k-step Estimation and Bootstrap-based Inference for Structural Discrete Markov Decision Models

Hiroyuki Kasahara
Department of Economics
University of Western Ontario

Katsumi Shimotsu
Department of Economics
Queen's University

October 6, 2005

Abstract

This paper analyzes the higher-order properties of k-step estimators and their practical implementations for the parametric discrete Markov decision models. We propose a new k-step estimator that can achieve the quadratic convergence without fully solving the fixed point problem in every step, thereby substantially reducing the computational cost. We then apply our k-step estimator to develop a k-step bootstrap procedure for the discrete Markov decision models and provide some Monte Carlo evidence based on a machine replacement model [c.f., Rust (1987); Cooper, Haltiwanger, and Power (1999)]. The proposed k-step bootstrap-t Wald tests perform well even with relatively small number of steps $k \geq 3$. Consistent with the theoretical results, the performance of our k-step bootstrap-t Wald tests are very similar to that of the k-step bootstrap-t Wald tests based on the NFXP algorithm developed by Rust (1987).

Preliminary and Incomplete

1 Introduction

This paper analyzes the higher-order properties of k-step estimators and their practical implementations for the parametric discrete Markov decision models in which the probability distribution is implicitly defined as a fixed point of the functional equation. We propose a new k-step estimator that can achieve the quadratic convergence without fully solving the fixed point problem in every step. We then apply our k-step estimator to develop a k-step bootstrap procedure for the discrete Markov decision models and provide some Monte Carlo evidence based on Rust's bus engine replacement model.

Studying the higher-order properties of k-step estimators and their implementations for the discrete Markov decision models is important for, at least, the following two reasons. First, we often encounter the situation in which the initial consistent but inefficient, perhaps in higher order, estimate is available. For example, Hotz and Millar (1993)'s conditional choice probability estimator is relatively easy to implement computationally, but the imprecise initial estimates of the conditional choice probabilities may lead to inaccurate estimates of the structural parameter. In such a situation, a k-step estimator may be used to achieve the higher-order asymptotic efficiency comparable to that of the maximum likelihood estimator [c.f, Andrews (1999b)]. Little is known, however, about the convergence properties of different k-step estimators in the discrete Markov decision models even though the higher-order efficiency of a k-step estimator crucially depends on its convergence rate.

Second, the simulation evidence in this paper reveals that the conventional asymptotic tests can be unreliable in finite samples in context of a discrete Markov decision model, indicating a serious need of a more reliable inference method such as bootstrapping. Applying the bootstrap to the discrete Markov decision models is often difficult, however; the standard bootstrap requires a model to be estimated repeatedly under different bootstrap samples while it is not unusual for each of the estimations of a discrete Markov decision model to take more than a day. Recently, Davidson and MacKinnon (1999) and Andrews (2002a) consider a computationally attractive k-step bootstrap estimator. Starting from the estimate based on the original sample, the k-step bootstrap estimator is obtained by taking the Newton-Raphson (NR) steps in the bootstrap sample. The k-step bootstrap may be feasible for many of the discrete Markov decision models when the number of NR steps, k, is small.

While the celebrated Nested Fixed Point (NFXP) algorithm developed by Rust (1987) can be applicable to develop a k-step estimator, it requires fully solving the fixed point problem for every step and its computational burden may limit its practical applicability. Building on the work by Hotz and Miller (1993) and Aguirregabiria and Mira (2002), we propose an alternative k-step estimator that *does not* require repeated solutions of fixed point problem and yet can achieve the same convergence rate as the NFXP estimator; in particular, with the NR steps, our k-step estimator achieves the quadratic convergence. Since the fixed point calculation occupies a large fraction of computational time, our k -step estimator is likely to be faster than the k-step estimator based on the NFXP algorithm.

Our k-step estimator is closely related to, but different from, the nested pseudo-likelihood (PL) estimator developed by Aguirregabiria and Mira (2002). The PL estimator avoids repeated solution of the dynamic programming problem but requires fully solving the pseudo-likelihood maximization problem for each iteration, which could be costly in models with many parameters. Furthermore, we show that the convergence rate of the PL estimator is less than quadratic. These properties contrast with those of our k-step estimator which requires neither repeated solution of the dynamic programming problem nor repeated solution of the pseudo-likelihood maximization problem and, yet, can achieve the quadratic convergence.

The superior performance of our estimator over the PL estimator is not without cost. The pseudo-likelihood of our estimator is defined in terms of 2 policy iterations while the pseudo-likelihood of the AM's PL estimator is defined in terms of 1 policy iteration. The computational cost for each NR step is larger in our estimator than in the PL estimator because our estimator involves one additional policy iteration. By defining the pseudo-likelihood in terms of 2 policy iterations, however, the NR search direction of our k-step estimator becomes the same as that of the NFXP estimator—of which pseudo-likelihood is defined in terms of infinitely many policy iterations—except for negligible higher-order terms.

Achieving better convergence properties is particularly important for developing k-step bootstrap estimators since, given the initial root-N consistent estimator together with some smoothness and moment conditions, the convergence rate determines the number of NR steps required to yield higher-order improvements over the first order asymptotics. A difference in required numbers of NR steps may lead to a substantial difference in computational time to implement k-step bootstrap.

We examine the performance of our k-step bootstrap estimator by conducting Monte Carlo experiments based on a machine replacement model [c.f., Rust (1987); Cooper, Haltiwanger, and Power (1999)]. The proposed k-step bootstrap-t Wald tests perform well even with a relatively small number of steps $k \geq 3$. Consistent with the theoretical results, the performance of our k-step bootstrap-t Wald tests are very similar to that of the k-step bootstrap-t Wald tests based on the NFXP algorithm.

The remainder of the paper is organized as follows. Section 2 introduces the model. In Section 3, we propose and analyze a modification to the PL estimator to achieve a faster rate of convergence. Section 4 describes our k-step estimation algorithm and proves its convergence properties. Section 5 analyzes the higher-order improvements from applying parametric bootstrapping the NFXP and k-step estimators. Section 6 reports the simulation results. Proofs and technical results are collected in Appendix A and B in Sections 7 and 8.

2 The Econometric Model

This section introduces a class of discrete Markov decision models we consider in this paper. We closely follows the setup and the notations of Aguirregabiria and Mira (2002) [AM, hereafter] and the reader is referred to AM for the details of the model. An agent maximizes the expected discounted sum of utilities, $E[\sum_{j=0}^{\infty} \beta^j U(s_{t+j}, a_{t+j})|a_t, s_t]$, where s_t is the vector of states and a_t is an action to choose from the discrete and finite set $A = \{1, 2, ..., J\}$. The transition probability of the state is given by $p(s_{t+1}|s_t, a_t)$. The Bellman equation for this dynamic optimization problem is written as

$$W(s_t) = \max_{a \in A} \left\{ u(s_t, a) + \beta \int W(s_{t+1}) p(ds_{t+1}|s_t, a) \right\}.$$

From the viewpoint of an econometrician, the state vector can be partitioned as $s_t = (x_t, \epsilon_t)$, where x_t is observable and ϵ_t is unobservable. We consider the following assumptions.

Assumption 1 (Additive Separability): The unobservable state variable ϵ_t is additively separable in utility function so that $U(s_t, a_t) = u(x_t, a_t) + \epsilon_t(a_t)$, where $\epsilon_t(a_t)$ is the a_t^{th} element of the unobservable state vector $\epsilon_t = \{\epsilon_t(a) : a \in A\}$.

Assumption 2 (Conditional Independence): The transition probability of the state variables can be written as $p(s_{t+1}|s_t, a_t) = g(\epsilon_{t+1}|x_{t+1}) f(x_{t+1}|x_t, a_t)$, where $g(\cdot)$ has finite first moments and is twice differentiable in ϵ; the support of $\epsilon(a)$ is the real line for all a.

Assumption 3: The observable state variable x_t has a compact support $X \subset \mathbb{R}^d$.

Assumptions 1 and 2 are identical to Assumptions 1 and 2 in AM. They are first introduced by Rust (1987) and widely used in the literature. Assumption 3 admits x_t to have a continuous distribution, relaxing Assumption 3 in AM that assumes x_t has a finite support.

Define the integrated value function $V(x) = \int W(x, \epsilon) g(d\epsilon|x)$, and let B_V be the space of $V \equiv \{V(x) : x \in X\}$. The Bellman equation can be rewritten in terms of this integrated value function as:

$$V(x) = \int \max_{a \in A} \left\{ u(x,a) + \epsilon(a) + \beta \int_X V(x') f(dx'|x,a) \right\} g(d\epsilon|x). \tag{1}$$

Let $\Gamma(\cdot)$ be the Bellman operator defined by the right-hand side of the above Bellman equation. The Bellman equation is compactly written as $V = \Gamma(V)$.

Let $P(a|x)$ denote the conditional choice probabilities of the action a given the observable state x, and let B_P be the space of $\{P(a|x) : x \in X\}$. Given the value function V, $P(a|x)$ is expressed as

$$P(a|x) = \int I \left\{ a = \arg\max_{j \in A} [v(x,j;V) + \epsilon(j)] \right\} g(d\epsilon|x), \tag{2}$$

where

$$v(x,a;V) = u(x,a) + \beta \int_X V(x') f(dx'|x,a)$$

is the choice-specific value function and $I(\cdot)$ is an indicator function. The right-hand side of the equation (2) can be viewed as a mapping from one Banach (B-) space B_V to another B-space B_P. Define this mapping $\Lambda(V) : B_V \to B_P$ as

$$[\Lambda(V)](a|x) \equiv \int I \left\{ a = \arg\max_{j \in A} [v(x,j;V) + \epsilon(j)] \right\} g(d\epsilon|x). \tag{3}$$

We now derive the mapping from choice probabilities to value functions based on Hotz and Miller (1993). First, the Bellman equation (1) can be rewritten as

$$V(x) = \sum_{a \in A} P(a|x) \left\{ u(x,a) + E[\epsilon(a)|x,a;\tilde{v},P] + \beta \int_X V(x') f(dx'|x,a) \right\} \tag{4}$$

where

$$E[\epsilon(a)|x,a;\tilde{v}_x,P(a|x)] = [P(a|x)]^{-1} \int \epsilon(a) I\{\tilde{v}(x,a;V) + \epsilon(a) \geq \tilde{v}(x,j;V) + \epsilon(j),\ j \in A\} g(d\epsilon|x),$$

where $\tilde{v}(x,a;V) = v(x,a;V) - v(x,1;V)$ and $\tilde{v}_x \equiv \{\tilde{v}(x,a;V) : a > 1\}$.

Define $P_x \equiv \{P(a|x) : a > 1\}$. For each x, there exists a mapping from the utility differences $\tilde{v}_x$ to the conditional choice probabilities P_x. Denote this mapping as $P_x = Q_x(\tilde{v}_x)$. Hotz and Miller (1993) showed that this mapping is invertible so that the utility differences can be expressed in terms of the conditional choice probabilities: $\tilde{v}_x = Q_x^{-1}(P_x)$. Then, we may express the conditional expectations of $\epsilon(a)$ in terms of the choice probabilities P_x as $e_x(a, P_x) \equiv E[\epsilon(a)|x, a; Q_x^{-1}(P_x), P(a|x)]$.[1]

By substituting these functions into (4), for each x, we obtain

$$V(x) = u_P(x) + \beta E_P V(x), \tag{5}$$

where

$$\begin{aligned} u_P(x) &= \sum_{a\in A} P(a|x)[u(x,a) + e_x(a, P_x)], \\ E_P V(x) &= \sum_{a\in A} P(a|x) \int_X V(x') f(dx'|x,a). \end{aligned}$$

Here, u_P is the expected utility function implied by the conditional choice probability P_x while E_P is the conditional expectation operator for the stochastic process $\{x_t, a_t\}$ induced by the conditional choice probability $P(a_t|x_t)$ and the transition density $f(x_{t+1}|x_t, a_t)$.

Define $P \equiv \{P_x : x \in X\}$. The value function implied by the conditional choice probability P is a unique solution to the linear operator equation (5):[2] $V = (I - \beta E_P)^{-1} u_P$. The right-hand side of this equation can be viewed as a mapping from the probability space to the value function space. Define this mapping as

$$\varphi(P) \equiv (I - \beta E_P)^{-1} u_P.$$

We may define a policy iteration operator Ψ as a composite operator of $\varphi(\cdot)$ and $\Lambda(\cdot)$:

$$P = \Psi(P) \equiv \Lambda(\varphi(P)).$$

The fixed point of the Bellman equation (1) can be expressed in terms of the fixed point of the policy iteration operator as $V = \varphi(P)$. If we consider the pair of sequences $\{V^K, P^K\}$

---

[1] When unobservables are independently distributed across alternatives and have extreme value distributions, we have $P(a|x) = Q_x^a(\tilde{v}) = \exp(\tilde{v}(x,a)/\sigma)/[1 + \sum_{j=2}^J \exp(\tilde{v}(x,j)/\sigma)]$ for $a > 1$. Then, inverting Q_x, we get $\tilde{v}(x,a) = Q_x^{a^{-1}}(P) = \ln[P(a|x)/\{1 - \sum_{j=2}^J P(j|x)\}]$. Also, we have $e(x,a;P) = \gamma - \ln(P(a|x))$ where γ is Euler's constant.

