha}\frac{\partial\Psi_{jax}(\alpha^*,g^*,P^*)}{\partial\lambda'}\\ +\frac{1}{\Psi_{j0x}(\alpha^*,g^*,P^*)}\sum_{\check{a}\in\tilde{A}}\frac{\partial\Psi_{\check{a}jx}(\alpha^*,g^*,P^*)}{\partial\alpha}\sum_{\tilde{a}\in\tilde{A}}\frac{\partial\Psi_{\tilde{a}jx}(\alpha^*,g^*,P^*)}{\partial\lambda'}\end{array}\right]\\ &= -\left\{\begin{array}{c}\{\frac{\partial\Psi_{jax}(\alpha^*,g^*,P^*)}{\partial\alpha} : (j,a,x)\in J\times\tilde{A}\times X\}'\times\\ diag\{m^*(x)(diag\{1/P^*_{jax} : a\in\tilde{A}\} + \mathbf{1}_{|\tilde{A}|\times|\tilde{A}|}/P^*_{j0x}) : (j,x)\in J\times X\}\times\\ \{\frac{\partial\Psi_{jax}(\alpha^*,g^*,P^*)}{\partial\lambda} : (j,a,x)\in J\times\tilde{A}\times X\}\end{array}\right\} = -\Psi'_\alpha\Omega^{-1}_{PP}\Psi_\lambda,\end{aligned}$$

where the first equality uses that Assumptions A.2 and A.3, the second equality follows from Assumption A.2, and the final equality follows from the following argument. By Eq. (2.6), $\Omega_{PP} = diag\{\Sigma_{jx} : (j, x) \in J \times X\}$ with $\Sigma_{jx} \equiv (diag\{P^*_{jx}\} - P^*_{jx}P^{*\prime}_{jx})/m^*(x)$ and $P^*_{jx} \equiv \{P^*_{jax} : a \in \tilde{A}\}$, and so $\Omega^{-1}_{PP} = diag\{\Sigma^{-1}_{jx} : (j, x) \in J \times X\}$ with $\Sigma^{-1}_{jx} = m^*(x)(diag\{1/P^*_{jax} : a \in \tilde{A}\} + \mathbf{1}_{|\tilde{A}|\times|\tilde{A}|}/P^*_{j0x})$.

To apply Theorem A.1, we first verify Assumption A.7.

Part (a). For any $(\tilde{g},\tilde{P}) = (g^*,P^*) + o_p(1)$,

$$\sup_{\alpha\in\Theta_\alpha} |\hat{Q}_{ML}(\alpha,\tilde{g},\tilde{P}) - Q_{ML}(\alpha,g^*,P^*)| \leq \left[\begin{array}{c} \sup_{\alpha\in\Theta_\alpha} |\sum_{(j,a,x)\in J\times A\times X} \frac{n_{jax}}{n} \ln(\Psi_{jax}(\alpha,\tilde{g},\tilde{P})/\Psi_{jax}(\alpha,g^*,P^*))| \\ +\sup_{\alpha\in\Theta_\alpha} |\sum_{(j,a,x)\in J\times A\times X} (\frac{n_{jax}}{n} - m^*(x)P^*_{jax}) \ln \Psi_{jax}(\alpha,g^*,P^*)| \end{array}\right]$$
$$\leq \sum_{(j,a,x)\in J\times A\times X} \left[\begin{array}{c} \frac{n_{jax}}{n} \sup_{\alpha\in\Theta_\alpha} |\ln \Psi_{jax}(\alpha,\tilde{g},\tilde{P}) - \ln \Psi_{jax}(\alpha,g^*,P^*)| \\ +(\frac{n_{jax}}{n} - m^*(x)P^*_{jax})| \ln(\inf_{\alpha\in\Theta_\alpha} |\Psi_{jax}(\alpha,g^*,P^*)|) \end{array}\right] = o_p(1),$$

where the second inequality uses Assumptions A.1, A.3, and the intermediate value theorem.

Part (b). Define the function:

$$G(\alpha) \equiv Q_{ML}(\alpha,g^*,P^*) - Q_{ML}(\alpha^*,g^*,P^*) = \sum_{(j,a,x)\in J\times A\times X} m^*(x)P^*_{jax} \ln\left(\frac{\Psi_{jax}(\alpha,g^*,P^*)}{\Psi_{jax}(\alpha^*,g^*,P^*)}\right),$$

which is properly defined by Assumptions A.2 and A.3. By definition, $G(\alpha^*) = 0$. On the other hand, consider any $\alpha \neq \alpha^*$. Assumption A.2 implies that $\Psi(\alpha,g^*,P^*) \neq \Psi(\alpha^*,g^*,P^*)$. This and Assumption A.3 then implies that $\Psi_{jax}(\alpha,g^*,P^*)/\Psi_{jax}(\alpha^*,g^*,P^*) \neq 1$ for some $(j,a,x) \in J\times A\times X$. Then,

$$G(\alpha) < \ln\left(\sum_{(j,a,x)\in J\times A\times X} m^*(x)P^*_{jax}\frac{\Psi_{jax}(\alpha,g^*,P^*)}{\Psi_{jax}(\alpha^*,g^*,P^*)}\right) = \ln\left(\sum_{(j,a,x)\in J\times A\times X} m^*(x)\Psi_{jax}(\alpha,g^*,P^*)\right) = 0,$$

where inequality follows from Jensen's inequality, the strict convexity of the logarithm, and $\Psi_{jax}(\alpha,g^*,P^*)/\Psi_{jax}(\alpha^*,g^*,P^*) \neq 1$ for some $(j,a,x) \in J\times A\times X$, the first equality follows from Assumption A.3, and the final equality follows from $\sum_{(j,a,x)\in J\times A\times X} m^*(x)\Psi_{jax}(\alpha,g^*,P^*) = 1$ for any $\alpha \in \Theta_\alpha$. Therefore, $G(\alpha)$ and $Q_{ML}(\alpha,g^*,P^*)$ are uniquely maximized at $\alpha = \alpha^*$.

Part (c). Consider the following derivation for any $(\alpha,g,P) \in \Theta_\alpha \times \Theta_g \times \Theta_P$ s.t. $\Psi(\alpha,g,P)$ is positive and differentiable.

$$\begin{aligned}
\frac{\partial \hat{Q}_{ML}(\alpha,g,P)}{\partial\alpha} &= \frac{\partial}{\partial\alpha}\left\{\sum_{(j,x)\in J\times X} \frac{n_x}{n}\left[\sum_{a\in\tilde{A}} \hat{P}_{jax} \ln \Psi_{jax}(\alpha,g,P) + \hat{P}_{j0x} \ln\left(1 - \sum_{a\in\tilde{A}} \Psi_{jax}(\alpha,g,P)\right)\right]\right\} \\
&= \sum_{(j,a,x)\in J\times\tilde{A}\times X} \frac{n_x}{n}\left[\frac{\hat{P}_{jax}}{\Psi_{jax}(\alpha,g,P)} - \frac{\hat{P}_{j0x}}{\Psi_{j0x}(\alpha,g,P)}\right]\frac{\partial\Psi_{jax}(\alpha,g,P)}{\partial\alpha} \\
&= \sum_{(j,x)\in J\times X} \frac{n_x}{n}\left[\sum_{a\in\tilde{A}} \frac{\hat{P}_{jax} - \Psi_{jax}(\alpha,g,P)}{\Psi_{jax}(\alpha,g,P)} + \frac{\sum_{\tilde{a}\in\tilde{A}}(\hat{P}_{\tilde{a}jx} - \Psi_{\tilde{a}jx}(\alpha,g,P))}{\Psi_{j0x}(\alpha,g,P)}\right]\frac{\partial\Psi_{jax}(\alpha,g,P)}{\partial\alpha} \\
&= \sum_{(j,a,x)\in J\times\tilde{A}\times X} \frac{n_x}{n}\left[\frac{\hat{P}_{jax} - \Psi_{jax}(\alpha,g,P)}{\Psi_{jax}(\alpha,g,P)}\frac{\partial\Psi_{jax}(\alpha,g,P)}{\partial\alpha} + \frac{\hat{P}_{jax} - \Psi_{jax}(\alpha,g,P)}{\Psi_{j0x}(\alpha,g,P)}\sum_{\tilde{a}\in A}\frac{\partial\Psi_{\tilde{a}jx}(\alpha,g,P)}{\partial\alpha}\right], \qquad \text{(A.1)}
\end{aligned}$$

where we have used that $\hat{P}_{j0x} = 1 - \sum_{a\in\tilde{A}} \hat{P}_{jax}$ and $\Psi_{j0x}(\alpha,g,P) = 1 - \sum_{a\in\tilde{A}} \Psi_{jax}(\alpha,g,P)$, and so

$\partial\Psi_{j0x}(\alpha,g,P)/\partial\alpha = -\sum_{a\in\tilde{A}}\partial\Psi_{jax}(\alpha,g,P)/\partial\alpha$. Then,

$$\begin{aligned}
&\sqrt{n}\frac{\partial\hat{Q}_{ML}(\alpha^*,g^*,P^*)}{\partial\alpha}\\
&= \sum_{(j,a,x)\in J\times\tilde{A}\times X}\frac{n_x}{n}\sqrt{n}(\hat{P}_{jax}-P^*_{jax})\left[\frac{1}{P^*_{jax}}\frac{\partial\Psi_{jax}(\alpha^*,g^*,P^*)}{\partial\alpha}+\frac{1}{P^*_{j0x}}\sum_{\tilde{a}\in A}\frac{\partial\Psi_{\tilde{a}jx}(\alpha^*,g^*,P^*)}{\partial\alpha}\right]\\
&= \sum_{(j,a,x)\in J\times\tilde{A}\times X} m^*(x)\sqrt{n}(\hat{P}_{jax}-P^*_{jax})\left[\frac{1}{P^*_{jax}}\frac{\partial\Psi_{jax}(\alpha^*,g^*,P^*)}{\partial\alpha}+\frac{1}{P^*_{j0x}}\sum_{\tilde{a}\in A}\frac{\partial\Psi_{\tilde{a}jx}(\alpha^*,g^*,P^*)}{\partial\alpha}\right]+o_p(1)\\
&= \left[\begin{array}{c}\{\frac{\partial\Psi_{jax}(\alpha^*,g^*,P^*)}{\partial\alpha}:(j,a,x)\in J\times\tilde{A}\times X\}'\times\\ diag\{m^*(x)(diag\{1/P^*_{jax}:a\in\tilde{A}\}+\mathbf{1}_{|\tilde{A}|\times|\tilde{A}|}/P^*_{j0x}):(j,x)\in J\times X\}\times\\ \{\sqrt{n}(\hat{P}_{jax}-P^*_{jax}):(a,x,j)\in\tilde{A}\times X\times J\}\end{array}\right]+o_p(1)\\
&= \Psi'_\alpha\Omega^{-1}_{PP}\sqrt{n}(\hat{P}-P^*)+o_p(1),
\end{aligned}$$

where the first equality holds by Eq. (A.1) and Assumption A.3, the second equality holds by Assumption A.1, and the final equality follows $\Omega^{-1}_{PP} = diag\{\Sigma^{-1}_{jx}:(j,x)\in J\times X\}$ with $\Sigma^{-1}_{jx} = m^*(x)(diag\{1/P^*_{jax}:a\in\tilde{A}\}+\mathbf{1}_{|\tilde{A}|\times|\tilde{A}|}/P^*_{j0x})$.