[2] Note that, for any $V_0 \in B_V$,

$$(I - \beta E_P)^{-1} V_0 = V_0 + \beta E_P V_0 + \beta^2 E_P^2 V_0 + ...$$

where the right hand side is converging because $\beta \in (0,1)$ and $E_P^k V_0$ is bounded for any integer k and any $V_0 \in B_V$.

recursively defined by $V^K = \varphi(P^K)$ and $P^{K+1} = \Lambda(V^K)$, then the sequence $\{V^K\}$ is the sequence of Newton iterations monotonically converging to the optimal solution to the Bellman equation (1).[3]

Before proceeding, we collect some definitions. Because P and V have an infinite dimension when x_t is continuously distributed, the derivatives of Ψ, Λ, and φ need to be defined as Fréchet (F-) derivatives. For a map $g : X \to Y$, where X and Y are B-spaces, g is F-differentiable at x iff there exists a linear and continuous map T such that

$$g(x+h) - g(x) = Th + o(||h||), \quad h \to 0$$

for all h in some neighborhood of zero, where $|| \cdot ||$ is an appropriate norm (e.g. sup norm, Euclidean norm if $g \in \mathbb{R}^M$). If it exists, this T is called the F-derivative of g at x, and we let $Dg(x)$ denote the F-derivative of g. Note that $Dg(x)$ is an operator. When x belongs to an Euclidean space, the F-derivative coincides with the standard derivative $dg(x)/dx$. Concepts such as the chain rule, product rule, higher-order and partial derivatives, and Taylor expansion are defined in an analogous manner to the corresponding concepts defined for the functions in Euclidean spaces. See Zeider (1986) for further discussion. Ichimura and Lee (2004) provide a concise summary on F-derivatives. Let $D^j g(x,y)$ denote the jth order F-derivative of $g(x,y)$, and let $D_x g(x,y)$ denote the partial F-derivative of $g(x,y)$ with respect to x. If x is a finite dimensional parameter, $D_x g(x,y)$ is equal to the standard partial derivative $\partial g(x,y)/\partial x$.

For matrix and (nonnegative) scalar sequences of random variables $\{X_N, N \geq 1\}$ and $\{Y_N, N \geq 1\}$, respectively, we write $X_N = O_p(Y_N)(o_p(Y_N))$ if $||X_N|| \leq CY_N$ for some (all) $C > 0$ with probability arbitrarily close to one for sufficiently large N.

We assume $P(a|x)$ is a continuous function of the continuously distributed elements of x. Let B_P denote the space of $P(a|x)$. One of the important properties of the policy iteration operator Ψ is that the derivative of Ψ in P is zero at the fixed point. AM prove this property in the case where the support of x_t is finite. As established in the following proposition, this zero-Jacobian property also holds when the support of x_t is not finite and V does not belong to a Euclidean space.

---

[3] Monotonicity can be proved as follows [c.f., Proposition 6.4.1 of Puterman (1994)]. Note that $u_{P^{K+1}} + \beta E_{P^{K+1}} V^K \geq u_P^K + \beta E_{P^K} V^K = V^K$. Then, $u_{P^{K+1}} \geq (I - \beta E_{P^{K+1}})V^K$. Multiplying both sides by $(I - \beta E_{P^{K+1}})^{-1}$, we have $V^{K+1} = (I - \beta E_{P^{K+1}})^{-1} u_{P^{K+1}} \geq V^K$.

Proposition 1 *Suppose Assumptions 1 - 3 hold. Then* $\varphi(\cdot)$ *is F-differentiable at the fixed point* P*. If* $\Psi(\cdot)$ *is F-differentiable at* P*, then* $D\varphi(\cdot) = D\Psi(\cdot) = 0$ *(zero operator) if evaluated at the fixed point* P*. In other words,* $D\varphi(P)\xi = D\Psi(P)\xi = 0$ *for any* $\xi \in B_P$*.*

3 Conditional Maximum Likelihood Estimator and its Variants

We consider a parametric model by assuming that the utility function and the transition probabilities are unknown up to a $L_\theta \times 1$ parameter vector $\theta \equiv (\theta_u, \theta_g, \theta_f)$, where θ_u, θ_g, and θ_f are parameter vectors in the utility function u, the density of unobservables g, and the conditional transition probability function f, respectively. Consequently, the policy iteration operator Ψ is parameterized as $\Psi(P, \theta) = \Lambda(\varphi(P, \theta), \theta)$. This corresponds to AM's notation $\Psi_\theta(P)$.

Let P_θ denote the fixed point of the policy iteration operator so that $P_\theta = \Psi(P_\theta, \theta)$. Let $\{(a_i, x_i) : i = 1, 2, ..., N\}$ be a random sample of (x, a) from the population. Under Assumption 2, the log-likelihood function can be decomposed into conditional choice probability and transition probability terms as:

$$l(\theta) = l_1(\theta) + l_2(\theta) = \sum_{i=1}^{N} \ln P_\theta(a_i|x_i) + \sum_{i=1}^{N} \ln f_{\theta_f}(x_i'|x_i, a_i).$$

Since θ_f can be estimated consistently without having to solve the Markov decision model, we focus on the estimation of $\alpha \equiv (\theta_u, \theta_g)$ given initial consistent estimates of θ_f from the likelihood $l_2(\theta)$. Thus, $\Psi(P, \theta) = \Psi(P, \alpha, \theta_f)$, and we use both $\Psi(P, \theta)$ and $\Psi(P, \alpha, \theta_f)$ henceforth.

The conditional maximum likelihood estimator solves the following constrained maximization problem:

$$\max_{\alpha} \frac{1}{N} \sum_{i=1}^{N} \ln P(a_i|x_i) \quad s.t. \quad P = \Psi(P, \alpha, \hat{\theta}_f). \tag{6}$$

Rust (1987) formulates the parameter restriction in terms of Bellman's equation. As shown in the previous section, we may rewrite the fixed point problem in the value function space as the fixed point problem in the policy function space (or the probability space) and define the estimator as (6). Denote the NFXP estimator, the solution to this maximization problem, and the associated conditional choice probability estimate by $\hat{\alpha}$ and $\hat{P}$.

3.1 Nested Pseudo-likelihood (PL) estimator

The nested pseudo-likelihood (PL) estimator developed by AM is recursively defined as follows. Assume an initial consistent estimator $\hat{P}_0$ is available. Given $\hat{P}^{PL}_{j-1}$, $\hat{\alpha}^{PL}_j$ is computed by

$$\hat{\alpha}^{PL}_j = \arg\max_{\alpha} \frac{1}{N}\sum_{i=1}^{N} \ln \Psi(\alpha, \hat{P}^{PL}_{j-1}, \hat{\theta}_f)(a_i|x_i).$$

Then $\hat{P}^{PL}_j$ is updated using the obtained estimate $\hat{\alpha}^{PL}_j$ as $\hat{P}^{PL}_j = \Psi(\hat{P}^{PL}_{j-1}, \hat{\alpha}^{PL}_j, \hat{\theta}_f)$. This process is iterated for $j = 1, \ldots, k$.

Let P^0 be the true set of conditional choice probabilities, and let f^0 be the true conditional transition probability of x. Let Θ_α and Θ_f be the set of possible values of α and θ_f, and define $\Theta = \Theta_\alpha \times \Theta_f$. Consider the following regularity conditions:

Assumption 4.

1. Θ_α and Θ_f are compact.
2. $\Psi(P, \alpha, \theta_f)$ is three times continuously F-differentiable.
3. $\Psi(P, \alpha, \theta_f)(a|x) > 0$ for any $(a, x) \in A \times X$ and any $\{P, \alpha, \theta_f\} \in B_P \times \Theta_\alpha \times \Theta_f$.
4. $W_i = \{a_i, x'_i, x_i\}$, for $i = 1, 2, \ldots, N$, are independently and identically distributed, and $dF(x) > 0$ for any x in the support of x_i, where $F(x)$ is the distribution function of x_i.
5. There is a unique $\theta^0_f \in$int(Θ_f) such that $f_{\theta^0_f}(x'|x, a) = f^0(x'|x, a)$ for all (x, x', a).
6. There is a unique $\alpha^0 \in$int(Θ_α) such that, for any $(a, x) \in A \times X$, $P_{(\alpha^0,\theta^0_f)}(a|x) = P^0(a|x)$. Furthermore, for any $\alpha \neq \alpha^0$, the set $\{(a, x) : \Psi(P^0, \alpha, \theta^0_f)(a|x) \neq P^0(a|x)\}$ is nonempty.
7. $E \sup_{(P,\alpha,\theta_f)} ||D^s\Psi(P, \alpha, \theta_f)(a_i|x_i)||^2 < \infty$ for $s = 0, \ldots, 3$.
8. $\hat{\theta}_f - \theta^0_f = O_p(N^{-1/2})$, $\hat{P}^{PL}_0 - P^0 = o_p(1)$, and the NFXP estimator $\hat{\alpha}$ satisfies $\sqrt{N}(\hat{\alpha} - \alpha^0) \to_d N(0, \Omega)$.

Assumptions 4(1)–4(6) are similar to the regularity conditions (a)-(f) in AM. The following proposition establishes that the PL estimator converges to the MLE, $\hat{\alpha}$, but its convergence rate is slower than quadratic.

Proposition 2 *Suppose Assumptions 1-4 hold. Then, for* $k = 1, 2, \ldots$

$$\begin{aligned}\hat{\alpha}_k^{PL} - \hat{\alpha} &= O_p(N^{-1/2}||\hat{P}_{k-1}^{PL} - \hat{P}|| + ||\hat{P}_{k-1}^{PL} - \hat{P}||^2), \\ \hat{P}_k^{PL} - \hat{P} &= O_p(N^{-1/2}||\hat{P}_{k-1}^{PL} - \hat{P}|| + ||\hat{P}_{k-1}^{PL} - \hat{P}||^2).\end{aligned}$$

Note that $\hat{P}_0^{PL} - P^0 = O_p(N^{-b})$ with $b > 1/4$ suffices for $\sqrt{N}(\hat{\alpha}_k^{PL} - \alpha^0) \to_d N(0, \Omega)$ for all $k \geq 1$. This weakens assumption (g) of Proposition 4 of AM and also implies that the PL estimator is valid even if x_t has a continuous support and a kernel-based estimator is used to estimate P^0. If $\hat{P}_0^{PL} - P^0 = O_p(N^{-b})$ with $b \in (1/4, 1/2]$, repeated substitution gives

$$\hat{\alpha}_k^{PL} - \hat{\alpha} = O_p(N^{-(k-1)/2-2b}), \quad \hat{P}_k^{PL} - \hat{P} = O_p(N^{-(k-1)/2-2b}).$$

In particular, if the support of x_t is finite and we can use $\hat{P}_0^{PL}$ such that $\hat{P}_0^{PL} - P^0 = O_p(N^{-1/2})$, then the convergence rate becomes $N^{-(k+1)/2}$.

3.2 The modified PL estimator

We introduce the nested modified pseudo-likelihood (MPL) estimator that achieves a faster rate of convergence than the PL estimator.

First, for a given $\hat{P}_{j-1}^{MPL}$, the MPL estimator $\hat{\alpha}_j^{MPL}$ is defined by

$$\hat{\alpha}_j^{MPL} = \arg\max_{\alpha} N^{-1} \sum_{i=1}^{N} \ln \Psi^2(\hat{P}_{j-1}^{MPL}, \alpha, \hat{\theta}_f)(a_i|x_i),$$

where $\Psi^2(P, \alpha, \theta_f)(a_i|x_i) \equiv \Psi(\Psi(P, \alpha, \theta_f), \alpha, \theta_f)(a_i|x_i)$. Then $\hat{P}_j^{MPL}$ is updated using the obtained estimate $\hat{\alpha}_j^{MPL}$ as $\hat{P}_j^{MPL} = \Psi(\hat{P}_{j-1}^{MPL}, \hat{\alpha}_j^{MPL}, \hat{\theta}_f)$. This process is iterated for $j = 1, \ldots, k$.

Using $\Psi^2(P, \alpha, \theta_f)$ instead of $\Psi(P, \alpha, \theta_f)$ achieves a faster rate of convergence. Of course, the MPL estimator requires more policy iterations than the PL estimator for computing each $\hat{\alpha}_j$. Therefore, the overall computational cost for achieving a given rate of convergence may be higher with the MPL estimator.