Part (d). Consider the following derivation for any $(\alpha,g,P)\in\Theta_\alpha\times\Theta_g\times\Theta_P$ s.t. $\Psi(\alpha,g,P)$ is positive and twice differentiable.

$$\begin{aligned}
\frac{\partial^2\hat{Q}_{ML}(\alpha,g,P)}{\partial\alpha\partial\lambda'} &= \frac{\partial}{\partial\lambda'}\sum_{(j,a,x)\in J\times\tilde{A}\times X}\frac{n_x}{n}\left[\frac{\hat{P}_{jax}-\Psi_{jax}(\alpha,g,P)}{\Psi_{jax}(\alpha,g,P)}-\frac{\hat{P}_{j0x}-\Psi_{j0x}(\alpha,g,P)}{\Psi_{j0x}(\alpha,g,P)}\right]\frac{\partial\Psi_{jax}(\alpha,g,P)}{\partial\alpha}\\
&= \sum_{(j,a,x)\in J\times\tilde{A}\times X}\frac{n_x}{n}\left\{\begin{array}{c}\left[\frac{\hat{P}_{jax}-\Psi_{jax}(\alpha,g,P)}{\Psi_{jax}(\alpha,g,P)}-\frac{\hat{P}_{j0x}-\Psi_{j0x}(\alpha,g,P)}{\Psi_{j0x}(\alpha,g,P)}\right]\frac{\partial\Psi_{jax}(\alpha,g,P)}{\partial\alpha\partial\lambda'}\\ -\left[\frac{\Psi_{jax}(\alpha,g,P)+(\hat{P}_{jax}-\Psi_{jax}(\alpha,g,P))}{\Psi_{jax}(\alpha,g,P)^2}\right]\frac{\partial\Psi_{jax}(\alpha,g,P)}{\partial\alpha}\frac{\partial\Psi_{jax}(\alpha,g,P)}{\partial\lambda'}\\ +\left[\frac{\Psi_{j0x}(\alpha,g,P)+(\hat{P}_{j0x}-\Psi_{j0x}(\alpha,g,P))}{\Psi_{j0x}(\alpha,g,P)^2}\right]\frac{\partial\Psi_{jax}(\alpha,g,P)}{\partial\alpha}\frac{\partial\Psi_{j0x}(\alpha,g,P)}{\partial\lambda'}\end{array}\right\}. \quad (A.2)
\end{aligned}$$

Then, for any $\lambda\in\{\alpha,g,P\}$ and $(\tilde{\alpha},\tilde{g},\tilde{P}) = (\alpha^*,g^*,P^*)+o_p(1)$,

$$\begin{aligned}
\frac{\partial^2\hat{Q}_{ML}(\tilde{\alpha},\tilde{g},\tilde{P})}{\partial\alpha\partial\lambda'} &\xrightarrow{p} \sum_{(j,a,x)\in J\times\tilde{A}\times X} m^*(x)\left[\frac{1}{P^*_{0xj}}\frac{\partial\Psi_{jax}(\alpha^*,g^*,P^*)}{\partial\alpha}\frac{\partial\Psi_{j0x}(\alpha^*,g^*,P^*)}{\partial\lambda'}-\frac{1}{P^*_{axj}}\frac{\partial\Psi_{jax}(\alpha^*,g^*,P^*)}{\partial\alpha}\frac{\partial\Psi_{jax}(\alpha^*,g^*,P^*)}{\partial\lambda'}\right]\\
&= -\sum_{(j,x)\in J\times X} m^*(x)\left[\begin{array}{c}\sum_{a\in\tilde{A}}\frac{1}{P^*_{axj}}\frac{\partial\Psi_{jax}(\alpha^*,g^*,P^*)}{\partial\alpha}\frac{\partial\Psi_{jax}(\alpha^*,g^*,P^*)}{\partial\lambda'}+\\ \frac{1}{P^*_{0xj}}\sum_{\check{a}\in\tilde{A}}\frac{\partial\Psi_{\check{a}jx}(\alpha^*,g^*,P^*)}{\partial\alpha}\sum_{\tilde{a}\in\tilde{A}}\frac{\partial\Psi_{\tilde{a}jx}(\alpha^*,g^*,P^*)}{\partial\lambda'}\end{array}\right]\\
&= -\left[\begin{array}{c}\{\frac{\partial\Psi_{axj}(\alpha^*,g^*,P^*)}{\partial\alpha}:(a,x,j)\in A\times X\times\{1,\ldots,J\}\}'\times\\ diag\{m^*(x)[diag\{1/P^*_{axj}:a\in A\}+\mathbf{1}_{|A|\times|A|}\times 1/P^*_{0xj}]:(x,j)\in X\times\{1,\ldots,J\}\}\times\\ \{\frac{\partial\Psi_{axj}(\alpha^*,g^*,P^*)}{\partial\lambda}:(a,x,j)\in A\times X\times\{1,\ldots,J\}\}\end{array}\right]\\
&= -\Psi'_\alpha\Omega^{-1}_{PP}\Psi_\lambda = \frac{\partial^2 Q_{ML}(\alpha^*,g^*,P^*)}{\partial\alpha\partial\lambda'},
\end{aligned}$$

where the first equality holds by Eq. (A.2) and Assumptions A.1 and A.3, the second equality holds by $\Psi_{j0x}(\alpha,g,P) = 1-\sum_{a\in\tilde{A}}\Psi_{jax}(\alpha,g,P)$ and so $\partial\Psi_{j0x}(\alpha,g,P)/\partial\lambda' = -\sum_{a\in\tilde{A}}\partial\Psi_{jax}(\alpha,g,P)/\partial\lambda'$, and the final equality follows $\Omega^{-1}_{PP} = diag\{\Sigma^{-1}_{jx}:(j,x)\in J\times X\}$ with $\Sigma^{-1}_{jx} = m^*(x)(diag\{1/P^*_{jax}:a\in\tilde{A}\}+\mathbf{1}_{|\tilde{A}|\times|\tilde{A}|}/P^*_{j0x})$.

Part (e). $\partial^2 Q_{ML}(\alpha^*,g^*,P^*)/\partial\alpha\partial\alpha' = -\Psi'_\alpha\Omega^{-1}_{PP}\Psi_\alpha$ is nonsingular because Ψ_α has full rank.

This completes the verification of Assumption A.7. Since we also assume Assumptions A.3 and A.4, Theorem

A.1 applies. In particular, Eq. (A.21) yields:

$$\sqrt{n}(\hat{\theta}_{K-ML}-\theta^*) = \begin{pmatrix} \sqrt{n}(\hat{\alpha}_{K-ML}-\alpha^*) \\ \sqrt{n}(\hat{g}-g^*) \end{pmatrix}$$

$$= \begin{bmatrix} A_K + B_K\Upsilon_{K,P} & B_K\Upsilon_{K,0} & B_K\Upsilon_{K,g}+C_K \\ \mathbf{0}_{d_g\times d_P} & \mathbf{0}_{d_g\times d_P} & \mathbf{I}_{d_g} \end{bmatrix} \sqrt{n} \begin{bmatrix} \hat{P}-P^* \\ \hat{P}_0 - P^* \\ \hat{g}-g^* \end{bmatrix} + o_p(1), \quad \text{(A.3)}$$

with A_K, B_K, and C_K determined according to Eq. (A.22), and $\Upsilon_{K,P}$, $\Upsilon_{K,0}$, and $\Upsilon_{K,g}$ determined according to Eq. (A.23). As a next step, we now work out these constants.

For $k \leq K$, $\Xi_k = \Psi_\alpha'\Omega_{PP}^{-1}$ and $\partial^2 Q_{ML}(\alpha^*, g^*, P^*)/\partial\alpha\partial\lambda' = -\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\lambda$, and so, according to Eq. (A.22), $A_k = (\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1}\Psi_\alpha'\Omega_{PP}^{-1}$, $B_k = -(\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1}\Psi_\alpha'\Omega_{PP}^{-1}\Psi_P$, and $C_k = -(\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1}\Psi_\alpha'\Omega_{PP}^{-1}\Psi_g$. In addition, according to Eq. (A.23), $\{\Upsilon_{k,P} : k \leq K\}$, $\{\Upsilon_{k,g} : k \leq K\}$, and $\{\Upsilon_{k,0} : k \leq K\}$ are as follows. Set $\Upsilon_{1,0} \equiv \mathbf{I}_{d_P}$, $\Upsilon_{1,g} \equiv \mathbf{0}_{d_P\times d_g}$, $\Upsilon_{1,P} \equiv \mathbf{0}_{d_P\times d_P}$ and, for any $k = 1,\ldots,K-1$,

$$\begin{aligned}
\Upsilon_{k+1,P} &= (\mathbf{I}_{d_P} - \Psi_\alpha(\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1}\Psi_\alpha'\Omega_{PP}^{-1})\Psi_P\Upsilon_{k,P} + \Psi_\alpha(\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1}\Psi_\alpha'\Omega_{PP}^{-1} \\
\Upsilon_{k+1,0} &= (\mathbf{I}_{d_P} - \Psi_\alpha(\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1}\Psi_\alpha'\Omega_{PP}^{-1})\Psi_P\Upsilon_{k,0} \\
\Upsilon_{k+1,g} &= (\mathbf{I}_{d_P} - \Psi_\alpha(\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1}\Psi_\alpha'\Omega_{PP}^{-1})\Psi_P\Upsilon_{k,g} + (\mathbf{I}_{d_P} - \Psi_\alpha(\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1}\Psi_\alpha'\Omega_{PP}^{-1})\Psi_g.
\end{aligned}$$

Then, Eq. (3.1) follows from setting $\Upsilon_{k,P} \equiv \Phi_{k,P}$, and $\Upsilon_{k,g} \equiv \Phi_{k,g}\Psi_g$, $\Upsilon_{k,0} \equiv \Phi_{k,0}$ for all $k \leq K$.

If we plug this information into Eq. (A.3) and combine with Assumption A.4, we deduce that:

$$\sqrt{n}(\hat{\theta}_{K-ML}-\theta^*) \;=\; \begin{pmatrix} \sqrt{n}(\hat{\alpha}_{K-ML}-\alpha^*) \\ \sqrt{n}(\hat{g}-g^*) \end{pmatrix} \xrightarrow{d} N\left(\begin{pmatrix} \mathbf{0}_{d_\alpha} \\ \mathbf{0}_{d_g} \end{pmatrix}, \begin{pmatrix} \Sigma_{K-ML}(\hat{P}_0) & \Sigma_{\alpha g,K-ML} \\ \Sigma_{\alpha g,K-ML}' & \Omega_{gg} \end{pmatrix}\right), \quad \text{(A.4)}$$

where $\Sigma_{K-ML}(\hat{P}_0)$ is as defined in Theorem 3.1 and

$$\Sigma_{\alpha g,K-ML} \;\equiv\; (\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1}\Psi_\alpha'\Omega_{PP}^{-1}[(\mathbf{I}_{d_P} - \Psi_P\Phi_{K,P})\Omega_{Pg} - \Psi_P\Phi_{K,0}\Omega_{0g} - (\Psi_P\Phi_{K,g} + \mathbf{I}_{d_P})\Psi_g\Omega_{gg}].$$

The desired result is a corollary of Eq. (A.4). □

Proof of Theorem 4.1. As we now show, this result is a consequence of applying Theorem A.1 with $\hat{Q}_k \equiv \hat{Q}_{k-MD}$ and $Q_k \equiv Q_{k-MD}$ where, for any $(\alpha, g, P) \in \Theta_\alpha \times \Theta_g \times \Theta_P$,

$$Q_{k-MD}(\alpha, g, P) \;\equiv\; -(P^* - \Psi(\alpha, g, P))'W_k(P^* - \Psi(\alpha, g, P)).$$

For any $\lambda \in \{\alpha, g, P\}$, notice that $\partial^2 Q_{k-MD}(\alpha^*, g^*, P^*)/\partial\alpha\partial\lambda' = -2\Psi_\alpha' W_k\Psi_\lambda$.

To apply this result, we first verify Assumption A.7.

Part (a). For any $(\tilde{g}, \tilde{P}) = (g^*, P^*) + o_p(1)$,

$$\sup_{\alpha\in\Theta_\alpha} |\hat{Q}_{k-MD}(\alpha, \tilde{g}, \tilde{P}) - Q_{k-MD}(\alpha, g^*, P^*)| \leq \begin{bmatrix} \|\hat{W}_k - W_k\| + \|\tilde{P} - P^*\|^2\|W_k\| + 2\|\tilde{P} - P^*\|\|W_k\| \\ +2\|W_k\|\sup_{\alpha\in\Theta_\alpha}\|\Psi(\alpha, g^*, P^*) - \Psi(\alpha, \tilde{g}, \tilde{P})\| \end{bmatrix} = o_p(1),$$

where the last equality uses Assumption A.6.

Part (b). First, consider $\alpha = \alpha^*$. Then, Assumption A.2 implies $\Psi(\alpha, g^*, P^*) = P^*$ and this, in turn, implies that $Q_{k-MD}(\alpha^*, g^*, P^*) = 0$. Second, consider $\alpha \neq \alpha^*$. Then, Assumption A.2 implies $\Psi(\alpha, g^*, P^*) \neq P^*$. This and Assumption A.6 imply that $Q_{k-MD}(\alpha, g^*, P^*) < 0$. Then, $Q_{k-MD}(\alpha, g^*, P^*)$ is uniquely maximized at $\alpha = \alpha^*$, as required.

Part (c). Consider the following derivation:

$$\begin{aligned}\sqrt{n}\frac{\partial \hat{Q}_{k-MD}(\alpha^*,g^*,P^*)}{\partial\alpha} &= 2(\hat{P}-\Psi(\alpha^*,g^*,P^*))'\hat{W}_k\frac{\partial\Psi(\alpha^*,g^*,P^*)}{\partial\alpha}\\ &= 2\Psi_\alpha' W_k\sqrt{n}(\hat{P}-P^*)+o_p(1),\end{aligned}$$

where the second line uses Assumptions A.3 and A.6.

Part (d). For any $\lambda\in\{\alpha,g,P\}$ and $(\tilde{\alpha},\tilde{g},\tilde{P})=(\alpha^*,g^*,P^*)+o_p(1)$,

$$\begin{aligned}\frac{\partial^2\hat{Q}_{k-MD}(\tilde{\alpha},\tilde{g},\tilde{P})}{\partial\alpha\partial\lambda'} &= -2\frac{\partial\Psi(\tilde{\alpha},\tilde{g},\tilde{P})'}{\partial\lambda'}\hat{W}_k\frac{\partial\Psi(\tilde{\alpha},\tilde{g},\tilde{P})}{\partial\alpha}+2(\hat{P}-\Psi(\tilde{\alpha},\tilde{g},\tilde{P}))'\hat{W}_k\frac{\partial\Psi(\tilde{\alpha},\tilde{g},\tilde{P})}{\partial\alpha\partial\lambda'}\\ &\overset{p}{\to} \frac{\partial^2 Q_{k-MD}(\alpha^*,g^*,P^*)}{\partial\alpha\partial\lambda'}=-2\Psi_\alpha'W_k\Psi_\lambda,\end{aligned}$$

where the convergence uses Assumptions A.3 and A.6.

Part (e). $\partial^2 Q_{K-MD}(\alpha^*,g^*,P^*)/\partial\alpha\partial\alpha'=-2\Psi_\alpha'W_k\Psi_\alpha$ is nonsingular by Assumptions A.3 and A.6.

This completes the verification of Assumption A.7. Since we also assume Assumptions A.3 and A.4, Theorem A.1 applies. In particular, Eq. (A.21) yields:

$$\begin{aligned}\sqrt{n}(\hat{\theta}_{K-MD}-\theta^*) &= \left(\begin{array}{c}\sqrt{n}(\hat{\alpha}_{K-MD}-\alpha^*)\\ \sqrt{n}(\hat{g}-g^*)\end{array}\right)\\ &= \left[\begin{array}{ccc}A_K+B_K\Upsilon_{K,P} & B_K\Upsilon_{K,0} & B_K\Upsilon_{K,g}+C_K\\ \mathbf{0}_{d_g\times d_P} & \mathbf{0}_{d_g\times d_P} & \mathbf{I}_{d_g}\end{array}\right]\sqrt{n}\left[\begin{array}{c}\hat{P}-P^*\\ \hat{P}_0-P^*\\ \hat{g}-g^*\end{array}\right]+o_p(1), \qquad \text{(A.5)}\end{aligned}$$

with A_K, B_K, and C_K determined according to Eq. (A.22), and $\Upsilon_{K,P}$, $\Upsilon_{K,0}$, and $\Upsilon_{K,g}$ determined according to Eq. (A.23). As a next step, we now work out these constants.

For $k\leq K$, $\Xi_k=2\Psi_\alpha'W_k$ and $\partial^2Q_{k-MD}(\alpha^*,g^*,P^*)/\partial\alpha\partial\lambda'=-2\Psi_\alpha'W_k\Psi_\lambda$, and so $A_k=(\Psi_\alpha'W_k\Psi_\alpha)^{-1}\Psi_\alpha'W_k$, $B_k=-(\Psi_\alpha'W_k\Psi_\alpha)^{-1}\Psi_\alpha'W_k\Psi_P$, and $C_k=-(\Psi_\alpha'W_k\Psi_\alpha)^{-1}\Psi_\alpha'W_k\Psi_g$. Then, $\{\Upsilon_{k,P}:k\leq K\}$, $\{\Upsilon_{k,g}:k\leq K\}$, and $\{\Upsilon_{k,0}:k\leq K\}$ are as follows. Set $\Upsilon_{1,0}\equiv\mathbf{I}_{d_P}$, $\Upsilon_{1,g}\equiv\mathbf{0}_{d_P\times d_g}$, $\Upsilon_{1,P}\equiv\mathbf{0}_{d_P\times d_P}$ and, for any $k=1,\ldots,K-1$,

$$\begin{aligned}\Upsilon_{k+1,P}&=(\mathbf{I}_{d_P}-\Psi_\alpha(\Psi_\alpha'W_k\Psi_\alpha)^{-1}\Psi_\alpha'W_k)\Psi_P\Upsilon_{k,P}+\Psi_\alpha(\Psi_\alpha'W_k\Psi_\alpha)^{-1}\Psi_\alpha'W_k\\ \Upsilon_{k+1,0}&=(\mathbf{I}_{d_P}-\Psi_\alpha(\Psi_\alpha'W_k\Psi_\alpha)^{-1}\Psi_\alpha'W_k)\Psi_P\Upsilon_{k,0}\\ \Upsilon_{k+1,g}&=(\mathbf{I}_{d_P}-\Psi_\alpha(\Psi_\alpha'W_k\Psi_\alpha)^{-1}\Psi_\alpha'W_k)\Psi_P\Upsilon_{k,g}+(\mathbf{I}_{d_P}-\Psi_\alpha(\Psi_\alpha'W_k\Psi_\alpha)^{-1}\Psi_\alpha'W_k)\Psi_g.\end{aligned}$$

Then, Eq. (4.1) follows from setting $\Upsilon_{k,P}\equiv\Phi_{k,P}$, and $\Upsilon_{k,g}\equiv\Phi_{k,g}\Psi_g$, $\Upsilon_{k,0}\equiv\Phi_{k,0}$ for all $k\leq K$.

If we plug this information into Eq. (A.5) and combine with Assumption A.4, we deduce that:

$$\sqrt{n}(\hat{\theta}_{K-MD}-\theta^*) = \left(\begin{array}{c}\sqrt{n}(\hat{\alpha}_{K-MD}-\alpha^*)\\ \sqrt{n}(\hat{g}-g^*)\end{array}\right)\overset{d}{\to}N\left(\left(\begin{array}{c}\mathbf{0}_{d_\alpha}\\ \mathbf{0}_{d_g}\end{array}\right),\left(\begin{array}{cc}\Sigma_{K-MD}(\hat{P}_0,\{W_k:k\leq K\}) & \Sigma_{\alpha g,K-MD}\\ \Sigma_{\alpha g,K-MD}' & \Omega_{gg}\end{array}\right)\right), \qquad \text{(A.6)}$$

where $\Sigma_{K-MD}(\hat{P}_0,\{W_k:k\leq K\})$ is as defined in Theorem 4.1 and

$$\Sigma_{\alpha g,K-MD} \equiv (\Psi_\alpha'W_K\Psi_\alpha)^{-1}\Psi_\alpha'W_K[(\mathbf{I}_{d_P}-\Psi_P\Phi_{K,P})\Omega_{Pg}-\Psi_P\Phi_{K,0}\Omega_{0g}-(\Psi_P\Phi_{K,g}+\mathbf{I}_{d_P})\Psi_g\Omega_{gg}].$$