Assumption 5. (a) The set $\{(a, x) : \Psi(P, \alpha, \theta_f^0)(a|x) \neq P^0(a|x)\}$ is nonempty for any $(\alpha, P) \neq (\alpha^0, P^0)$. (b) $E \sup_{(P,\alpha,\theta_f)\in\Theta_0} ||D^s \Psi^2(P, \alpha, \theta_f)(a_i|x_i)||^2 < \infty$ for $s = 0, \ldots, 3$. (c) $\hat{P}_0^{MPL} - P^0 = o_p(1)$.

Assumption 5(a) is necessary for the identification of the MPL estimator.

Proposition 3 *Suppose Assumptions 1-5 hold. Then, for* $k = 1, 2, \ldots$

$$\hat{\alpha}_k^{MPL} - \hat{\alpha} = O_p(||\hat{P}_{k-1}^{MPL} - \hat{P}||^2), \quad \hat{P}_k^{MPL} - \hat{P} = O_p(||\hat{P}_{k-1}^{MPL} - \hat{P}||^2).$$

If $\hat{P}_0^{MPL} - P^0 = O_p(N^{-b})$ with $b \in (1/4, 1/2]$, then the convergence rate is given by

$$\hat{\alpha}_k^{MPL} - \hat{\alpha} = O_p(N^{-b2^k}), \quad \hat{P}_k^{MPL} - \hat{P} = O_p(N^{-b2^k}).$$

In particular, if $\hat{P}_0^{MPL} - P^0 = O_p(N^{-1/2})$, which is possible when X is finite or $\hat{P}_0^{MPL}$ is constructed as the fixed point of P_θ at a root-N consistent estimate of θ^0, then $\hat{\alpha}_k^{MPL} - \hat{\alpha} = O_p(N^{-2^{k-1}})$ and the MPL estimator achieves the quadratic convergence.

Remark 1 *Note that* $N^{-1}\sum_{i=1}^N (\partial/\partial\alpha) \ln \Psi(\hat{P}, \hat{\alpha}, \hat{\theta}_f)(a_i|x_i) = N^{-1}\sum_{i=1}^N (\partial/\partial\alpha) \ln \Psi^2(\hat{P}, \hat{\alpha}, \hat{\theta}_f)(a_i|x_i) = 0$ *because the NFXP estimator solves the maximization problem. Therefore, heuristically speaking, the distance between the NFXP estimator and the PL/MPL estimator is related with* $\hat{P}_{k-1}^{PL} - \hat{P}$, $\hat{P}_{k-1}^{MPL} - \hat{P}$, *and the behavior of the second derivatives* $N^{-1}\sum_{i=1}^N D_{P\alpha} \ln \Psi(P, \alpha, \hat{\theta}_f)(a_i|x_i)$ *and* $N^{-1}\sum_{i=1}^N D_{P\alpha} \ln \Psi^2(P, \alpha, \hat{\theta}_f)(a_i|x_i)$. *In the case of the MPL estimator, the Appendix shows* $D_{P\alpha} \ln \Psi^2(\hat{P}, \hat{\alpha}, \hat{\theta}_f)(a_i|x_i) = 0$, *and hence* $\hat{P}_{k-1}^{MPL} - \hat{P}$ *has a second-order effect on* $\hat{\alpha}_k^{MPL} - \hat{\alpha}$. *In the case of the PL estimator,* $D_{P\alpha} \ln \Psi^2(\hat{P}, \hat{\alpha}, \hat{\theta}_f)(a_i|x_i) \neq 0$ *in general but* $N^{-1}\sum_{i=1}^N D_{P\alpha} \ln \Psi(\hat{P}, \hat{\alpha}, \hat{\theta}_f)(a_i|x_i) = ED_{P\alpha} \ln \Psi(P^0, \alpha^0)(a_i|x_i) + O_p(N^{-1/2}) = O_p(N^{-1/2})$. *Therefore,* $\hat{P}_{k-1}^{PL} - \hat{P}$ *has a negligible effect on* $\hat{\alpha}_k^{PL} - \hat{\alpha}$ *in the limit but its convergence rate is slower.*

3.3 Covariance matrix estimation and test statistics

Suppose $\hat{\theta}_f$ is obtained by maximizing $l_2(\theta_f)$. Suppress $(a|x)$ and $(x'|x, a)$ from $P_\theta(a|x)$ and $f_{\theta_f}(x'|x, a)$. Expanding the first order condition for $\hat{\alpha}$ and $\hat{\theta}_f$ gives the asymptotic covariance matrix, $\Sigma(\theta^0)$, of $\hat{\theta} = (\hat{\alpha}', \hat{\theta}_f')'$ as

$$\Sigma(\theta^0) = D(\theta^0)^{-1} V(\theta^0) (D(\theta^0)^{-1})',$$

where

$$\begin{aligned} D(\theta) &= \begin{bmatrix} D_{11}(\theta) & D_{12}(\theta) \\ 0 & D_{22}(\theta) \end{bmatrix} = - \begin{bmatrix} E(\partial^2/\partial\alpha'\partial\alpha) \ln P_\theta & E(\partial^2/\partial\alpha'\partial\theta_f) \ln P_\theta \\ 0 & E(\partial^2/\partial\theta_f'\partial\theta_f) \ln f_{\theta_f} \end{bmatrix}, \\ V(\theta) &= \begin{bmatrix} V_{11}(\theta) & V_{12}(\theta) \\ V_{21}(\theta) & V_{22}(\theta) \end{bmatrix} = E\left[\begin{pmatrix} (\partial/\partial\alpha') \ln P_\theta \\ (\partial/\partial\theta_f') \ln f_{\theta_f} \end{pmatrix} \begin{pmatrix} (\partial/\partial\alpha') \ln P_\theta \\ (\partial/\partial\theta_f') \ln f_{\theta_f} \end{pmatrix}' \right]. \end{aligned}$$

The information matrix equality from the MLE based on $l_1(\theta)$ alone (when θ_f is known) and $l_2(\theta)$ alone implies $D_{11}(\theta^0) = V_{11}(\theta^0)$ and $D_{22}(\theta^0) = V_{22}(\theta^0)$, respectively. Furthermore, the information matrix from the full MLE based on $l(\theta)$ implies $-E(\partial^2/\partial\alpha'\partial\theta_f)\ln P_\theta = E(\partial/\partial\alpha')\ln P_\theta(\partial/\partial\theta_f)(\ln P_\theta + \ln f_{\theta_f})$.

There are several ways to estimate $\Sigma(\theta^0)$ consistently. Let

$$\begin{aligned} D_N(\theta) &= -\frac{1}{N}\sum_{i=1}^N \begin{bmatrix} (\partial^2/\partial\alpha'\partial\alpha)\ln P_\theta & (\partial^2/\partial\alpha'\partial\theta_f)\ln P_\theta \\ 0 & (\partial^2/\partial\theta_f'\partial\theta_f)\ln f_{\theta_f} \end{bmatrix}, \\ D_N^O(\theta) &= \frac{1}{N}\sum_{i=1}^N \begin{bmatrix} (\partial/\partial\alpha')\ln P_\theta(\partial/\partial\alpha)\ln P_\theta & (\partial/\partial\alpha')\ln P_\theta(\partial/\partial\theta_f)(\ln P_\theta + \ln f_{\theta_f}) \\ 0 & (\partial/\partial\theta_f')\ln f_{\theta_f}(\partial/\partial\theta_f)\ln f_{\theta_f} \end{bmatrix}, \\ V_N(\theta) &= \frac{1}{N}\sum_{i=1}^N \left[\begin{pmatrix} (\partial/\partial\alpha')\ln P_\theta \\ (\partial/\partial\theta_f')\ln f_{\theta_f} \end{pmatrix} \begin{pmatrix} (\partial/\partial\alpha')\ln P_\theta \\ (\partial/\partial\theta_f')\ln f_{\theta_f} \end{pmatrix}' \right]. \end{aligned}$$

$D_N^O(\theta)$ is an OPG estimator of $D(\theta)$, which does not require the calculation of the second derivatives of $\ln P_\theta$ and $\ln f_{\theta_f}$. Then one can use $\Sigma_N = \Sigma_N(\bar{\theta})$, where $\bar{\theta}$ is a consistent estimate of θ^0 and

$$\begin{aligned} \Sigma_N(\theta) &= D_N(\theta)^{-1}V_N(\theta)(D_N(\theta)^{-1})', \quad \text{or} \\ \Sigma_N(\theta) &= D_N^O(\theta)^{-1}V_N(\theta)(D_N^O(\theta)^{-1})'. \end{aligned} \tag{7}$$

$\Sigma_N(\bar{\theta}) \to_p \Sigma(\theta^0)$ from the standard argument, but this potentially requires a large number of policy iterations, being based on the full solution of the fixed point problem.

Alternatively, we may estimate $V(\theta)$ and $D(\theta)$ using the pseudo-likelihood function defining the PL and MPL estimators. Define $D_N^{PL}(P,\theta)$ and $D_N^{MPL}(P,\theta)$ by replacing P_θ in the definition of $D_N(\theta)$ with $\Psi(P,\theta)$ and $\Psi^2(P,\theta)$, respectively, and define $D_N^{O,PL}(P,\theta)$, $D_N^{O,MPL}(P,\theta)$, $V_N^{PL}(P,\theta)$, and $V_N^{MPL}(P,\theta)$ analogously. As shown below, we can estimate $\Sigma(\theta^0)$ consistently using these estimates with the PL and MPL estimators of P and α.

Proposition 4 *Let $\bar{P}$ and $\bar{\theta}$ denote estimators that converge to P^0 and θ^0 in probability. Then, $D_N^r(\bar{P},\bar{\theta}), D_N^{O,r}(\bar{P},\bar{\theta}) \to_p D(\theta^0)$ and $V_N^r(\bar{P},\bar{\theta}) \to_p V(\theta^0)$ for $r = \{PL, MPL\}$.*

This result is essential for developing inference procedures with PL and MPL estimation; using $V_N^{PL}(\hat{P}_k^{PL},\hat{\alpha}_k^{PL},\hat{\theta}_f)$ and $V_N^{MPL}(\hat{P}_k^{MPL},\hat{\alpha}_k^{MPL},\hat{\theta}_f)$ in place of $V_N(\hat{\alpha},\hat{\theta}_f)$, we may construct t- and Wald statistics with a limited number of policy iterations.

Let θ_r, θ_r^0, and $\hat{\theta}_r$ denote the r -th elements of θ, θ^0, and $\hat{\theta}$ respectively. Let $(\Sigma_N)_{rr}$ denote the (r,r)-th element of Σ_N. The t-statistic for testing the null hypothesis $H_0 : \theta_r = \theta_r^0$ is

$$T_N(\theta_r^0) = N^{1/2}(\hat{\theta}_r - \theta_r^0)/(\Sigma_N)_{rr}^{1/2}.$$

Let $\eta(\theta)$ be an R^{L_η}-valued function that is continuously differentiable at θ^0. The Wald statistic for testing $H_0 : \eta(\theta^0) = 0$ versus $H_A : \eta(\theta^0) \neq 0$ is

$$\begin{aligned} \mathcal{W}_N(\theta^0) &= H_N(\hat{\theta},\theta^0)' H_N(\hat{\theta},\theta^0), \quad \text{where} \\ H_N(\theta,\theta^0) &= \left(\frac{\partial}{\partial\theta}\eta(\theta)\Sigma_N(\theta)\frac{\partial}{\partial\theta'}\eta(\theta)\right)^{-1/2} N^{1/2}\eta(\theta). \end{aligned}$$

Then $T_N(\theta_r^0) \to_d N(0,1)$ and $\mathcal{W}_N(\theta^0) \to_d \chi^2_{L_\eta}$ under the null hypotheses.

4 k-step Estimation

We propose a k-step version of the PL and MPL estimators, which update the parameter α using one quasi-NR step. This reduces the computational cost of the corresponding estimators, because one does not have to fully solve the optimization problem. Let $L_N(P,\alpha,\theta_f)$ denote the objective function of the PL estimator as

$$L_N(P,\alpha,\theta_f) = \frac{1}{N}\sum_{i=1}^{N} \ln \Psi(P,\alpha,\theta_f)(a_i|x_i).$$

The k-step PL estimator, $(\tilde{\alpha}_k^{PL}, \tilde{P}_k^{PL})$, is defined recursively as:

Step 1: Given $(\tilde{P}_{j-1}^{PL}, \tilde{\alpha}_{j-1}^{PL}, \hat{\theta}_f)$, update α by

$$\tilde{\alpha}_j^{PL} = \tilde{\alpha}_{j-1}^{PL} - (Q_{N,j-1})^{-1}\frac{\partial}{\partial\alpha'}L_N(\tilde{P}_{j-1}^{PL}, \tilde{\alpha}_{j-1}^{PL}, \hat{\theta}_f), \tag{8}$$

where $Q_{N,j-1} = Q_N(\tilde{P}_{j-1}^{PL}, \tilde{\alpha}_{j-1}^{PL}, \hat{\theta}_f)$.