The desired result is a corollary of Eq. (A.6). □

Proof of Theorem 4.2. The asymptotic distribution with arbitrary $\hat{P}_0$ and W_1 follows from Theorem 4.1. As a corollary of Lemma A.2,

$$\Sigma_{1-MD}(\hat{P}_0,W_1)-\Sigma_{1-MD}(\hat{P},W_1)\text{ is PSD}, \qquad \text{(A.7)}$$

where

$$\Sigma_{1-MD}(\hat{P},W_1)$$
$$= (\Psi_\alpha' W_1 \Psi_\alpha)^{-1}\Psi_\alpha' W_1 \left[\begin{pmatrix} \mathbf{I}_{d_P} & -\Psi_P & -\Psi_g \end{pmatrix} \begin{pmatrix} \Omega_{PP} & \Omega_{PP} & \Omega_{Pg} \\ \Omega_{PP} & \Omega_{PP} & \Omega_{Pg} \\ \Omega_{Pg}' & \Omega_{Pg}' & \Omega_{gg} \end{pmatrix} \begin{pmatrix} \mathbf{I}_{d_P} \\ -\Psi_P' \\ -\Psi_g' \end{pmatrix} \right] W_1'\Psi_\alpha(\Psi_\alpha' W_1' \Psi_\alpha)^{-1}$$
$$= (\Psi_\alpha' W_1 \Psi_\alpha)^{-1}\Psi_\alpha' W_1 \left[\begin{pmatrix} \mathbf{I}_{d_P} - \Psi_P & -\Psi_g \end{pmatrix} \begin{pmatrix} \Omega_{PP} & \Omega_{Pg} \\ \Omega_{Pg}' & \Omega_{gg} \end{pmatrix} \begin{pmatrix} \mathbf{I}_{d_P} - \Psi_P' \\ -\Psi_g' \end{pmatrix} \right] W_1'\Psi_\alpha(\Psi_\alpha' W_1' \Psi_\alpha)^{-1}. \quad \text{(A.8)}$$

By Assumptions A.3 and A.5, the expression in brackets in RHS of Eq. (A.8) is non-singular. Then, standard arguments in GMM estimation (e.g. McFadden and Newey, 1994, page 2165) imply that W_1^* in Eq. (4.2) minimizes $\Sigma_{1-MD}(\hat{P},W_1)$. In other words,

$$\Sigma_{1-MD}(\hat{P},W_1) - \Sigma_{1-MD}(\hat{P},W_1^*) \text{ is PSD.} \quad \text{(A.9)}$$

By combining Eqs. (A.7) and (A.9), we conclude that $\Sigma_{1-MD}(\hat{P}_0,W_1) - \Sigma_{1-MD}(\hat{P},W_1^*)$ is PSD, as desired. Finally, Eq. (4.3) follows from plugging in this information. □

Proof of Theorem 4.3. As in the statement of optimality, let $\hat{\alpha}_{K-MD}$ denote the K-MD estimator with arbitrary initial CCP estimator $\hat{P}_0$ and arbitrary weight matrices $\{W_k : k \leq K\}$. By Theorem 4.1, $\sqrt{n}(\hat{\alpha}_{K-MD} - \alpha^*) \xrightarrow{d} N(\mathbf{0}_{d_\alpha}, \Sigma_{K-MD}(\hat{P}_0, \{W_k : k \leq K\}))$. As a corollary of Lemma A.2,

$$\Sigma_{K-MD}(\hat{P}_0, \{W_k : k \leq K\}) - \Sigma_{K-MD}(\hat{P}, \{W_k : k \leq K\}) \quad \text{is PSD.} \quad \text{(A.10)}$$

As in the statement of invariance, let $\hat{\alpha}^*_{K-MD}$ denote the K-MD estimator with initial CCP estimator $\hat{P}$ and weight matrices $\{W_k : k \leq K-1\}$ for steps $1,\ldots,K-1$ (if $K>1$), and the corresponding optimal weight matrix in step K. By Theorem 4.1, $\sqrt{n}(\hat{\alpha}^*_{K-MD} - \alpha^*) \xrightarrow{d} N(\mathbf{0}_{d_\alpha}, \Sigma_{K-MD}(\hat{P}, \{\{W_k : k \leq K-1\}, W_K^*\}))$. By definition of an optimal choice of W_K,

$$\Sigma_{K-MD}(\hat{P}, \{\{W_k : k \leq K\}\}) - \Sigma_{K-MD}(\hat{P}, \{\{W_k : k \leq K-1\}, W_K^*\}) \quad \text{is PSD.} \quad \text{(A.11)}$$

As a next step, we provide an explicit formula the optimal choice of W_K and we compute the resulting asymptotic variance $\Sigma_{K-MD}(\hat{P}, \{\{W_k : k \leq K-1\}, W_K^*\})$. To this end, consider the following derivation.

$$\Sigma_{K-MD}(\hat{P}, \{W_k : k \leq K\})$$
$$= \left\{ \begin{array}{c} (\Psi_\alpha' W_K \Psi_\alpha)^{-1}\Psi_\alpha' W_K \times \\ \begin{bmatrix} (\mathbf{I}_{d_P} - \Psi_P\Phi_{K,P})' \\ -(\Psi_P\Phi_{K,0})' \\ -((\mathbf{I}_{d_P} + \Psi_P\Phi_{K,g})\Psi_g)' \end{bmatrix}' \begin{pmatrix} \Omega_{PP} & \Omega_{PP} & \Omega_{Pg} \\ \Omega_{PP}' & \Omega_{PP} & \Omega_{Pg} \\ \Omega_{Pg}' & \Omega_{Pg}' & \Omega_{gg} \end{pmatrix} \begin{bmatrix} (\mathbf{I}_{d_P} - \Psi_P\Phi_{K,P})' \\ -(\Psi_P\Phi_{K,0})' \\ -((\mathbf{I}_{d_P} + \Psi_P\Phi_{K,g})\Psi_g)' \end{bmatrix} \\ \times W_K'\Psi_\alpha(\Psi_\alpha' W_K'\Psi_\alpha)^{-1} \end{array} \right\}$$
$$= (\Psi_\alpha' W_K \Psi_\alpha)^{-1}\Psi_\alpha' W_K \Delta_K W_K' \Psi_\alpha (\Psi_\alpha' W_K' \Psi_\alpha)^{-1},$$

where

$$\Delta_K \equiv \begin{bmatrix} (\mathbf{I}_{d_P} - \Psi_P\Phi_{K,P0})' \\ -((\mathbf{I}_{d_P} + \Psi_P\Phi_{K,g})\Psi_g)' \end{bmatrix}' \begin{pmatrix} \Omega_{PP} & \Omega_{Pg} \\ \Omega_{Pg}' & \Omega_{gg} \end{pmatrix} \begin{bmatrix} (\mathbf{I}_{d_P} - \Psi_P\Phi_{K,P0})' \\ -((\mathbf{I}_{d_P} + \Psi_P\Phi_{K,g})\Psi_g)' \end{bmatrix} \quad \text{(A.12)}$$

and $\{\Phi_{k,P0} : k \leq K\}$ is defined by $\Phi_{k,P0} \equiv \Phi_{k,P} + \Phi_{k,0}$ for $k \leq K$.

We now derive W_K^*. To this end, we first show that Δ_K is non-singular. Notice that $(\mathbf{I}_{d_P} - \Psi_P\Phi_{1,P0}) = (\mathbf{I}_{d_P} - \Psi_P)$.

In addition, for any $k \leq K-1$,

$$\begin{aligned}\mathbf{I}_{d_P} - \Psi_P \Phi_{k+1,P0} &= (\mathbf{I}_{d_P} - \Psi_P \Psi_\alpha (\Psi_\alpha' W_k \Psi_\alpha)^{-1} \Psi_\alpha' W_k) - \Psi_P (\mathbf{I}_{d_P} - \Psi_\alpha (\Psi_\alpha' W_k \Psi_\alpha)^{-1} \Psi_\alpha' W_k) \Psi_P \Phi_{k,P0} \\ &= (\mathbf{I}_{d_P} - \Psi_P) + \Psi_P (\mathbf{I}_{d_P} - \Psi_\alpha (\Psi_\alpha' W_k \Psi_\alpha)^{-1} \Psi_\alpha' W_k)(\mathbf{I}_{d_P} - \Psi_P \Phi_{k,P0}).\end{aligned}$$

From this and some algebra, we conclude that

$$\begin{aligned}(\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0}) &= \Lambda_K (\mathbf{I}_{d_P} - \Psi_P) \\ (\mathbf{I}_{d_P} + \Psi_P \Phi_{K,g}) \Psi_g &= \Lambda_K \Psi_g.\end{aligned} \tag{A.13}$$

By combining Eqs. (A.12) and (A.13), we deduce that

$$\Delta_K = \Lambda_K \begin{bmatrix} (\mathbf{I}_{d_P} - \Psi_P)' \\ -\Psi_g' \end{bmatrix}' \begin{bmatrix} \Omega_{PP} & \Omega_{Pg} \\ \Omega_{Pg}' & \Omega_{gg} \end{bmatrix} \begin{bmatrix} (\mathbf{I}_{d_P} - \Psi_P)' \\ -\Psi_g' \end{bmatrix} \Lambda_K'. \tag{A.14}$$

By Assumptions A.3 and A.5, and that Λ_K is non-singular, it follows that $\Delta_K \in \mathbb{R}$ is non-singular. Then, classical results in GMM estimation imply that the optimal weight matrix is

$$W_K^* \;=\; \Delta_K^{-1}, \tag{A.15}$$

resulting in an (optimal) asymptotic variance $\Sigma_{K-MD}(\hat{P}, \{\{W_k : k \leq K-1\}, W_K^*\}) = (\Psi_\alpha' \Delta_K^{-1} \Psi_\alpha)^{-1}$.

Since the choice of $\{W_k : k \leq K-1\}$ was completely arbitrary, the proof of invariance follows from showing that

$$\Psi_\alpha' \Delta_K^{-1} \Psi_\alpha = (\Sigma^*)^{-1}, \tag{A.16}$$

where we are using Assumptions A.3 and A.5 to deduce that Σ^* is non-singular. To this end, define the following matrices

$$\begin{aligned}A_K &\equiv (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0}) \Omega_{PP} (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0})' \\ B_K &\equiv \begin{bmatrix} (\mathbf{I}_{d_P} + \Psi_P \Phi_{K,g}) \Psi_g \Omega_{gg} \Psi_g' (\mathbf{I}_{d_P} + \Psi_P \Phi_{K,g})' \\ -(\mathbf{I}_{d_P} + \Psi_P \Phi_{K,g}) \Psi_g \Omega_{Pg}' (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0})' \\ -(\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0}) \Omega_{Pg} \Psi_g' (\mathbf{I}_{d_P} + \Psi_P \Phi_{K,g})' \end{bmatrix} \\ C_K &\equiv (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0})^{-1} B_K (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0})'^{-1},\end{aligned} \tag{A.17}$$

where we have used that $(\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0})$ is non-singular, which follows from Eq. (A.13), Assumption A.3, and that Λ_K is non-singular. In turn, this implies that A_K is non-singular and, in fact,

$$A_K^{-1} = (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0})'^{-1} \Omega_{PP}^{-1} (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0})^{-1}. \tag{A.18}$$

The following derivation proves that $C_K = C_1$.

$$\begin{aligned}C_K &= \begin{bmatrix} (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0})^{-1} (\mathbf{I}_{d_P} + \Psi_P \Phi_{K,g}) \Psi_g \Omega_{gg} \Psi_g' (\mathbf{I}_{d_P} + \Psi_P \Phi_{K,g})' (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0})'^{-1} \\ -(\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0})^{-1} (\mathbf{I}_{d_P} + \Psi_P \Phi_{K,g}) \Psi_g \Omega_{Pg}' - \Omega_{Pg} \Psi_g' (\mathbf{I}_{d_P} + \Psi_P \Phi_{K,g})' (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0})'^{-1} \end{bmatrix} \\ &= \begin{bmatrix} (\mathbf{I}_{d_P} - \Psi_P)^{-1} \Psi_g \Omega_{gg} \Psi_g' (\mathbf{I}_{d_P} - \Psi_P)'^{-1} \\ -(\mathbf{I}_{d_P} - \Psi_P)^{-1} \Psi_g \Omega_{Pg}' - \Omega_{Pg} \Psi_g' (\mathbf{I}_{d_P} - \Psi_P)'^{-1} \end{bmatrix} = C_1.\end{aligned}$$

where the first equality uses Eq. (A.17), the second equality uses Lemma A.1(b), and the final equality holds by Eq. (A.17) with $K = 1$.

We are now ready to prove Eq. (A.16), which follows immediately from the next derivation.

$$\begin{aligned}
\Psi_\alpha'\Delta_K^{-1}\Psi_\alpha &= \Psi_\alpha'(A_K+B_K)^{-1}\Psi_\alpha\\
&= \Psi_\alpha'(A_K^{-1}-(\mathbf{I}_{d_P}+A_K^{-1}B_K)^{-1}A_K^{-1}B_KA_K^{-1})\Psi_\alpha\\
&= \Psi_\alpha'(\mathbf{I}_{d_P}-\Psi_P\Phi_{K,P0})'^{-1}\Omega_{PP}^{-1}\{\Omega_{PP}-\Omega_{PP}(\mathbf{I}_{d_P}+\Omega_{PP}^{-1}C_K)^{-1}\Omega_{PP}^{-1}C_K\}\Omega_{PP}^{-1}(\mathbf{I}_{d_P}-\Psi_P\Phi_{K,P0})^{-1}\Psi_\alpha\\
&= \Psi_\alpha'(\mathbf{I}_{d_P}-\Psi_P)'^{-1}\Omega_{PP}^{-1}\{\Omega_{PP}-\Omega_{PP}(\mathbf{I}_{d_P}+\Omega_{PP}^{-1}C_1)^{-1}\Omega_{PP}^{-1}C_1\}\Omega_{PP}^{-1}(\mathbf{I}_{d_P}-\Psi_P)^{-1}\Psi_\alpha\\
&= \Psi_\alpha'\left(\begin{bmatrix}(\mathbf{I}_{d_P}-\Psi_P)'\\ -\Psi_g'\end{bmatrix}'\begin{bmatrix}\Omega_{PP} & \Omega_{Pg}\\ \Omega_{Pg}' & \Omega_{gg}\end{bmatrix}\begin{bmatrix}(\mathbf{I}_{d_P}-\Psi_P)'\\ -\Psi_g'\end{bmatrix}\right)^{-1}\Psi_\alpha = \Psi_\alpha'\Delta_1^{-1}\Psi_\alpha = (\Sigma^*)^{-1},
\end{aligned}$$

where the first equality follows from $\Delta_K = A_K + B_K$, which is implied by combining Eqs. (A.12) and (A.17), the second equality follows from Δ_K and A_K being non-singular, the third equality follows from Eqs. (A.17) and (A.18), the fourth equality is based on Lemma A.1(a) and $C_K = C_1$, the fifth equality follows from algebra and the final equality holds by Eq. (4.3).

Since the choice of $\{W_k : k \le K-1\}$ was completely arbitrary, the proof of optimality follows from showing the following argument. Note that

$$\Sigma_{K-MD}(\hat P_0,\{W_k:k\le K\})-\Sigma^* = \begin{bmatrix} (\Sigma_{K-MD}(\hat P_0,\{W_k:k\le K\})-\Sigma_{K-MD}(\hat P,\{W_k:k\le K\}))\\ +(\Sigma_{K-MD}(\hat P,\{W_k:k\le K\})-\Sigma_{K-MD}(\hat P,\{\tilde W_k:k\le K\}))\\ +\Sigma_{K-MD}(\hat P,\{\tilde W_k:k\le K\})-\Sigma^*\end{bmatrix}.$$

The RHS is the sum of three terms. The first term is PSD by Eq. (A.10), the second term is PSD by Eq. (A.11), and the third bracket is zero by Eq. (A.16). Then, $\Sigma_{K-MD}(\hat P_0,\{W_k : k\le K\})-\Sigma^*$ is PSD, as desired. □

A.2 Additional auxiliary results

Theorem A.1 (General result for iterative estimators). *Fix $K \ge 1$ arbitrarily. Assume Assumptions A.3, A.4, and A.7. Then, for all $k \le K$,*

$$\sqrt{n}(\hat\alpha_k-\alpha^*) = \begin{bmatrix}(A_k+B_k\Upsilon_{k,P}) & B_k\Upsilon_{k,0} & (B_k\Upsilon_{k,g}+C_k)\end{bmatrix}\sqrt{n}\begin{bmatrix}\hat P-P^*\\ \hat P_0-P^*\\ \hat g-g^*\end{bmatrix}+o_p(1) \tag{A.19}$$

$$\sqrt{n}(\hat P_{k-1}-P^*) = \begin{bmatrix}\Upsilon_{k,P} & \Upsilon_{k,0} & \Upsilon_{k,g}\end{bmatrix}\sqrt{n}\begin{bmatrix}\hat P-P^*\\ \hat P_0-P^*\\ \hat g-g^*\end{bmatrix}+o_p(1), \tag{A.20}$$

$$\sqrt{n}(\hat\theta_k-\theta^*) = \begin{bmatrix}(A_k+B_k\Upsilon_{k,P}) & B_k\Upsilon_{k,0} & (B_k\Upsilon_{k,g}+C_k)\\ \mathbf{0}_{d_g\times d_P} & \mathbf{0}_{d_g\times d_P} & \mathbf{I}_{d_g}\end{bmatrix}\sqrt{n}\begin{bmatrix}\hat P-P^*\\ \hat P_0-P^*\\ \hat g-g^*\end{bmatrix}+o_p(1), \tag{A.21}$$

where $\hat\theta_k\equiv(\hat\alpha_k,\hat g)$, $\theta^*\equiv(\alpha^*,g^*)$, $\{A_k : k\le K\}$, $\{B_k : k\le K\}$, and $\{C_k : k\le K\}$ are defined by:

$$\begin{aligned}
A_k &\equiv -\left(\frac{\partial^2 Q_k(\alpha^*,g^*,P^*)}{\partial\alpha\partial\alpha'}\right)^{-1}\Xi_k\\
B_k &\equiv -\left(\frac{\partial^2 Q_k(\alpha^*,g^*,P^*)}{\partial\alpha\partial\alpha'}\right)^{-1}\left(\frac{\partial^2 Q_k(\alpha^*,g^*,P^*)}{\partial\alpha\partial P'}\right)\\
C_k &\equiv -\left(\frac{\partial^2 Q_k(\alpha^*,g^*,P^*)}{\partial\alpha\partial\alpha'}\right)^{-1}\left(\frac{\partial^2 Q_k(\alpha^*,g^*,P^*)}{\partial\alpha\partial g'}\right),
\end{aligned} \tag{A.22}$$

and $\{\Upsilon_{k,P} : k \leq K\}$, $\{\Upsilon_{k,g} : k \leq K\}$, *and* $\{\Upsilon_{k,0} : k \leq K\}$ *are iteratively defined as follows. Set* $\Upsilon_{1,P} \equiv \mathbf{0}_{d_P \times d_P}$, $\Upsilon_{1,g} \equiv \mathbf{0}_{d_P \times d_g}$, $\Upsilon_{1,0} \equiv \mathbf{I}_{d_P}$ *and, for any* $k \leq K-1$,

$$\begin{aligned} \Upsilon_{k+1,P} &\equiv (\Psi_P + \Psi_\alpha B_k)\Upsilon_{k,P} + \Psi_\alpha A_k \\ \Upsilon_{k+1,0} &\equiv (\Psi_P + \Psi_\alpha B_k)\Upsilon_{k,0}. \\ \Upsilon_{k+1,g} &\equiv (\Psi_P + \Psi_\alpha B_k)\Upsilon_{k,g} + \Psi_g + \Psi_\alpha C_k. \end{aligned} \tag{A.23}$$

Proof. We divide the proof into three steps.

Step 1. Show that $(\hat{\alpha}_k, \hat{P}_{k-1}) = (\alpha^*, P^*) + o_p(1)$ for any $k \leq K$. We prove this by induction.

We begin with the initial step, i.e., show that the result holds for $k = 1$. First, $\hat{P}_0 = P^* + o_p(1)$ follows directly from Assumption A.4. Assumptions A.4 and A.7 imply that $\sup_{\alpha \in \Theta_\alpha} |\hat{Q}_1(\alpha, \hat{g}, \hat{P}_0) - Q_1(\alpha, g^*, P^*)| = o_p(1)$, $Q_1(\alpha, g^*, P^*)$ is upper semi-continuous function of α, and $Q_1(\alpha, g^*, P^*)$ is uniquely maximized at α^*. From these conditions, $\hat{\alpha}_1 = \alpha^* + o_p(1)$ follows from standard consistency results for extremum estimators (e.g. McFadden and Newey (1994)).

We next show the inductive step, i.e., assume that the result holds for $k \leq K-1$ and show that it holds for $k+1$. First, notice that:

$$\begin{aligned} \hat{P}_k - P^* &= \Psi(\hat{\alpha}_k, \hat{g}, \hat{P}_{k-1}) - \Psi(\alpha^*, g^*, P^*) \\ &= \Psi_\alpha(\alpha^*, g^*, P^*)(\hat{\alpha}_k - \alpha^*) + \Psi_g(\alpha^*, g^*, P^*)(\hat{g} - g^*) + \Psi_P(\alpha^*, g^*, P^*)(\hat{P}_{k-1} - P^*) + o_p(1) = o_p(1), \end{aligned}$$

where the second line follows from the intermediate value theorem, the inductive hypothesis, and Assumptions A.3 and A.4. Assumptions A.4 and A.7 imply that $\sup_{\alpha \in \Theta_\alpha} |\hat{Q}_{k+1}(\alpha, \hat{g}, \hat{P}_0) - Q_{k+1}(\alpha, g^*, P^*)| = o_p(1)$, $Q_{k+1}(\alpha, g^*, P^*)$ is upper semi-continuous function of α, and $Q_{k+1}(\alpha, g^*, P^*)$ is uniquely maximized at α^*. By repeating previous arguments, $\hat{\alpha}_{k+1} = \alpha^* + o_p(1)$ follows.

Step 2. Derive an expansion for $\sqrt{n}(\hat{\alpha}_k - \alpha^*)$ for any $k \leq K$.