Step 2: Update P using the policy iteration operator evaluated at the updated $\tilde{\alpha}_j^{PL}$:

$$\tilde{P}_j^{PL} = \Psi(\tilde{P}_{j-1}^{PL}, \tilde{\alpha}_j^{PL}, \hat{\theta}_f).$$

Iterate Steps 1-2 until $j = k$.

The matrix $Q_{N,j-1}$ determines whether the k-step PL estimator uses the NR, default NR, line-search NR, or Gauss-Newton (GN) steps. The NR choice of $Q_{N,j-1}$ is $Q^{NR}_{N,j-1} = (\partial^2/\partial\alpha\partial\alpha')L_N(\tilde{P}^{PL}_{j-1}, \tilde{\alpha}^{PL}_{j-1}, \hat{\theta}_f$ The default NR choice of $Q_{N,j-1}$, denoted $Q^{D}_{N,j-1}$ equals $Q^{NR}_{N,j-1}$ if $\tilde{\alpha}^{PL}_{j}$ defined in (8) satisfies $L_N(\tilde{P}^{PL}_{j-1}, \tilde{\alpha}^{PL}_{j}, \hat{\theta}_f) \geq L_N(\tilde{P}^{PL}_{j-1}, \tilde{\alpha}^{PL}_{j-1}, \hat{\theta}_f)$, but equals some other matrix otherwise. Typically, $(1/\varepsilon)I_{\dim(\alpha)}$ for some small $\varepsilon > 0$ is used. The line-search NR choice, $Q^{LS}_{N,j-1}$, computes $\tilde{\alpha}^{PL,\alpha}_{j}$ for $\alpha \in (0,1]$ using $(1/\alpha)Q^{NR}_{N,j-1}$ and chooses the one that maximizes the objective function.

The GN choice, denoted $Q^{GN}_{N,j-1}$, uses a matrix that approximates the NR matrix $Q^{NR}_{N,j-1}$. A popular choice is the OPG estimator

$$Q^{OPG}_{N,j-1} = \frac{1}{N}\sum_{i=1}^{N} \frac{\partial}{\partial\alpha} \ln \Psi(\tilde{P}^{PL}_{j-1}, \tilde{\alpha}^{PL}_{j-1}, \hat{\theta}_f)(a_i|x_i) \frac{\partial}{\partial\alpha'} \ln \Psi(\tilde{P}^{PL}_{j-1}, \tilde{\alpha}^{PL}_{j-1}, \hat{\theta}_f)(a_i|x_i),$$

because this does not require the calculation of the second derivative of the objective function.

The following propositions establish that the k-step version of the PL estimator achieves the same rate of convergence as the original PL estimator.

Proposition 5 *Assume the regularity conditions of Proposition 2 hold and the initial estimates* $(\tilde{\alpha}^{PL}_{0}, \tilde{P}^{PL}_{0})$ *are consistent. Then, for* $k = 1, 2, \ldots,$

$$\begin{aligned}
\tilde{\alpha}^{PL}_{k} - \hat{\alpha} &= O_p(N^{-1/2}||\tilde{P}^{PL}_{k-1} - \hat{P}|| + ||\tilde{\alpha}^{PL}_{k-1} - \hat{\alpha}||^2 + ||\tilde{P}^{PL}_{k-1} - \hat{P}||^2), \\
&\quad + \begin{cases} 0, & \textit{for NR, default NR, line-search NR} \\ O_p(N^{-1/2}||\hat{\alpha} - \tilde{\alpha}^{PL}_{k-1}||), & \textit{for OPG,} \end{cases} \\
\tilde{P}^{PL}_{k} - \hat{P} &= O_p(||\tilde{\alpha}^{PL}_{k} - \hat{\alpha}||).
\end{aligned}$$

When the initial estimates satisfy $\tilde{\alpha}^{PL}_{0} - \alpha^0 = O_p(N^{-b})$ and $\tilde{P}^{PL}_{0} - P^0 = O_p(N^{-b})$ with $b \in (1/4, 1/2]$, repeated substitution gives

$$\tilde{\alpha}^{PL}_{k} - \hat{\alpha} = O_p(N^{-(k-1)/2-2b}), \quad \tilde{P}^{PL}_{k} - \hat{P} = O_p(N^{-(k-1)/2-2b}),$$

and the k-step PL estimator achieves the same convergence rate as the PL estimator.

The k-step MPL estimator $(\tilde{\alpha}^{MPL}_{k}, \tilde{P}^{MPL}_{k})$ is defined analogously using $N^{-1}\sum_{i=1}^{N} \ln \Psi^2(P, \alpha, \theta)(a_i|x_i)$ as $L_N(P, \alpha, \theta)$. As shown in the following proposition, it achieves the same rate of convergence as the MPL estimator when the NR, default NR, or line-search NR is used. When the OPG is used, its convergence rate reduces to that of the k-step PL estimator due to the error of $Q^{OPG}_{N,j-1}$ in approximating $(\partial^2/\partial\alpha\partial\alpha')L_N$.

Proposition 6 *Assume the regularity conditions of Proposition 3 hold and the initial estimates* $(\tilde{\alpha}_0^{MPL}, \tilde{P}_0^{MPL})$ *are consistent. Then, for* $k = 1, 2, \ldots,$

$$\begin{aligned}
\tilde{\alpha}_k^{MPL} - \hat{\alpha} &= O_p(||\hat{\alpha} - \tilde{\alpha}_{k-1}^{MPL}||^2 + ||\hat{P} - \tilde{P}_{k-1}^{MPL}||^2) \\
&\quad + \begin{cases} 0, & \textit{for NR, default NR, line-search NR} \\ O_p(N^{-1/2}||\hat{\alpha} - \tilde{\alpha}_{k-1}^{MPL}||), & \textit{for OPG,} \end{cases} \\
\tilde{P}_k^{MPL} - \hat{P} &= O_p(||\tilde{\alpha}_k^{MPL} - \hat{\alpha}||).
\end{aligned}$$

When the initial estimates satisfy $\tilde{\alpha}_0^{PL} - \alpha^0 = O_p(N^{-b})$ and $\tilde{P}_0^{PL} - P^0 = O_p(N^{-b})$ with $b \in (1/4, 1/2]$, repeated substitution gives

$$\begin{aligned}
\tilde{\alpha}_k^{MPL} - \hat{\alpha} &= O_p(N^{-b2^k}), \quad \tilde{P}_k^{MPL} - \hat{P} = O_p(N^{-b2^k}), \text{ for NR, default NR, line-search NR} \\
\tilde{\alpha}_k^{MPL} - \hat{\alpha} &= O_p(N^{-(k-1)/2-2b}), \quad \tilde{P}_k^{MPL} - \hat{P} = O_p(N^{-(k-1)/2-2b}), \text{ for OPG.}
\end{aligned}$$

5 Parametric Bootstrap and Higher-order Improvements

In this section, we analyze the higher-order improvements from applying parametric bootstrapping to the parametric discrete Markov decision models.

First, consider bootstrapping the NFXP estimator. The parametric bootstrap sample $\{W_i^* : i = 1, \ldots, n\}$ is generated using the parametric density at the (unrestricted) conditional NFXP estimator $\hat{\alpha}$ and the MLE $\hat{\theta}_f$. The conditional distribution of the bootstrap sample given $\hat{\theta} = (\hat{\alpha}', \hat{\theta}_f')'$ is the same as the distribution of the original sample except that the true parameter is $\hat{\theta}$ rather than $\theta^0 = (\alpha^{0\prime}, \theta_f^{0\prime})'$.

The bootstrap estimator $\theta^* = (\alpha^{*\prime}, \theta_f^{*\prime})'$ is defined exactly as the original estimator $\hat{\theta}$ but using the bootstrap sample $\{W_i^* : i = 1, \ldots, n\}$. Specifically,

$$\begin{aligned}
\theta_f^* &= \underset{\theta_f \in \Theta_f}{\arg\max}\, l_2^*(\theta_f), \quad \text{where } l_2^*(\theta_f) = \sum_{i=1}^N \ln f_{\theta_f}(x_i'^*|x_i^*, a_i^*), \qquad (9) \\
\alpha^* &= \underset{\alpha \in \Theta_\alpha}{\arg\max}\, \frac{1}{N}\sum_{i=1}^N \ln P(a_i^*|x_i^*) \quad s.t. \quad P = \Psi(P, \alpha, \theta_f^*).
\end{aligned}$$

The bootstrap covariance matrix estimator, Σ_N^*, is defined to be $\Sigma_N^*(\theta^*)$ where $\Sigma_N^*(\theta)$ has the same definition as $\Sigma_N(\theta)$ in (7) but with the bootstrap sample in place of the original sample.

The bootstrap t and Wald statistics are defined as

$$\begin{aligned} T_N^*(\hat{\theta}_r) &= N^{1/2}(\theta_r^* - \hat{\theta}_r)/(\Sigma_N^*)_{rr}^{1/2}, \\ \mathcal{W}_N^*(\hat{\theta}) &= H_N^*(\theta^*, \hat{\theta})' H_N^*(\theta^*, \hat{\theta}), \quad \text{where} \\ H_N^*(\theta, \hat{\theta}) &= \left(\frac{\partial}{\partial \theta} \eta(\theta) \Sigma_N^*(\theta) \frac{\partial}{\partial \theta'} \eta(\theta) \right)^{-1/2} N^{1/2}(\eta(\theta) - \eta(\hat{\theta})), \end{aligned} \tag{10}$$

where θ_r^* denotes the r-th element of θ^*, and $(\Sigma_N^*)_{rr}$ denotes the (r,r)-th element of Σ_N^*.

Let $z^*_{|T|,\alpha}$, $z^*_{T,\alpha}$, and $z^*_{\mathcal{W},\alpha}$ denote the $1-\alpha$ quantiles of $|T_N^*(\hat{\theta}_r)|$, $T_N^*(\hat{\theta}_r)$, and $\mathcal{W}_N^*(\hat{\theta})$, respectively. The symmetric two-sided bootstrap CI for θ_r^0 of confidence level $100(1-\alpha)\%$ is

$$CI_{SYM} = [\hat{\theta}_r - z^*_{|T|,\alpha}(\Sigma_N)_{rr}^{1/2}/N^{1/2}, \hat{\theta}_r + z^*_{|T|,\alpha}(\Sigma_N)_{rr}^{1/2}/N^{1/2}]. \tag{11}$$

The equal-tailed two-sided bootstrap CI for θ_r^0 of confidence level $100(1-\alpha)\%$ is

$$CI_{ET} = [\hat{\theta}_r - z^*_{T,\alpha/2}(\Sigma_N)_{rr}^{1/2}/N^{1/2}, \hat{\theta}_r - z^*_{T,1-\alpha/2}(\Sigma_N)_{rr}^{1/2}/N^{1/2}]. \tag{12}$$

The symmetric two-sided bootstrap t test of $H_0 : \theta_r = \theta_r^0$ versus $H_1 : \theta_r \neq \theta_r^0$ of significance level α rejects H_0 if $|T_N(\theta_r^0)| > z^*_{|T|,\alpha}$. The equal-tailed two-sided bootstrap t test of significance level α for the same hypotheses rejects H_0 if $T_N(\theta_r^0) < z^*_{T,1-\alpha/2}$ or $T_N(\theta_r^0) > z^*_{T,\alpha/2}$. The bootstrap Wald test of $H_0 : \eta(\theta^0) = 0$ versus $H_A : \eta(\theta^0) \neq 0$ rejects H_0 if $\mathcal{W}_N(\theta^0) > z^*_{\mathcal{W},\alpha}$.

We introduce the technical conditions that are used in establishing the higher-order improvements. They mainly consist of the conditions on the higher-order differentiability, the existence of the higher-order moments, and the Cramér condition. They are essentially the same as Assumptions 2-4 in Andrews (2001). Let a be a non-negative constant such that $2a$ is an integer. Let $\ell_\theta(x,a) = \ln P_\theta(a|x) + \ln f_{\theta_f}(x'|x,a)$, and let $h(W_i,\theta) \in R^{L_h}$ denote the vector containing the unique components of $(\partial/\partial\theta')\ell_\theta(a|x)$ and $(\partial/\partial\theta')\ell_\theta(a|x)(\partial/\partial\theta)\ell_\theta(a|x)$ and their partial derivatives with respect to θ through order $d = \max\{2a+2, 3\}$. Let $\lambda_{\min}(A)$ denote the smallest eigenvalue of a matrix A. Let $d(\theta, B)$ denote the distance between a point θ and a set B.