For any $k \leq K$, consider the following derivation.

$$\begin{aligned} \mathbf{0}_{d_\alpha \times 1} &= \sqrt{n} \frac{\partial \hat{Q}_k(\hat{\alpha}_k, \hat{g}, \hat{P}_{k-1})}{\partial \alpha} + o_p(1) \\ &= \begin{bmatrix} \sqrt{n} \frac{\partial \hat{Q}_k(\alpha^*, g^*, P^*)}{\partial \alpha} + \frac{\partial^2 Q_k(\alpha^*, g^*, P^*)}{\partial \alpha \partial \alpha'} \sqrt{n}(\hat{\alpha}_k - \alpha^*) + \\ \frac{\partial^2 Q_k(\alpha^*, g^*, P^*)}{\partial \alpha \partial P'} \sqrt{n}(\hat{P}_{k-1} - P^*) + \frac{\partial^2 Q_k(\alpha^*, g^*, P^*)}{\partial \alpha \partial g'} \sqrt{n}(\hat{g} - g^*) \end{bmatrix} + o_p(1) \\ &= \begin{bmatrix} \Xi_k \sqrt{n}(\hat{P} - P^*) + \frac{\partial^2 Q_k(\alpha^*, g^*, P^*)}{\partial \alpha \partial \alpha'} \sqrt{n}(\hat{\alpha}_k - \alpha^*) + \\ \frac{\partial^2 Q_k(\alpha^*, g^*, P^*)}{\partial \alpha \partial P'} \sqrt{n}(\hat{P}_{k-1} - P^*) + \frac{\partial^2 Q_k(\alpha^*, g^*, P^*)}{\partial \alpha \partial g'} \sqrt{n}(\hat{g} - g^*), \end{bmatrix} + o_p(1), \end{aligned} \tag{A.24}$$

where the first line holds because $(\hat{\alpha}_k, \hat{g}, \hat{P}_{k-1}) = (\alpha^*, g^*, P^*) + o_p(1)$ (due to the step 1 and Assumption A.4), $\hat{\alpha}_k$ is the maximizer of $\hat{Q}_k(\alpha, \hat{g}, \hat{P}_{k-1})$ in Θ_α, and $\hat{\alpha}_k$ belongs to the interior of Θ_α with probability approaching one (due to the preliminary result and Assumption A.3), the second line holds by the intermediate value theorem and elementary convergence arguments based on Assumption A.7, and the third line holds by Assumption A.7.

We are now ready to derive the desired expansion.

$$\begin{aligned} \sqrt{n}(\hat{\alpha}_k - \alpha^*) &= \\ -\left(\frac{\partial^2 Q_k(\alpha^*, g^*, P^*)}{\partial \alpha \partial \alpha'}\right)^{-1} &\left[\Xi_k \sqrt{n}(\hat{P} - P^*) + \frac{\partial^2 \hat{Q}_k(\alpha^*, g^*, P^*)}{\partial \alpha \partial P'} \sqrt{n}(\hat{P}_{k-1} - P^*) + \frac{\partial^2 \hat{Q}_k(\alpha^*, g^*, P^*)}{\partial \alpha \partial g'} \sqrt{n}(\hat{g} - g^*)\right] \\ = \; & A_k \sqrt{n}(\hat{P} - P^*) + B_k \sqrt{n}(\hat{P}_{k-1} - P^*) + C_k \sqrt{n}(\hat{g} - g^*) + o_p(1), \end{aligned} \tag{A.25}$$

where the first line holds by Eq. (A.24) and Assumption A.7, and the second line holds by Eq. (A.22) and Assumption A.7.

Step 3. Show Eqs. (A.19), (A.20), and (A.21). Eq. (A.21) follows immediately from Eq. (A.19). Eqs. (A.19) and (A.20) are the result of the following inductive argument.

We begin with the initial step, i.e., show that the result holds for $k = 1$. By $\Upsilon_{1,P} \equiv \mathbf{0}_{d_P \times d_P}$, $\Upsilon_{1,g} \equiv \mathbf{0}_{d_P \times d_g}$, $\Upsilon_{1,0} \equiv \mathbf{I}_{d_P}$, Eq. (A.20) holds for $k = 1$. By the same argument and step 2, Eq. (A.19) holds for $k = 1$.

We next show the inductive step, i.e., assume that the result holds for $k \leq K - 1$ and show that it holds for $k + 1$. First, consider the following derivation:

$$\begin{aligned}
\sqrt{n}(\hat{P}_k - P^*) &= \Psi_\alpha \sqrt{n}(\hat{\alpha}_k - \alpha^*) + \Psi_g \sqrt{n}(\hat{g} - g^*) + \Psi_P \sqrt{n}(\hat{P}_{k-1} - P^*) + o_p(1) \\
&= \left[\begin{array}{c} [\Psi_\alpha A_k + (\Psi_P + \Psi_\alpha B_k)\Upsilon_{k,P}]\sqrt{n}(\hat{P} - P^*) + [\Psi_\alpha C_k + \Psi_g + (\Psi_P + \Psi_\alpha B_k)\Upsilon_{k,g}]\sqrt{n}(\hat{g} - g^*) \\ +(\Psi_P + \Psi_\alpha B_k)\Upsilon_{k,0}\sqrt{n}(\hat{P}_0 - P^*) \end{array} \right] + o_p(1) \\
&= \Upsilon_{k+1,P}\sqrt{n}(\hat{P} - P^*) + \Upsilon_{k+1,g}\sqrt{n}(\hat{g} - g^*) + \Upsilon_{k+1,0}\sqrt{n}(\hat{P}_0 - P^*) + o_p(1), \qquad \text{(A.26)}
\end{aligned}$$

where the first equality holds by $\hat{P}_k \equiv \Psi(\hat{\alpha}_k, \hat{g}, \hat{P}_{k-1})$, $P^* = \Psi(\alpha^*, g^*, P^*)$, Assumption A.3, and the intermediate value theorem, the second line holds by step 2 and the inductive hypothesis, and the last equality holds by Eq. (A.23). This verifies Eq. (A.20) for $k + 1$. Second, consider the following derivation:

$$\begin{aligned}
&\sqrt{n}(\hat{\alpha}_{k+1} - \alpha^*) = A_{k+1}\sqrt{n}(\hat{P} - P^*) + B_{k+1}\sqrt{n}(\hat{P}_k - P^*) + C_{k+1}\sqrt{n}(\hat{g} - g^*) + o_p(1) \\
&= (A_{k+1} + B_{k+1}\Upsilon_{k+1,P})\sqrt{n}(\hat{P} - P^*) + (C_{k+1} + B_{k+1}\Upsilon_{k+1,g})\sqrt{n}(\hat{g} - g^*) + B_{k+1}\Upsilon_{k+1,0}\sqrt{n}(\hat{P}_0 - P^*) + o_p(1),
\end{aligned}$$

where the first equality holds by step 2 and the second equality follows from Eq. (A.26). This verifies Eq. (A.19) for $k + 1$, and completes the proof. □

Lemma A.1. *Assume the conditions in Theorem 4.3. Let* $\{\Phi_{k,0} : k \leq K\}$, $\{\Phi_{k,P} : k \leq K\}$, *and* $\{\Phi_{k,g} : k \leq K\}$ *defined as in Eq. (4.1), and let* $\Phi_{k,P0} \equiv \Phi_{k,P} + \Phi_{k,0}$ *for all* $k \leq K$. *Then,*

1. $(\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0})^{-1}\Psi_\alpha = (\mathbf{I}_{d_P} - \Psi_P)^{-1}\Psi_\alpha$.
2. $(\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P0})^{-1}(\mathbf{I}_{d_P} + \Psi_P \Phi_{K,g}) = (\mathbf{I}_{d_P} - \Psi_P)^{-1}$.

Proof. Throughout this proof, denote $\Pi_k \equiv \Psi_\alpha(\Psi'_\alpha W_k \Psi_\alpha)^{-1}\Psi'_\alpha W_k$ for all $k \leq K$.

Part 1. It suffices to show that $(\mathbf{I}_{d_P} - \Psi_P \Phi_{k,P0})(\mathbf{I}_{d_P} - \Psi_P)^{-1}\Psi_\alpha = \Psi_\alpha$ for $k \leq K$. We show this by induction. The initial step follows from $\Phi_{k,P0} = \mathbf{I}_{d_P}$. We next show the inductive step, i.e., assume the result holds for $k \leq K-1$ and show it also holds for $k + 1$. Consider the following derivation.

$$\begin{aligned}
(\mathbf{I}_{d_P} - \Psi_P \Phi_{k+1,P0})(\mathbf{I}_{d_P} - \Psi_P)^{-1}\Psi_\alpha &= (\mathbf{I}_{d_P} - \Psi_P + \Psi_P(\mathbf{I}_{d_P} - \Pi_k)(\mathbf{I}_{d_P} - \Psi_P \Phi_{k,P0}))(\mathbf{I}_{d_P} - \Psi_P)^{-1}\Psi_\alpha \\
&= \Psi_\alpha + \Psi_P(\mathbf{I}_{d_P} - \Pi_k)(\mathbf{I}_{d_P} - \Psi_P \Phi_{k,P0})(\mathbf{I}_{d_P} - \Psi_P)^{-1}\Psi_\alpha \\
&= \Psi_\alpha + \Psi_P(\mathbf{I}_{d_P} - \Pi_k)\Psi_\alpha = \Psi_\alpha,
\end{aligned}$$

as required, where the first equality follows from Eq. (4.1) and some algebra, the second equality follows from the inductive hypothesis, and the final equality follows from $\Pi_k \Psi_\alpha = \Psi_\alpha$.

Part 2. It suffices to show that $(\mathbf{I}_{d_P} + \Psi_P \Phi_{k,g})(\mathbf{I}_{d_P} - \Psi_P) = (\mathbf{I}_{d_P} - \Psi_P \Phi_{k,P0})$ for $k \leq K$. We show this by induction. The initial step follows from $\Phi_{0,g} = \mathbf{0}_{d_P \times d_g}$ and $\Phi_{0,P} = \mathbf{I}_{d_P}$. We next show the inductive step, i.e., assume the result holds for $k \leq K - 1$ and show it also holds for $k + 1$. Consider the following derivation.

$$\begin{aligned}
(\mathbf{I}_{d_P} + \Psi_P \Phi_{k+1,g})(\mathbf{I}_{d_P} - \Psi_P) &= (\mathbf{I}_{d_P} - \Psi_P) + \Psi_P(\mathbf{I}_{d_P} - \Pi_k)(\Psi_P \Phi_{k,g} + \mathbf{I}_{d_P})(\mathbf{I}_{d_P} - \Psi_P) \\
&= (\mathbf{I}_{d_P} - \Psi_P) + \Psi_P(\mathbf{I}_{d_P} - \Pi_k)(\mathbf{I}_{d_P} - \Psi_P \Phi_{k,P0}) \\
&= -\Psi_P(\mathbf{I}_{d_P} - \Pi_k)\Psi_P \Phi_{k,P0} + \mathbf{I}_{d_P} - \Psi_P \Pi_k \\
&= \mathbf{I}_{d_P} - \Psi_P \Phi_{k+1,P0},
\end{aligned}$$

where the first and fourth equalities follows from Eq. (4.1), the second equality follows from the inductive hypothesis, and the third equality follows from algebra. □

Lemma A.2. *Under Assumptions A.4 and A.5,*

$$\begin{pmatrix} \Omega_{PP} & \Omega_{P0} & \Omega_{Pg} \\ \Omega_{P0}' & \Omega_{00} & \Omega_{0g} \\ \Omega_{Pg}' & \Omega_{0g}' & \Omega_{gg} \end{pmatrix} - \begin{pmatrix} \Omega_{PP} & \Omega_{PP} & \Omega_{Pg} \\ \Omega_{PP} & \Omega_{PP} & \Omega_{Pg} \\ \Omega_{Pg}' & \Omega_{Pg}' & \Omega_{gg} \end{pmatrix}$$

is PSD.

Proof. First, note that Assumption A.4 implies

$$\sqrt{n}\begin{pmatrix} \hat{P}_0 - P^* \\ \hat{g} - g^* \end{pmatrix} \xrightarrow{d} N\left(\begin{pmatrix} \mathbf{0}_{d_P} \\ \mathbf{0}_{d_g} \end{pmatrix}, \begin{pmatrix} \Omega_{00} & \Omega_{0g} \\ \Omega_{0g}' & \Omega_{gg} \end{pmatrix}\right).$$

Second, note that Assumption A.5 implies that

$$\begin{pmatrix} \Omega_{00} & \Omega_{0g} \\ \Omega_{0g}' & \Omega_{gg} \end{pmatrix} - \begin{pmatrix} \Omega_{PP} & \Omega_{Pg} \\ \Omega_{Pg}' & \Omega_{gg} \end{pmatrix}$$

is PSD, i.e, for any $\gamma_P \in \mathbb{R}^{d_P}$ and $\gamma_g \in \mathbb{R}^{d_g}$,

$$\gamma_P'(\Omega_{00} - \Omega_{PP})\gamma_P + 2\gamma_P'(\Omega_{0g} - \Omega_{Pg})\gamma_g \geq 0. \tag{A.27}$$

The remainder of this proof follows arguments similar to those used to show Hausman (1978, Lemma 2.1). Fix $r \in \mathbb{R}$ and $A \in \mathbb{R}^{d_P \times d_P}$ arbitrarily. By Assumption A.4,

$$\sqrt{n}\begin{pmatrix} \hat{P} + rA(\hat{P}_0 - \hat{P}) - P^* \\ \hat{g} - g^* \end{pmatrix} = \sqrt{n}\begin{pmatrix} (\mathbf{I}_{d_P} - rA)(\hat{P} - P^*) + rA(\hat{P}_0 - P^*) \\ \hat{g} - g^* \end{pmatrix} \xrightarrow{d} N\left(\begin{pmatrix} \mathbf{0}_{d_P} \\ \mathbf{0}_{d_g} \end{pmatrix}, \Sigma\right),$$

and

$$\Sigma \equiv \begin{pmatrix} \begin{pmatrix} r^2 A\Omega_{00}A' + (\mathbf{I}_{d_P} - rA)\Omega_{PP}(\mathbf{I}_{d_P} - rA') \\ +r(\mathbf{I}_{d_P} - rA)\Omega_{P0}A' + rA\Omega_{P0}'(\mathbf{I}_{d_P} - rA') \end{pmatrix} & rA(\Omega_{0g} - \Omega_{Pg}) + \Omega_{Pg} \\ r(\Omega_{0g} - \Omega_{Pg})'A' + \Omega_{Pg}' & \Omega_{gg} \end{pmatrix}.$$

Assumption A.5 implies that

$$\Sigma - \begin{pmatrix} \Omega_{PP} & \Omega_{Pg} \\ \Omega_{Pg}' & \Omega_{gg} \end{pmatrix} = \begin{pmatrix} \begin{pmatrix} r^2 A\Omega_{00}A' + (\mathbf{I}_{d_P} - rA)\Omega_{PP}(\mathbf{I}_{d_P} - rA') + \\ r(\mathbf{I}_{d_P} - rA)\Omega_{P0}A' + rA\Omega_{P0}'(\mathbf{I}_{d_P} - rA') - \Omega_{PP} \end{pmatrix} & rA(\Omega_{0g} - \Omega_{Pg}) \\ r(\Omega_{0g} - \Omega_{Pg})'A' & \mathbf{0}_{d_g \times d_g} \end{pmatrix}$$

is PSD, i.e., for any $\lambda_P \in \mathbb{R}^{d_P}$ and $\lambda_g \in \mathbb{R}^{d_g}$,

$$H(r) \equiv \begin{pmatrix} \lambda_P'(r^2 A\Omega_{00}A' + (\mathbf{I}_{d_P} - rA)\Omega_{PP}(\mathbf{I}_{d_P} - rA') + r(\mathbf{I}_{d_P} - rA)\Omega_{P0}A' + rA\Omega_{P0}'(\mathbf{I}_{d_P} - rA') - \Omega_{PP})\lambda_P \\ +2r\lambda_P'(\Omega_{0g} - \Omega_{Pg})'A'\lambda_g \end{pmatrix} \geq 0. \tag{A.28}$$

Note that $H(0) = 0$, i.e., $H(r)$ achieves a minimum at $r = 0$. Then, the first order condition for a minimization has to be satisfied at $r = 0$, which implies

$$H'(0) = \lambda_P'(\Omega_{P0}A' + A\Omega_{P0}' - A\Omega_{PP} - \Omega_{PP}A')\lambda_P + 2\lambda_g'(\Omega_{0g} - \Omega_{Pg})'A'\lambda_P = 0. \tag{A.29}$$

Since Eq. (A.29) has to hold for $\lambda_g = \mathbf{0}_{d_g}$, $A = \mathbf{I}_{d_P}$, and all $\lambda_P \in \mathbb{R}^{d_P}$, we deduce that

$$2\Omega_{PP} = \Omega_{P0} + \Omega_{P0}'. \tag{A.30}$$

Plugging this information into Eq. (A.29) yields

$$H'(0) = \lambda_P'((\Omega_{P0} - \Omega_{P0}')A' + A(\Omega_{P0}' - \Omega_{P0}))\lambda_P/2 + 2\lambda_P'(\Omega_{0g} - \Omega_{Pg})'A'\lambda_g = 0. \tag{A.31}$$

Since Eq. (A.31) has to hold for $\lambda_g = \mathbf{0}_{d_g}$, $A = \Omega_{P0} - \Omega_{P0}'$ and all $\lambda_P \in \mathbb{R}^{d_P}$, we deduce that $\Omega_{P0}' = \Omega_{P0}$. If we combine this with Eq. (A.30), we conclude that

$$\Omega_{PP} = \Omega_{P0} = \Omega_{P0}'. \tag{A.32}$$

Plugging this information into Eq. (A.31) yields

$$H'(0) = 2\lambda_P'(\Omega_{0g} - \Omega_{Pg})'A'\lambda_g = 0. \tag{A.33}$$

Since Eq. (A.33) has hold for $A = \mathbf{I}_{d_P}$ and all $\lambda_P \in \mathbb{R}^{d_P}$ and $\lambda_g \in \mathbb{R}^{d_g}$, we conclude that

$$\Omega_{0g} = \Omega_{Pg}. \tag{A.34}$$

For any $\mu = (\mu_P', \mu_0', \mu_g')'$ with $\mu_P, \mu_0 \in \mathbb{R}^{d_P}$ and $\mu_g \in \mathbb{R}^{d_g}$, consider the following argument.

$$\mu' \left[\begin{pmatrix} \Omega_{PP} & \Omega_{P0} & \Omega_{Pg} \\ \Omega_{P0}' & \Omega_{00} & \Omega_{0g} \\ \Omega_{Pg}' & \Omega_{0g}' & \Omega_{gg} \end{pmatrix} - \begin{pmatrix} \Omega_{PP} & \Omega_{PP} & \Omega_{Pg} \\ \Omega_{PP} & \Omega_{PP} & \Omega_{Pg} \\ \Omega_{Pg}' & \Omega_{Pg}' & \Omega_{gg} \end{pmatrix} \right] \mu = \mu_0'(\Omega_{00} - \Omega_{PP})\mu_0 \geq 0,$$

where the equality uses Eqs. (A.32) and (A.34), and the inequality uses Eq. (A.27) for $\gamma_P = \mu_0$ and $\gamma_g = \mathbf{0}_{d_g}$. Since the choice of μ was arbitrary, the desired result follows. □

A.3 Single-step estimation

By definition, single-step K-stage PI estimation is a special case of the two-step version in which the estimation of g^* is removed from the first step and is incorporated into the second step. To capture this within the notation of the paper, this section uses α^* to refer to $\theta^* = (\alpha^*, g^*)$ in the main text.

A.3.1 Estimation procedure

The estimation procedure for single-step estimation is as follows.

- **Step 1:** Estimate P^* with preliminary or 0-step estimators of the CCPs denoted by $\hat{P}_0$.
- **Step 2:** Estimate α^* with $\hat{\alpha}_K$, computed using the following algorithm. Initialize $k = 1$ and then:

  (a) Compute:

  $$\hat{\alpha}_k \equiv \arg\max_{\alpha \in \Theta_\alpha} \hat{Q}_k(\alpha, \hat{P}_{k-1}),$$

  where $\hat{Q}_k : \Theta_\alpha \times \Theta_P \to \mathbb{R}$ is the k-th step sample objective function. If $k = K$, exit the algorithm. If $k < K$, go to (b).

  (b) Estimate P^* with the k-step estimator of the CCPs, given by:

  $$\hat{P}_k \equiv \Psi(\hat{\alpha}_k, \hat{P}_{k-1}).$$

  Then, increase k by one unit and return to (a).

A.3.2 Assumptions

The assumptions for the single-step estimation are similar to the ones discussed in Section 2.3. In particular, Assumptions A.1 and A.6 remain the same, while Assumptions A.2-A.5 are modified as follows.

Assumption A.2'. (Identification) $\Psi(\alpha, P^*) = P^*$ *if and only if* $\alpha = \alpha^*$.

Assumption A.3'. (Regularity conditions) *Assume the following conditions:*

(i) α^* *belongs to the interior of* Θ_α.

(ii) $\sup_{\alpha\in\Theta_\alpha} |\Psi(\alpha, \tilde{P}) - \Psi(\alpha, P^*)| = o_p(1)$, *provided that* $\tilde{P} = P^* + o_p(1)$.

(iii) $\inf_{\alpha\in\Theta_\alpha} \Psi_{jax}(\alpha, \tilde{P}) > 0$ *for all* $(j, a, x) \in J \times A \times X$, *provided that* $\tilde{P} = P^* + o_p(1)$.

(iv) $\Psi(\alpha, P)$ *is twice continuously differentiable in a neighborhood of* (α^*, P^*). *We use* $\Psi_\lambda \equiv \partial\Psi(\alpha^*, P^*)/\partial\lambda$ *for* $\lambda \in \{\alpha, P\}$.