We assume the true parameter θ^0 lies in a subset Θ_0 of Θ and establish asymptotic refinements that hold uniformly for $\theta^0 \in \Theta_0$. For some $\delta > 0$, let $\Theta_1 = \{\theta \in \Theta : d(\theta, \Theta_0) < \delta/2\}$ and $\Theta_2 = \{\theta \in \Theta : d(\theta, \Theta_0) < \delta\}$ be slightly larger sets than Θ_0. For the reasons why these sets need to be considered, see Andrews (2001).

Assumption 6. (a) Θ_1 is an open set. (b) Given any $\varepsilon > 0$, there exists $\eta > 0$ such that $||\theta - \theta^0|| > \varepsilon$ implies that $E_{\theta^0} \ln P_{\theta^0}(a|x) - E_{\theta^0} \ln P_\theta(a|x) > \eta$ and $E_{\theta^0} \ln f_{\theta_f}(x'|x,a) -$

$E_{\theta^0} \ln f_{\theta_f}(x'|x,a) > \eta$ for all $\theta \in \Theta$ and $\theta^0 \in \Theta_1$. (c) $\sup_{\theta^0 \in \Theta_1} E_{\theta^0} \sup_{\theta \in \Theta} ||(\partial/\partial\theta) \ln P_\theta(a|x)||^{q_0} < \infty$, $\sup_{\theta^0 \in \Theta_1} E_{\theta^0} \sup_{\theta \in \Theta} ||(\partial/\partial\theta) \ln f_{\theta_f}(x'|x,a)||^{q_0} < \infty$, $\sup_{\theta_0 \in \Theta_1} E_{\theta^0} \sup_{\theta \in \Theta} |\ln P_\theta(a|x)|^{q_0} < \infty$, and $\sup_{\theta_0 \in \Theta_1} E_{\theta^0} \sup_{\theta \in \Theta} |\ln f_{\theta_f}(x'|x,a)|^{q_0} < \infty$ for all $\theta \in \Theta$ for $q_0 = \max\{2a+1, 2\}$.

Assumption 7. (a) $\ln P_\theta(a|x)$ and $\ln f_{\theta_f}(x'|x,a)$ are $d = \max\{2a+2, 3\}$ times partially differentiable with respect to θ on Θ_2 for all (x, x', a) in the support of (x_i, x_i', a_i). (b) $\sup_{\theta^0 \in \Theta_1} E_{\theta^0} ||h(W_i, \theta)||^{q_1} < \infty$ for some $q_1 > 2a+2$. (c) $\inf_{\theta^0 \in \Theta_1} \lambda \min(V(\theta_0)) > 0$, $\inf_{\theta^0 \in \Theta_1} \lambda \min(D(\theta_0)) > 0$. (d) There is a function $C_h(W_i)$ such that $||h(W_i,\theta) - h(W_i,\theta^0)|| \leq C_h(W_i)||\theta - \theta^0||$ for all $\theta \in \Theta_2$ and $\theta^0 \in \Theta_1$ such that $||\theta - \theta^0|| < \delta$ and $\sup_{\theta^0 \in \Theta_1} E_{\theta^0} C_h^{q_1}(W_i) < \infty$ for some $q_1 > 2a+2$.

Assumption 8. (a) For all $\varepsilon > 0$, there exists a positive δ such that for all $t \in R^{L_h}$ with $||t|| > \varepsilon$, $|E_{\theta^0} \exp(it'h(W_i, \theta^0))| \leq 1 - \delta$ for all $\theta^0 \in \Theta_1$. (b) $\Omega(\theta^0) = \text{Var}_{\theta^0}(h(W_i, \theta^0))$ has smallest eigenvalue bounded away from 0 over $\theta^0 \in \Theta_1$.

The higher-order differentiability of $\ln P_\theta(a|x)$ and $\ln f_{\theta_f}(x'|x,a)$ are satisfied if the density function of the unobserved state variable, ϵ, is sufficiently many times differentiable. Note that Assumption 2(b) of Andrews (2001) is satisfied by the definition of $\hat{\alpha}$ and $\hat{\theta}_f$. Assumption 2(c) of Andrews (2001) is satisfied with $\rho(\theta, \theta_0) = E_{\theta^0} \ln P_\theta(a|x)$ and $E_{\theta^0} \ln f_{\theta_f}(x'|x,a)$. Assumption 2(d) of Andrews (2001) is satisfied by the information inequality, the compactness of Θ, and the continuity of $\rho(\theta, \theta_0)$, which holds by our Assumption 6(b). See Lemma 2.4 of Newey and McFadden (1994). Because W_i is iid, Assumption 4 (a), (b), and (d) of Andrews (2001) are trivially satisfied, and his Assumption 4(c) reduces to the standard Cramér condition. Assumption 4 (f) of Andrews (2001) follows from our Assumption 8 (b) since W_i is iid.

The following Lemma establishes the higher-order improvements of the bootstrap NFXP estimator.

Lemma 7 *Suppose Assumptions 1-8 hold with a in Assumptions 6 and 7 as specified below. Then,*

(a) $\sup_{\theta^0 \in \Theta_0} |\Pr_{\theta^0}(\theta^0 \in CI_{SYM}) - (1-\alpha)| = O(N^{-2})$ *for* $a = 2$,

(b) $\sup_{\theta^0 \in \Theta_0} |\Pr_{\theta^0}(\theta^0 \in CI_{ET}) - (1-\alpha)| = o(N^{-1} \ln N)$ *for* $a = 1$,

(c) $\sup_{\theta^0 \in \Theta_0} |\Pr_{\theta^0}(\mathcal{W}_N(\theta^0) \leq z^*_{\mathcal{W},\alpha}) - (1-\alpha)| = o(N^{-3/2} \ln N)$ *for* $a = 3/2$.

The errors in coverage probability of standard delta method CIs are $O(N^{-1})$ and $O(N^{-1/2})$ for symmetric CIs and equal-tailed CIs, respectively. The errors in the standard Wald test is $O(N^{-1})$.

5.1 k-step parametric bootstrap

Bootstrapping the NFXP estimator is computationally very costly, because typically one has to estimate the model more than 150 times. We propose the k-step bootstrap PL and MPL estimators, which are defined as $\theta_k^{*PL} = ((\alpha_k^{*PL})', (\theta_f^*)')'$ and $\theta_k^{*MPL} = ((\alpha_k^{*MPL})', (\theta_f^*)')'$, where θ_f^* is defined in (9) and $(\alpha_k^{*PL}, P_k^{*PL}, \alpha_k^{*MPL}, P_k^{*MPL})$ are defined exactly as $(\tilde{\alpha}_k^{PL}, \tilde{P}_k^{PL}, \tilde{\alpha}_k^{MPL}, \tilde{P}_k^{MPL})$ but using the bootstrap sample $\{W_i^* : i = 1, \ldots, n\}$.

We estimate θ by the NFXP estimator in the original sample and use $P_{\hat{\theta}}$ (the fixed point at the NFXP estimator) as the initial estimate of P for the k-step estimation with the bootstrapped samples. Using the NFXP and $P_{\hat{\theta}}$ would not increase the computational burden significantly, because we have to estimate θ from the original sample and compute $P_{\hat{\theta}}$ only once.

We use the derivatives of the pseudo-likelihood function defining the PL or MPL estimator to construct the covariance matrix estimate. This saves the computational burden substantially, because we need to compute the covariance matrix estimates as many times as the number of bootstraps. However, a care must be exercised; using the second derivatives of the pseudo-likelihood function defining the PL estimator does not yield the higher-order refinement, because the second derivatives of $\ln P_\theta$ and $\ln \Psi(P, \theta)$ with respect to θ do not agree with each other even evaluated at the fixed point.

With θ_k^{*PL}, we compute the bootstrap covariance matrix estimator as

$$\Sigma_N^*(\theta_k^{*PL}) = D_N^{*O,PL}(\theta_k^{*PL})^{-1} V_N^{*PL}(\theta_k^{*PL}) (D_N^{*O,PL}(\theta_k^{*PL})^{-1})',$$

where $D_N^{*O,PL}(\theta)$ and $V_N^{*PL}(\theta)$ are the same as $D_N^{O,PL}(\theta)$ and $V_N^{PL}(\theta)$but constructed with the bootstrapped sample. With θ_k^{*MPL}, we use

$$\begin{aligned}
\Sigma_N^*(\theta_k^{*MPL}) &= D_N^{*MPL}(\theta_k^{*MPL})^{-1} V_N^{*MPL}(\theta_k^{*MPL}) (D_N^{*MPL}(\theta_k^{*MPL})^{-1})', \quad \text{or} \\
\Sigma_N^*(\theta_k^{*MPL}) &= D_N^{*O,MPL}(\theta_k^{*MPL})^{-1} V_N^{*MPL}(\theta_k^{*MPL}) (D_N^{*O,MPL}(\theta_k^{*MPL})^{-1})',
\end{aligned}$$

with analogous definitions for $D_N^{*MPL}(\theta)$, $V_N^{*MPL}(\theta)$, and $D_N^{*O,MPL}(\theta)$. $D_N^{*MPL}(\theta)$ must be used if $D_N(\theta)$ is used in forming $\Sigma_N(\theta)$, and $D_N^{*O,MPL}$ must be used if $D_N^O(\theta)$ is used in forming

$\Sigma_N(\theta)$.

The k-step bootstrap t- and Wald statistics, $T^*_{N,k}(\hat{\theta}_r)$ and $\mathcal{W}^*_{N,k}(\hat{\theta})$ are defined as in (10), but with (θ^*, Σ^*_N) replaced by $(\theta^{*PL}_k, \Sigma^*_N(\theta^{*PL}_k))$ or $(\theta^{*MPL}_k, \Sigma^*_N(\theta^{*MPL}_k))$. The k-step bootstrap confidence intervals, denoted $CI_{SYM,k}$, $CI_{ET,k}$, are defined as (11) and (12) but with $z^*_{|T|,\alpha}$, $z^*_{T,\alpha/2}$, and $z^*_{T,1-\alpha/2}$ replaced with $z^*_{|T|,k,\alpha}$, $z^*_{T,k,\alpha/2}$, and $z^*_{T,k,1-\alpha/2}$, respectively.

Define

$$\begin{aligned} \mu_{N,k} &= N^{-2^{k-1}} \ln^{2^k}(N) \text{ for the } k\text{-step MPL estimator with NR, defalult NR, and line-search NR,} \\ \mu_{N,k} &= N^{-(k+1)/2} \ln^{k+1}(N) \text{ for the } k\text{-step PL estimator and the } k\text{-step MPL estimator with OPG.} \end{aligned}$$

Lemma 8 establishes the higher-order equivalence of the k-step PL and MPL bootstrap estimators and NFXP bootstrap estimator. Lemma 9 shows, under suitable conditions on a and k, the difference between the bootstrap test statistic constructed using the k-step PL and MPL estimators and NFXP estimator are $o(N^{-a})$.

Lemma 8 *Suppose Assumptions 1-8 hold for some* $a > 0$ *with* $2a$ *an integer and* $\sup_{\theta\in\Theta} ||(\partial/\partial\theta)P_\theta(a_i|x_i)|| < \infty$ *with probability one. Then, for all* $\varepsilon > 0$ *and* $s = \{PL, MPL\}$,

$$\begin{aligned} \sup_{\theta^0\in\Theta_0} \Pr_{\theta_0}\left(\Pr^*_{\hat{\theta}}(||\theta^{*s}_k - \theta^*|| > \mu_{N,k}) > N^{-a}\varepsilon\right) &= o(N^{-a}), \\ \sup_{\theta^0\in\Theta_0} \Pr_{\theta_0}\left(\Pr^*_{\hat{\theta}}(|T^*_{N,k}(\hat{\theta}_r) - T^*_N(\hat{\theta}_r)| > N^{1/2}\mu_{N,k}) > N^{-a}\varepsilon\right) &= o(N^{-a}), \\ \sup_{\theta^0\in\Theta_0} \Pr_{\theta_0}\left(\Pr^*_{\hat{\theta}}(|\mathcal{W}^*_{N,k}(\hat{\theta}) - \mathcal{W}^*_{N,k}(\hat{\theta})| > N^{1/2}\mu_{N,k}) > N^{-a}\varepsilon\right) &= o(N^{-a}), \end{aligned}$$

Lemma 9 *Suppose the assumptions of Lemma 8 hold and* $\mu_{N,k} = o(N^{-(a+1/2)})$. *Then, for all* $\varepsilon > 0$,

$$\sup_{\theta^0\in\Theta_0} \Pr_{\theta_0}\left(\sup_z |\Xi_k(z)| > N^{-a}\varepsilon\right) = o(N^{-a}),$$

for $\Xi_k(z) = \Pr^*_{\hat{\theta}}(N^{1/2}(\theta^{*s}_k - \hat{\theta}) \leq z) - \Pr^*_{\hat{\theta}}(N^{1/2}(\theta^* - \hat{\theta}) \leq z)$ *with* $s = \{PL, MPL\}$, $\Pr^*_{\hat{\theta}}(T^*_{N,k}(\hat{\theta}_r) \leq z) - \Pr^*_{\hat{\theta}}(T^*_N(\hat{\theta}_r) \leq z)$, *or* $\Pr^*_{\hat{\theta}}(\mathcal{W}^*_{N,k}(\hat{\theta}) \leq z) - \Pr^*_{\hat{\theta}}(\mathcal{W}^*_N(\hat{\theta}) \leq z)$.