(v) $\mathbf{I}_{d_P} - \Psi_P \in \mathbb{R}^{d_P\times d_P}$ *and* $\Psi_\alpha \in \mathbb{R}^{d_P\times d_\alpha}$ *are full rank matrices.*

Assumption A.4'. (Baseline convergence) $(\hat{P}, \hat{P}_0)$ *satisfies the following condition:*

$$\sqrt{n}\begin{pmatrix} \hat{P} - P^* \\ \hat{P}_0 - P^* \end{pmatrix} \overset{d}{\to} N\left(\begin{pmatrix} \mathbf{0}_{d_P} \\ \mathbf{0}_{d_P} \end{pmatrix}, \begin{pmatrix} \Omega_{PP} & \Omega_{P0} \\ \Omega_{P0}' & \Omega_{00} \end{pmatrix}\right).$$

Assumption A.5'. (Baseline convergence II) *For any* $M \in \mathbb{R}^{d_P\times d_P}$, $(\mathbf{I}_{d_P} - M)\hat{P} + M\hat{P}_0$ *is not asymptotically more efficient than* $\hat{P}$.

A.3.3 Results for K-ML estimation

Theorem 3.1 derives the asymptotic distribution of the two-step K-ML estimator. The analogue result for the single-step K-ML estimator is as follows.

Theorem A.2 (Single-step K-ML)**.** *Fix* $K \geq 1$ *arbitrarily. Assume Assumption A.1, A.2', A.3', and A.4'. Then,*

$$\sqrt{n}(\hat{\alpha}_{K-ML} - \alpha^*) \overset{d}{\to} N(\mathbf{0}_{d_\alpha\times 1}, \Sigma_{K-ML}(\hat{P}_0)),$$

where

$$\Sigma_{K-ML}(\hat{P}_0) \equiv \left\{ \begin{array}{c} (\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1}\Psi_\alpha'\Omega_{PP}^{-1}\times \\ \begin{bmatrix} (\mathbf{I}_{d_P} - \Psi_P\Phi_{k,P})' \\ -(\Psi_P\Phi_{k,0})' \end{bmatrix}' \begin{pmatrix} \Omega_{PP} & \Omega_{P0} \\ \Omega_{P0}' & \Omega_{00} \end{pmatrix} \begin{bmatrix} (\mathbf{I}_{d_P} - \Psi_P\Phi_{k,P})' \\ -(\Psi_P\Phi_{k,0})' \end{bmatrix} \\ \times\Omega_{PP}^{-1}\Psi_\alpha(\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1} \end{array} \right\},$$

and $\{\Phi_{k,P} : k \leq K\}$ *and* $\{\Phi_{k,0} : k \leq K\}$ *are defined as follows. Set* $\Phi_{1,P} \equiv \mathbf{0}_{d_P\times d_P}$, $\Phi_{1,0} \equiv \mathbf{I}_{d_P}$ *and, for any* $k \leq K - 1$,

$$\begin{aligned} \Phi_{k+1,P} &\equiv (\mathbf{I}_{d_P} - \Psi_\alpha(\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1}\Psi_\alpha'\Omega_{PP}^{-1})\Psi_P\Phi_{k,P} + \Psi_\alpha(\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1}\Psi_\alpha'\Omega_{PP}^{-1}, \\ \Phi_{k+1,0} &\equiv (\mathbf{I}_{d_P} - \Psi_\alpha(\Psi_\alpha'\Omega_{PP}^{-1}\Psi_\alpha)^{-1}\Psi_\alpha'\Omega_{PP}^{-1})\Psi_P\Phi_{k,0}. \end{aligned}$$

A.3.4 Results for K-MD estimation

Theorem 4.1 derives the asymptotic distribution of the two-step K-MD estimator. The analogue result for the single-step K-MD estimator is as follows.

Theorem A.3 (Sigle-step K-MD)**.** *Fix $K \geq 1$ arbitrarily. Assume Assumption A.1, A.2', A.3', A.4', and A.6. Then,*

$$\sqrt{n}(\hat{\alpha}_{K-MD} - \alpha^*) \ \xrightarrow{d} \ N(\mathbf{0}_{d_\alpha \times 1}, \Sigma_{K-MD}(\hat{P}_0, \{W_k : k \leq K\})),$$

where

$$\Sigma_{K-MD}(\hat{P}_0, \{W_k : k \leq K\}) \ \equiv \ \left\{ \begin{array}{c} (\Psi_\alpha' W_K \Psi_\alpha)^{-1} \Psi_\alpha' W_K \times \\ \left[\begin{array}{c} (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P})' \\ -(\Psi_P \Phi_{K,0})' \end{array} \right]' \left(\begin{array}{cc} \Omega_{PP} & \Omega_{P0} \\ \Omega_{P0}' & \Omega_{00} \end{array} \right) \left[\begin{array}{c} (\mathbf{I}_{d_P} - \Psi_P \Phi_{K,P})' \\ -(\Psi_P \Phi_{K,0})' \end{array} \right] \\ \times W_K' \Psi_\alpha (\Psi_\alpha' W_K' \Psi_\alpha)^{-1} \end{array} \right\},$$

and $\{\Phi_{k,0} : k \leq K\}$ and $\{\Phi_{k,P} : k \leq K\}$ defined as follows. Set $\Phi_{1,P} \equiv \mathbf{0}_{d_P \times d_P}$, $\Phi_{1,0} \equiv \mathbf{I}_{d_P}$ and, for any $k \leq K-1$,

$$\begin{aligned} \Phi_{k+1,P} &\equiv (\mathbf{I}_{d_P} - \Psi_\alpha (\Psi_\alpha' W_k \Psi_\alpha)^{-1} \Psi_\alpha' W_k) \Psi_P \Phi_{k,P} + \Psi_\alpha (\Psi_\alpha' W_k \Psi_\alpha)^{-1} \Psi_\alpha' W_k, \\ \Phi_{k+1,0} &\equiv (\mathbf{I}_{d_P} - \Psi_\alpha (\Psi_\alpha' W_k \Psi_\alpha)^{-1} \Psi_\alpha' W_k) \Psi_P \Phi_{k,0}. \end{aligned}$$

Theorems 4.2 and 4.3 study optimality for the two-step K-MD estimator. The analogue result for the single-step K-MD estimator is as follows.

Theorem A.4 (Single-step optimality with $K = 1$)**.** *Assume Assumptions A.1, A.2', A.3', A.4', A.5', and A.6. Let $\hat{\alpha}^*_{1-MD}$ denote the 1-MD estimator with $\hat{P}_0 = \hat{P}$ and $W_1 = W_1^* \ \equiv \ [(\mathbf{I}_{d_P} - \Psi_P)\Omega_{PP}(\mathbf{I}_{d_P} - \Psi_P')]^{-1}$. Then,*

$$\sqrt{n}(\hat{\alpha}^*_{1-MD} - \alpha^*) \ \xrightarrow{d} \ N(\mathbf{0}_{d_\alpha \times 1}, \Sigma^*),$$

with

$$\Sigma^* \ \equiv \ (\Psi_\alpha' [(\mathbf{I}_{d_P} - \Psi_P)\Omega_{PP}(\mathbf{I}_{d_P} - \Psi_P')]^{-1} \Psi_\alpha)^{-1}. \tag{A.35}$$

Furthermore, $\Sigma_{1-MD}(\hat{P}_0, W_1) - \Sigma^*$ is positive semidefinite for all $(\hat{P}_0, W_1)$, i.e., $\hat{\alpha}^*_{1-MD}$ is optimal among all 1-MD estimators that satisfy our assumptions.

Theorem A.5 (Single-step invariance and optimality)**.** *Fix $K \geq 1$ arbitrarily and assume Assumptions A.1, A.2', A.3', A.4', A.5', and A.6. In addition, assume that the sequence of weight matrices $\{W_k : k \leq K-1\}$ is such that the matrix*

$$\Lambda_K \ \equiv \ \mathbf{I}_{d_P} \ + \ 1[K > 1] \sum_{b=1}^{K-1} \prod_{c=1}^{b} \Psi_P (\mathbf{I}_{d_P} - \Psi_\alpha (\Psi_\alpha' W_{K-c} \Psi_\alpha)^{-1} \Psi_\alpha' W_{K-c}) \in \mathbb{R}^{d_P \times d_P} \tag{A.36}$$

is non-singular. Then, we have the following two results.

1. *Invariance. Let $\hat{\alpha}^*_{K-MD}$ denote the K-MD estimator with $\hat{P}_0 = \hat{P}$, weight matrices $\{W_k : k \leq K-1\}$ for steps $1, \ldots, K-1$ (if $K > 1$), and the corresponding optimal weight matrix in step K. Then,*

   $$\sqrt{n}(\hat{\alpha}^*_{K-MD} - \alpha^*) \ \xrightarrow{d} \ N(\mathbf{0}_{d_\alpha \times 1}, \Sigma^*),$$

   where Σ^* is as in Eq. (A.35).

2. *Optimality. Let $\hat{\alpha}_{K-MD}$ denote the K-MD estimator with $\hat{P}_0$ and weight matrices $\{W_k : k \leq K\}$. Then,*

   $$\sqrt{n}(\hat{\alpha}_{K-MD} - \alpha^*) \ \xrightarrow{d} \ N(\mathbf{0}_{d_\alpha \times 1}, \Sigma_{K-MD}(\hat{P}_0, \{W_k : k \leq K\})).$$

   Furthermore, $\Sigma_{K-MD}(\hat{P}_0, \{W_k : k \leq K\}) - \Sigma^*$ is positive semidefinite, i.e., $\hat{\alpha}^*_{K-MD}$ is optimal among all K-MD estimators that satisfy our assumptions.

References

Aguirregabiria, V. and P. Mira (2002): “Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models,” *Econometrica*, 70, 1519–1543.

——— (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, 75, 1–53.

Amemiya, T. (1985): *Advanced Econometrics*, Harvard University Press.

Grund, B. (1993): “Kernel Estimators for Cell Probabilities,” *Journal of Multivariate Analysis*, 46, 283–308.

Hausman, J. (1978): “Specification Tests in Econometrics,” *Econometrica*, 46, 1251–71.

Kasahara, H. and K. Shimotsu (2008): “Pseudo-likelihood Estimation and Bootstrap Inference for Structural Discrete Markov Decision Models,” *Journal of Econometrics*, 146, 92–106.

——— (2012): “Sequential Estimation of Structural Models with a Fixed Point Constraint,” *Econometrica*, 80, 2303–2319.

McFadden, D. and W. K. Newey (1994): “Large sample estimation and hypothesis testing,” in *Handbook of Econometrics*, ed. by R. F. Engle and D. L. McFadden, Elsevier, vol. 4 of *Handbook of Econometrics*, 2111–2245.

Pakes, A., M. Ostrovsky, and S. Berry (2007): “Simple estimators for the parameters of discrete dynamic games (wiht entry/exit examples),” *RAND Journal of Economics*, 38, 373–399.

Pesendorfer, M. and P. Schmidt-Dengler (2008): “Asymptotic Least Squares Estimators for Dynamic Games,” *Review of Economic Studies*, 75, 901–928.

——— (2010): “Sequential Estimation of Dynamic Discrete Games: A Comment,” *Econometrica*, 78, 833–842.