The condition on $\sup_{\theta\in\Theta} ||(\partial/\partial\theta)P_\theta(a_i|x_i)|| < \infty$ can be weakened to a condition in terms of its moment, but with a longer proof. The following Lemma shows that the errors in coverage probability of the k-step PL and MPL bootstrap CIs are the same as those of the NFXP bootstrap CIs. Therefore, one can use the k-step estimators to achieve the same level of higher-order refinement as the one obtained with the NFXP estimator.

Lemma 10 *Suppose the assumptions of Lemma 8 hold. Then*

(*a*) *If* $a = 2$ *and* $\mu_{N,k} = o(N^{-5/2})$, *then* $\sup_{\theta^0 \in \Theta_0} |\Pr_{\theta^0}(\theta^0 \in CI_{SYM,k}) - (1-\alpha)| = O(N^{-2})$,

(*b*) *If* $a = 1$ *and* $\mu_{N,k} = o(N^{-3/2})$, *then* $\sup_{\theta^0 \in \Theta_0} |\Pr_{\theta^0}(\theta^0 \in CI_{ET,k}) - (1-\alpha)| = o(N^{-1} \ln N)$,

(*c*) *If* $a = 3/2$ *and* $\mu_{N,k} = o(N^{-3/2})$, *then* $\sup_{\theta^0 \in \Theta_0} |\Pr_{\theta^0}(\mathcal{W}_N(\theta^0) \leq z^*_{\mathcal{W},\alpha}) - (1-\alpha)| = o(N^{-3/2} \ln N)$.

Using the (non-k-step) PL and MPL estimators achieves the same level of higher-order refinement. We omit the proof because it is very similar to the proof of Lemmas 8-10. The condition $\mu_{N,k} = o(N^{-5/2})$ requires $k \geq 3$ for the k-step MPL estimator with the NR, default NR, and line-search NR, and $k \geq 5$ for the k-step MPL estimator with the OPG and the k-step PL estimator. The k-step MPL estimator with $k = 3$ requires 6 policy iterations. On the other hand, the PL estimator with $k = 5$ requires 5 policy iterations and hence less computation. The fewest policy iterations with $\mu_{N,k} = o(N^{-5/2})$ are achieved if we use the PL estimator in the first and second steps and the MPL estimator in the third step; this yields $\mu_{N,k} = O(N^{-3} \ln^6(N))$ with 4 policy iterations.

6 Monte Carlo Experiments

6.1 Experimental Design

The model we consider is a version of machine replacement model [e.g., Rust (1987) and Cooper, Haltiwanger, and Power (1999)]. There are two observable state variables in the model: machine age $s_t \in \mathbf{N}$ and productivity shock $\omega_t \in \mathbf{R}$. The vector of the observed state variables is denoted by $x_t = (s_t, \omega_t)$. The machine replacement decision is indicated by $a_t \in \{0, 1\}$. The transition function of s_t is given by $s_t = a_{t-1} + (1 - a_{t-1})(s_{t-1} + 1)$ while the productivity shock ω_t follows an AR(1) process:

$$\omega_t = \rho \omega_{t-1} + \eta_t,$$

where $\eta_t \sim N(0, \sigma_\eta)$. The profit function is given by $u(x_t, a_t) + \epsilon(a_t)$ with

$$u(x_t, a_t) = y(s_t, \omega_t) - mc(s_t) - a_t rc(s_t),$$

where $y(s_t, \omega_t)$ is the revenue function; $c(s_t)$ is the machine maintenance cost; $rc(s_t)$ is the replacement cost; $\epsilon(a_t)$ is an unobserved state variable which is assumed to be extreme value distributed independently across alternatives, $a_t = 0, 1$. We consider the following parametric specification:

$$\begin{aligned} rc(s_t) &= \theta_0, \\ y(s_t, \omega_t) &= \exp(\theta_1 s_t + \omega_t), \\ mc(s_t) &= \theta_2 s_t. \end{aligned}$$

We estimate three structural parameters, θ_0, θ_1, and θ_2. Other parameters in the model, $(\beta, \rho, \sigma_\eta)$, are not estimated; instead, they are assumed to be known and are fixed at $(\beta, \rho, \sigma_\eta) = (0.96, 0.8, 0.2)$. The data are generated using the machine replacement model at the parameter values $(\theta_0, \theta_1, \theta_2) = (2, -0.1, 0.05)$.

To simulate the data from the model with continuous-state space, we first solve an approximated model with a discrete-state space using a finite number of grids and then use the "self-approximating" property of the Bellman operator [c.f., Rust (1996)] to evaluate conditional choice probabilities at any point outside of the grids.[4] This allows us to generate the data from the approximated model with continuous-state and to evaluate a likelihood function at any point outside of the grids. In practice, we approximate the state space of ω by 10 Gauss-Hermite quadrature grids while the state space of s_t is given by $\{1, 2, ..., 10\}$.[5]

6.2 The Performance of PL, MPL, and Various k-step Estimators and their t-Statistics

We first compare the performance of different k-step estimators in finite sample. We generate 1000 data sets, each of which has 500 observations, and then, for each of them, we obtain the sequence of estimates for various k-step estimators.[6] We experiment PL, MPL, and their k-step versions (k-PL and k-MPL) and compare the mean, the median, and the 95 percentile of their distances from the Maximum Likelihood (ML) estimate normalized by the standard error of the ML. In a k-step version of PL and MPL estimators, the parameter is updated in the first stage

---

[4] The "self-approximating" property of the Bellman operator is an example of the Nystrom interpolation formula [c.f., Atkinson (1997)].

[5] There is an important issue of how the choice of approximation methods affects estimation and inference. When we approximate the state space of ω using 5 quadrature grids instead of 10 grids, the results were similar.

[6] We use the conditional choice probabilities under the true parameter as an initial starting point.

using a Newton-Raphson method, where the hessian is approximated by the outer-products-of-gradient (OPG).

Table 1 reports the mean, the median, and the standard error of the percentage absolute distance between the ML estimate and the true parameter value. With this sample size, the MLE is biased and its standard errors are substantial.

Table 2 reports the mean, the median, and the 95 percentile of the percentage absolute distance between various estimators and the ML estimator normalized by the standard error of ML. The fixed point under the true parameter values are used as the initial conditional choice probabilities for PL, MPL, k-PL, and k-MPL. Comparing the statistics reported in Table 1 with those in Table 2 gives us the idea of how well the various k-step estimates approximate the ML estimate. As expected, the MPL perform best though the PL also perform fairly well. In terms of the 95 percentile of the percentage absolute distance ("Top 5 percent"), consistent with our theoretical prediction on convergence rate, the MPL for $k = 2$ is comparable to the PL for $k = 5$. On the other hand, k-PL and k-MPL do not perform well, especially for "Top 5 percent"; there are a substantial number of samples under which a simple NR-step does not lead to an increase in their values of likelihood, indicating that their initial starting points are not within the "domain of attraction."[7]

Table 3 reports the mean, the median, and the 95 percentile of the absolute distance in t-statistics between the various estimates and the ML estimate under the null of the true parameter values. The results are similar to those of Table 2. Both PL and MPL perform well while their k-step versions perform poorly especially for "Top 5 percent."

Table 4 reports the mean, the median, and the 5 and the 95 percentiles of t-statistics of MLE under the null of the true values of parameters. We notice that there is substantial asymmetry between left and right tails, especially for θ_1. This is indicative that asymptotic theory—which predicts a symmetric distribution—may not provide reliable inference in this context.

We repeat the similar exercise for various estimators with $k = 3$ as reported in Table 5; it compares the *finite sample distribution* of t-statistics constructed from MLE with those from other estimators at $k = 3$ in terms of mean, median, 5 and 95 percentiles. Here, we also consider k-step PL and MPL with a default NR method. A default NR uses the hessian if the updating

---

[7] In the future, we will consider introducing a line-search as well as repeating NR-steps until it achieves a decrease in the value of likelihood.

leads to an increase in the value of the (pseudo) likelihood function and, if not, it uses other matrix in place of the hessian. [8] The distributions of t-statistics constructed from PL and MPL approximate that from MLE very well. On the other hand, the distributions from k-PL and k-MPL approximate that from MLE less successfully, especially for the 5 and 95 percentile; again, this is because a simple NR-step may not work when an initial point is not within the domain of attraction; when we use a default NR-step instead, the distributions from k-PL and k-MPL approximate that from MLE reasonably well although the approximation error of the 95 percentile for θ_1 is still substantial.

6.3 k-step Parametric Bootstrapping

In this section, we apply our k-step estimator to the bootstrapping. We test the following composite null hypothesis:

$$H_0: \quad \theta_1 = 0.1 \quad \text{and} \quad \theta_2 = -0.05.$$

First, we examine how well the asymptotic Wald test performs. Table 6 reports rejection frequencies for asymptotic Wald test at .10, .05, and .01 levels for different sample sizes: $N = 200$, 500, and 1000. The number of replications is set to 500.[9] The parameters are estimated using the full fixed point solution method of the NFXP algorithm while we computed the variance-covariance matrices using the outer product of gradient (OPG) with the derivative of likelihood function.

The Wald test overrejects the null hypotheses at all three levels. While the severity of overrejection decreases with the sample size, it substantially overrejects at all levels even with the sample size of 1000. The result indicates, in context of the parametric discrete Markov decision model, that the inference based on the asymptotic Wald test could be misleading especially when the sample size is small.

---

[8] In practice, we use $(1/\epsilon)I$ with a small value of ϵ for this other matrix. We first try $\epsilon = 0.1$ and check if the likelihood value increases or not. If not, we try $\epsilon = (0.1)^2, (0.1)^3, \ldots$ until we find ϵ that leads to an increase in the value of the likelihood function.

[9] The panel data of a sample size N is generated as follows. We set the time dimension equal to $T = 5$ so that, given N, each sample contains N/T cross-sectional observations. Each of N/T cross-sectional observations is generated by first drawing the initial state x from the implied stationary distribution of x and then, starting from the initial state, we simulated T periods of observations.

Table 6 presents the coverage performance of the asymptotic 90% and 95% confidence intervals. For instance, the asymptotic 90% CI is constructed from a standard normal table as $\hat{\theta}_i \pm 1.645\hat{\sigma}_i$ for $i = 0, 1, 2$, where $\hat{\theta}_i$ and $\hat{\sigma}_i$ are the estimate of θ_i and its standard error, respectively. The variance-covariance matrices are estimated by OPG. The table indicates the frequencies that the confidence intervals missed the true values on the left and right sides in 500 simulated samples. In the case of the 90% CI, the true coverage is 0.9 so that the ideal values of "Miss Left" and "Miss Right" are 0.05.

For the parameter θ_1, both the 90 % and the 95% confidence intervals severely undercover on the right while they overcover on the left. On the other hand, the asymmetry of miscoverage for θ_0 and θ_2 is not as severe as that for θ_1.[10] The miscoverage for θ_1 is still substantial even at $N = 1000$ although, as we expect, the overall coverage performance of the asymptotic confidence interval gets better as the sample size increases from 200 to 1000.

We conduct parameteric bootstraps with 500 simulated samples, where each simulated sample contains $N = 500$ observations. For each simulated sample, we estimate the parameters by Maximum Likelihood (ML) using the NFXP algorithm and draw B=199 bootstrap samples from the parameteric model evaluated at the ML estimates. We estimate parameters for each bootstrap sample using PL, MPL, k-PL, and k-MPL estimators starting from the ML estimates and the conditional choice probabilities at the corresponding fixed points. Conducting simulations, we examine the performance of bootstrap Wald tests under PL, MPL, k-PL, and k-MPL relative to the asymptotic Wald tests.

Table 8 shows rejection frequencies for bootstrap Wald test at .10, .05, and .01 levels for ML, PL, MPL, k-PL, and k-MPL estimators. The variance-covariance matrices are estimated by OPG using the derivatives of a pseudo-likelihood function with one policy iteration for PL and k-PL while using the derivatives of a pseudo-likelihood function with two policy iteration for MPL and k-MPL. Comparing Table 8 with Table 6, we notice that the performance of bootstrap Wald tests using ML, PL, and MPL estimators are substantially better than that of asymptotic Wald test. Furthermore, bootstrap Wald tests using k-step PL and MPL perform better than asymptotic Wald test with a relatively small number of steps; k-PL with $k = 5$ and k-MPL with

---

[10] This may be due to the difference in the degree of nonlinearity. The parameter θ_1 enters into the profit function through exponential function while θ_0 and θ_1 are linearly related to the profit function; consequently, the degree of nonlinearly in θ_1 is larger than those in θ_0 and θ_2.

$k = 3$ lead to rejection frequencies at .10 and .05 that are much better than those of asymptotic Wald test.

Coverage performance of bootstrap 90% and 95% CIs for various estimators are reported in Tables 9-12. The miscoverage for θ_1 of bootstrap equal-tailed CIs is less severe than that of asymptotic CIs.

7 Appendix A: proofs

We give the proof for one dimensional case for both α, and θ_f. The generalization to a more general case is straightforward. We suppress the subscript N from the estimators of α and P to ease notation. For an n-linear operator $M(x_1, \ldots, x_n)$ such as an n -th F-derivative, the operator norm of M is defined as $||M|| = \sup_{||x_1|| = \cdots = ||x_n|| = 1} ||M(x_1, \ldots, x_n)||$.

To simplify the notation, let $\overline{\psi}_\alpha(P, \alpha, \theta_f) = N^{-1} \sum_{i=1}^N (\partial/\partial\alpha) \ln \Psi(P, \alpha, \theta_f)$ and $\overline{\psi}^2_\alpha(P, \alpha, \theta_f) = N^{-1} \sum_{i=1}^N (\partial/\partial\alpha) \ln \Psi^2(P, \alpha, \theta_f)$.

7.1 Proof of Proposition 1

For an arbitrary set of conditional choice probabilities P^0, define $H_x(P^0, V) = u_{P^0}(x) + \beta E P^0 V(x)$. Then, by definition,

$$\varphi(P^0)(x) = H_x(P^0, \varphi(P^0)) = u_P(x) + \beta \sum_{a \in A} P^0(a|x) \int_X \varphi(P^0)(x') f(dx'|x, a).$$

For $h \in B_P$, let $h(a|x)$ be the ath element of h corresponding to x. A little algebra gives

$$\begin{aligned} \varphi(P^0 + h)(x) - \varphi(P^0)(x) &= (1 - \beta E_{P^0})^{-1} \left[u_{P^0+h}(x) - u_{P^0}(x) + \beta \left(E_{P^0+h} - E_{P^0} \right) \varphi(P^0 + h)(x) \right] \\ &= (1 - \beta E_{P^0})^{-1} \left[u_{P^0+h}(x) - u_{P^0}(x) + \beta \left(E_{P^0+h} - E_{P^0} \right) \varphi(P^0)(x) \right], \end{aligned}$$

where the second line follows because $\varphi(\cdot)$ is continuous. Define $h_x^0 = (h(2|x), \ldots, h(J|x))'$, $p_x^0 = (P_x^0(2), \ldots, P_x^0(J))'$, where P_x is defined in page 5. Note that $P^0(1|x) = 1 - \sum_{a=2}^J P^0(a|x)$ and $(P^0 + h)(1|x) = 1 - \sum_{a=2}^J (P^0 + h)(a|x)$. Hence, it follows from Lemma 2 of AM that $u_{P^0+h}(x) - u_{P^0}(x) = \sum_{a \in A \setminus \{1\}} h(a|x)[u(x, a) - u(x, 1)] - Q_x^{-1}(p_x^0)' h_x^0 + o(||h||)$, where $o(||h||)$ is uniform in $x \in X$, because $Q_x^{-1}(z)$ is continuous in z uniformly in $x \in X$.

When P^0 is the fixed point of Ψ, $(E_{P^0+h} - E_{P^0})\,\varphi(P^0)(x)$ equals

$$\begin{aligned}
&\beta \sum_{a\in A\backslash\{1\}} h(a|x)\left\{\int_X V(x)f(dx'|x,a) - \int_X V(x)f(dx'|x,1)\right\} \\
= \;& \tilde{v}_x' h_x^0 - \sum_{a\in A\backslash\{1\}} h_x(a)[u(x,a)-u(x,1)].
\end{aligned}$$

because $\tilde{v}_x = Q_x(P_x)$ when P is the fixed point of Ψ, we obtain $\varphi(P+h)-\varphi(P) = o(||h||)$ and hence $D\varphi(P) = 0$. Since $\Psi = \Lambda\circ\varphi$, it follows from the chain rule in B-spaces that $D\Psi(P) = D\Lambda(\varphi(P))D\varphi(P) = 0$. $\square$

7.2 Proof of Proposition 2

Because the NFXP estimator maximizes the objective function of the PL estimator if $P = \hat{P}$ (c.f. equation (Ap.4) of AM p. 1540), it follows that

$$\overline{\psi}_\alpha(\hat{P}_{k-1}^{PL}, \hat{\alpha}_k^{PL}, \hat{\theta}_f) = \overline{\psi}_\alpha(\hat{P}, \hat{\alpha}, \hat{\theta}_f) = 0. \tag{13}$$

$\hat{\alpha}_k^{PL}$ is consistent for all $k \geq 1$, because the consistency proof in the proof of Proposition 4 of AM does not depend on the finiteness X. Applying the generalized Taylor's theorem to $\overline{\psi}_\alpha(\hat{P}_{k-1}^{PL}, \hat{\alpha}_k^{PL}, \hat{\theta}_f) - \overline{\psi}_\alpha(\hat{P}, \hat{\alpha}, \hat{\theta}_f)$ gives

$$\int_0^1 (\partial/\partial\alpha)\overline{\psi}_\alpha(P_\tau, \alpha_\tau, \hat{\theta}_f)(\hat{\alpha}_k^{PL} - \hat{\alpha})d\tau + \int_0^1 D_P\overline{\psi}_\alpha(P_\tau, \alpha_\tau, \hat{\theta}_f)(\hat{P}_{k-1}^{PL} - \hat{P})d\tau = 0 \tag{14}$$

where $P_\tau = \tau\hat{P}_{k-1}^{PL} + (1-\tau)\hat{P}$ and $\alpha_\tau = \tau\hat{\alpha}_k^{PL} + (1-\tau)\hat{\alpha}$. Note that $\partial P_\theta/\partial\theta = \partial\Psi(P_\theta,\theta)/\partial\theta$ thus $\hat{P} - P^0 = O_p(N^{-1/2})$ since $\hat{\alpha}$ and $\hat{\theta}_f$ are $N^{1/2}$-consistent. For the first term in the left of (14), $\int_0^1(\partial/\partial\alpha)\overline{\psi}_\alpha(P_\tau, \alpha_\tau, \hat{\theta}_f)d\tau \rightarrow_p E(\partial^2/\partial\alpha^2)\ln\Psi(P^0,\theta^0) < 0$ (the inequality follows from AM p.1541) from Lemma 11 and the consistency of $P_\tau, \hat{\theta}_f$, and $\hat{\alpha}_k^{PL}$. For the second term in the left of (14), view $D_P\overline{\psi}_\alpha(P_\tau, \alpha_\tau, \hat{\theta}_f)(\hat{P}_{k-1}^{PL} - \hat{P})$ as a function from (P, α, θ_f) to real line and use $\hat{\alpha} - \alpha^0, \hat{\theta}_f - \theta_f^0, \hat{P} - P^0 = O_p(N^{-1/2})$, then it follows that

$$\begin{aligned}
&\int_0^1 D_P\overline{\psi}_\alpha(P_\tau, \alpha_\tau, \hat{\theta}_f)(\hat{P}_{k-1}^{PL} - \hat{P})d\tau \\
= \;& D_P\overline{\psi}_\alpha(P^0,\theta^0)(\hat{P}_{k-1}^{PL} - \hat{P}) + O_p(||\hat{P}_{k-1}^{PL} - \hat{P}||^2) \\
& + O_p(||\hat{\alpha}_k^{PL} - \hat{\alpha}||||\hat{P}_{k-1}^{PL} - \hat{P}||) + O_p(N^{-1/2}||\hat{P}_{k-1}^{PL} - \hat{P}||).
\end{aligned}$$

Furthermore, from the information equality and Proposition 1, we can show $ED_p(\partial/\partial\alpha)\ln\Psi(P^0,\theta^0) = 0$ (zero operator). Since (a_i, x_i) is iid, it follows that

$$D_P\overline{\psi}_\alpha(P^0,\theta^0) = O_p(N^{-1/2}). \tag{15}$$

Consequently, rearranging the terms in (14) gives

$$\begin{aligned}&\Big[E(\partial^2/\partial\alpha^2)\ln\Psi(P^0,\theta^0)+O_p(||\hat{P}^{PL}_{k-1}-\hat{P}||)\Big](\hat{\alpha}^{PL}_k-\hat{\alpha})\\ =\;& O_p(N^{-1/2}||\hat{P}^{PL}_{k-1}-\hat{P}||)+O_p(||\hat{P}^{PL}_{k-1}-\hat{P}||^2),\end{aligned}$$

and hence $\hat{\alpha}^{PL}_k-\hat{\alpha}=O_p(N^{-1/2}||\hat{P}^{PL}_{k-1}-\hat{P}||+||\hat{P}^{PL}_{k-1}-\hat{P}||^2)$.

The proof completes by showing the convergence rate of $\hat{P}^{PL}_k=\Psi(\hat{P}^{PL}_{k-1},\hat{\alpha}^{PL}_k,\hat{\theta}_f)$. Expand $\Psi(\hat{P}^{PL}_{k-1},\hat{\alpha}^{PL}_k,\hat{\theta}_f)$ around $(\hat{P},\hat{\alpha},\hat{\theta}_f)$, apply $D_P\Psi(\hat{P},\hat{\alpha},\hat{\theta}_f)=0$, and use Lemma 11, then we obtain

$$\begin{aligned}\hat{P}^{PL}_k &= \Psi(\hat{P}^{PL}_{k-1},\hat{\alpha}^{PL}_k,\hat{\theta}_f)\\ &= \Psi(\hat{P},\hat{\alpha},\hat{\theta}_f)+D_P\Psi(\hat{P},\hat{\alpha},\hat{\theta}_f)(\hat{P}^{PL}_{k-1}-\hat{P})+D_\alpha\Psi(\hat{P},\hat{\alpha},\hat{\theta}_f)(\hat{\alpha}^{PL}_k-\hat{\alpha})\\ &\quad +O_p(||\hat{P}^{PL}_{k-1}-\hat{P}||^2)+O_p(||\hat{\alpha}^{PL}_k-\hat{\alpha}||^2)\\ &= \hat{P}+O_p(||\hat{\alpha}^{PL}_k-\hat{\alpha}||)+O_p(||\hat{P}^{PL}_{k-1}-\hat{P}||^2)+O_p(||\hat{\alpha}^{PL}_k-\hat{\alpha}||^2). \qquad (16)\end{aligned}$$

Therefore, $\hat{P}^{PL}_k-\hat{P}=O(||\hat{\alpha}^{PL}_k-\hat{\alpha}||+||\hat{P}^{PL}_{k-1}-\hat{P}||^2)=O_p(N^{-1/2}||\hat{P}^{PL}_{k-1}-\hat{P}||+||\hat{P}^{PL}_{k-1}-\hat{P}||^2)$.
□

7.3 Proof of Proposition 3

First, we collect some results on the derivatives of $\Psi(P,\theta)$ and P_θ. Recall that P_θ, the fixed point at θ, is defined implicitly as a function of θ as $P_\theta=\Psi(P_\theta,\theta)$. It follows from Proposition 1 that

$$DP_\theta=D_P\Psi(P_\theta,\theta)DP_\theta+D_\theta\Psi(P_\theta,\theta)=D_\theta\Psi(P_\theta,\theta). \qquad (17)$$

Note that DP_θ is a map from Θ to B_p. Since the maps from Θ to B_P constitutes a Banach space, we can define their F-derivatives. It follows from the chain rule and $D_P\Psi(P_\theta,\theta)=0$ that, for all $h\in\Theta$

$$\begin{aligned}D^2P_\theta h &= D_{PP}\Psi(P_\theta,\theta)DP_\theta h\cdot DP_\theta+D_{\theta P}\Psi(P_\theta,\theta)h\cdot DP_\theta\\ &\quad D_{P\theta}\Psi(P_\theta,\theta)\cdot DP_\theta h+D_{\theta\theta}\Psi(P_\theta,\theta)h. \qquad (18)\end{aligned}$$

Second, we collect some results on the derivatives of $\Psi^2(P,\theta)=\Psi(\Psi(P,\theta),\theta)$ evaluated at P that is *not necessarily* the fixed point of $\Psi(\cdot,\theta)$:

$$D_\theta\Psi^2(P,\theta)=D_P\Psi(\Psi(P,\theta),\theta)D_\theta\Psi(P,\theta)+D_\theta\Psi(\Psi(P,\theta),\theta), \qquad (19)$$

and, for all $h \in \Theta$

$$
\begin{aligned}
D_{\theta\theta}\Psi^2(P,\theta)h &= D_{PP}\Psi(\Psi(P,\theta),\theta)D_\theta\Psi(P,\theta)h \cdot D_\theta\Psi(P,\theta) \\
&\quad + D_{\theta P}\Psi(\Psi(P,\theta),\theta)h \cdot D_\theta\Psi(P,\theta) + D_P\Psi(\Psi(P,\theta),\theta)D_{\theta\theta}\Psi(P,\theta)h \\
&\quad + D_{P\theta}\Psi(\Psi(P,\theta),\theta)D_\theta\Psi(P,\theta)h + D_{\theta\theta}\Psi(\Psi(P,\theta),\theta)h. \qquad (20)
\end{aligned}
$$

Finally, the cross derivative of $\Psi^2(P,\theta)$ takes the form, for all $h \in B_p$

$$
\begin{aligned}
D_{P\theta}\Psi^2(P,\theta)h &= D_{PP}\Psi(\Psi(P,\theta),\theta)D_P\Psi(P,\theta)h \cdot D_\theta\Psi(P,\theta) \\
&\quad + D_P\Psi(\Psi(P,\theta),\theta)D_{P\theta}\Psi(P,\theta)h + D_{P\theta}\Psi(\Psi(P,\theta),\theta)D_P\Psi(P,\theta)h.
\end{aligned}
$$

Therefore, if evaluated at the fixed point $P = P_\theta$, we have

$$
D_\theta \ln \Psi^2(P_\theta,\theta) = D \ln P_\theta, \quad D_{\theta\theta} \ln \Psi^2(P_\theta,\theta) = D^2 \ln P_\theta, \quad D_{P\theta} \ln \Psi^2(P_\theta,\theta) = 0. \qquad (21)
$$

Now we proceed to prove the main result. First we prove the consistency of $\hat{\alpha}_k^{MPL}$. The identification follows from, if $\alpha \neq \alpha^0$,

$$
\begin{aligned}
\Pr(\Psi^2(P^0,\alpha,\theta_f^0) \neq P^0) &= \Pr(\Psi(\Psi(P^0,\alpha,\theta_f^0),\alpha,\theta_f^0) \neq P^0) \\
&\geq \Pr(\Psi(\Psi(P^0,\alpha,\theta_f^0),\alpha,\theta_f^0) \neq P^0 \text{ and } \Psi(P^0,\alpha,\theta_f^0) \neq P^0) \\
&= \Pr(\Psi(P,\alpha,\theta_f^0) \neq P^0 | P \neq P^0) \Pr(\Psi(P^0,\alpha,\theta_f^0) \neq P^0) \\
&> 0,
\end{aligned}
$$

and an argument similar to the proof of consistency of $\hat{\alpha}_k^{PL}$ by AM. From the first order condition for the NFXP estimator and (21), we have

$$
0 = \frac{1}{N}\sum_{i=1}^{N} \frac{\partial}{\partial\alpha} \ln P_{(\hat{\alpha},\hat{\theta}_f)} = \overline{\psi}_\alpha^2(\hat{P},\hat{\alpha},\hat{\theta}_f). \qquad (22)
$$

The MPL estimator $\hat{\alpha}_k^{MPL}$ satisfies the first order condition

$$
\overline{\psi}_\alpha^2(\hat{P}_{k-1}^{MPL},\hat{\alpha}_k^{MPL},\hat{\theta}_f) = 0. \qquad (23)
$$

Applying the generalized Taylor's theorem to (22) and (23) gives

$$
\int_0^1 (\partial/\partial\alpha)\overline{\psi}_\alpha^2(P_\tau,\alpha_\tau,\hat{\theta}_f)(\hat{\alpha}_k^{MPL} - \hat{\alpha})d\tau + \int_0^1 D_P\overline{\psi}_\alpha^2(P_\tau,\alpha_\tau,\hat{\theta}_f)(\hat{P}_{k-1}^{MPL} - \hat{P})d\tau = 0, \qquad (24)
$$

where $P_\tau = \tau \hat{P}_{k-1}^{MPL} + (1-\tau)\hat{P}$ and $\alpha_\tau = \tau \hat{\alpha}_k^{MPL} + (1-\tau)\hat{\alpha}$. For the first term on the left of (24), $\int_0^1 (\partial/\partial\alpha)\overline{\psi}_\alpha^2(P_\tau, \alpha_\tau, \hat{\theta}_f)d\tau \rightarrow_p E(\partial^2/\partial\alpha^2)\ln \Psi^2(P^0, \alpha^0, \theta_f^0) = E(\partial^2/\partial^2\alpha)\ln P_{\theta^0} < 0$ from the consistency of $\hat{\alpha}_k^{MPL}$ for all $k > 1$ and Lemma 11. For the second term in the left of (24), recall $D_{P\theta} \ln \Psi^2(\hat{P}, \hat{\alpha}, \hat{\theta}_f) = 0$ from (21) because $\hat{P}$ is the fixed point of $\Psi(\cdot, \hat{\alpha}, \hat{\theta}_f)$. Therefore, expanding $D_P\overline{\psi}_\alpha^2(P_\tau, \alpha_\tau, \hat{\theta}_f)$ around $(\hat{P}, \hat{\alpha}, \hat{\theta}_f)$ and using Lemma 11 , we obtain

$$\int_0^1 D_P\overline{\psi}_\alpha^2(P_\tau, \alpha_\tau, \hat{\theta}_f)(\hat{P}_{k-1}^{MPL} - \hat{P})d\tau = O_p(||\hat{\alpha}_k^{MPL} - \hat{\alpha}||||\hat{P}_{k-1}^{MPL} - \hat{P}||) + O_p(||\hat{P}_{k-1}^{MPL} - \hat{P}||^2).$$

Therefore, rearranging the terms in (24) gives

$$\left[E(\partial^2/\partial\alpha^2)\ln P_{\theta^0} + o_p(1)\right](\hat{\alpha}_k^{MPL} - \hat{\alpha}) = O_p(||\hat{P}_{k-1}^{MPL} - \hat{P}||^2),$$

and hence $\hat{\alpha}_k^{MPL} - \hat{\alpha} = O_p(||\hat{P}_{k-1}^{MPL} - \hat{P}||^2)$, establishing the quadratic convergence.

The proof completes by showing the convergence rate of $\hat{P}_k^{MPL} = \Psi(\hat{P}_{k-1}^{MPL}, \hat{\alpha}_k^{MPL}, \hat{\theta}_f)$. Using the same argument as (16), we obtain

$$\begin{aligned}\hat{P}_k^{MPL} &= \Psi(\hat{P}_{k-1}^{MPL}, \hat{\alpha}_k^{MPL}, \hat{\theta}_f) \\ &= \hat{P} + O_p(||\hat{\alpha}_k^{MPL} - \hat{\alpha}||) + O_p(||\hat{P}_{k-1}^{MPL} - \hat{P}||^2) + O_p(||\hat{\alpha}_k^{MPL} - \hat{\alpha}||^2).\end{aligned}$$

It follows that $\hat{P}_k^{MPL} - \hat{P} = O_p(||\hat{\alpha}_k^{MPL} - \hat{\alpha}|| + ||\hat{P}_{k-1}^{MPL} - \hat{P}||^2) = O_p(||\hat{P}_{k-1}^{MPL} - \hat{P}||^2)$. □

7.4 Proof of Proposition 4

First, consider an MLE based on $l_3(\theta) = N^{-1}\sum_{i=1}^N[\ln \Psi(P^0, \theta) + \ln f_{\theta_f}]$, which has the same limiting distribution as the full MLE. The information matrix equality associated with it implies $-E(\partial^2/\partial\alpha'\partial\theta_f)\ln \Psi(P^0, \theta^0) = E(\partial/\partial\alpha')\ln \Psi(P^0, \theta^0)(\partial/\partial\theta_f)(\ln \Psi(P^0, \theta^0) + \ln f_{\theta_f^0})$. Then, the required result for the $(1,2)$th block of $D_N^{PL}(\bar{P}, \bar{\theta})$ follows from Lemma 11, the information matrix equality and (17) as:

$$\begin{aligned}-\frac{1}{N}\sum_{i=1}^N \frac{\partial^2}{\partial\alpha'\partial\theta_f}\ln \Psi(\bar{P}, \bar{\theta}) &\rightarrow_p -E\frac{\partial^2}{\partial\alpha'\partial\theta_f}\ln \Psi(P^0, \theta^0) \\ &= E(\partial/\partial\alpha')\ln \Psi(P^0, \theta^0)(\partial/\partial\theta_f)(\ln \Psi(P^0, \theta^0) + \ln f_{\theta_f^0}) \\ &= E(\partial/\partial\alpha')\ln P_{\theta^0}(\partial/\partial\theta_f)(\ln P_{\theta^0} + \ln f_{\theta_f^0}) \\ &= -E(\partial^2/\partial\alpha'\partial\theta_f)\ln P_{\theta^0}.\end{aligned}$$

The proof for the $(1,1)$th block of $D_N^{PL}(\bar{P}, \bar{\theta})$ follows from the same argument, and the $(2,2)$ block $D_N^{PL}(\bar{P}, \bar{\theta})$ does not depend on $\bar{P}$. The proof for $D_N^{MPL}(\bar{P}, \bar{\theta})$ is similar, using (21) instead of (17). An analogous argument gives the proof for $D_N^{O,r}(\bar{P}, \bar{\theta})$ and $V_N^r(\bar{P}, \bar{\theta})$. □

7.5 Proof of Proposition 5

We prove the result for only the NR and OPG. The proof for the default NR and line-search NR is essentially same except for showing $\Pr(Q_N^D \neq Q_N^{NR}) \to 0$ and $\Pr(Q_N^{LS} \neq Q_N^{NR}) \to 0$; see the proof of Lemma 1 of Andrews (2001). We suppress the superscript PL from $\tilde{\alpha}_j^{PL}$ and $\tilde{P}_j^{PL}$, and we suppress $\hat{\theta}_f$ from $\overline{\psi}_\alpha(P,\alpha,\hat{\theta}_f)$ and $Q_N(P,\alpha,\hat{\theta}_f)$ when it does not lead to confusion.

Recall the NFXP estimator satisfies the first order condition $\overline{\psi}_\alpha(\hat{P},\hat{\alpha}) = 0$. Applying the generalized Taylor's theorem to $\overline{\psi}_\alpha(\hat{P},\hat{\alpha}) - \overline{\psi}_\alpha(\tilde{P}_{k-1},\tilde{\alpha}_{k-1})$ gives

$$\begin{aligned} 0 &= \overline{\psi}_\alpha(\tilde{P}_{k-1},\tilde{\alpha}_{k-1}) + D_\alpha\overline{\psi}_\alpha(\tilde{P}_{k-1},\tilde{\alpha}_{k-1})(\hat{\alpha}-\tilde{\alpha}_{k-1}) \\ &\quad + D_P\overline{\psi}_\alpha(\tilde{P}_{k-1},\tilde{\alpha}_{k-1})(\hat{P}-\tilde{P}_{k-1}) + R_{N,k} \\ &= \overline{\psi}_\alpha(\tilde{P}_{k-1},\tilde{\alpha}_{k-1}) + Q_N(\tilde{P}_{k-1},\tilde{\alpha}_{k-1})(\tilde{\alpha}_k-\tilde{\alpha}_{k-1}) + Q_N(\tilde{P}_{k-1},\tilde{\alpha}_{k-1})(\hat{\alpha}-\tilde{\alpha}_k) \\ &\quad + \left[D_\alpha\overline{\psi}_\alpha(\tilde{P}_{k-1},\tilde{\alpha}_{k-1}) - Q_N(\tilde{P}_{k-1},\tilde{\alpha}_{k-1})\right](\hat{\alpha}-\tilde{\alpha}_{k-1}) \\ &\quad + D_P\overline{\psi}_\alpha(\tilde{P}_{k-1},\tilde{\alpha}_{k-1})(\hat{P}-\tilde{P}_{k-1}) + R_{N,k}, \end{aligned} \tag{25}$$

where $R_{N,k} = O_p(||\hat{P}-\tilde{P}_{k-1}||^2 + ||\hat{\alpha}-\tilde{\alpha}_{k-1}||^2)$ from Lemma 11 (b). The first two terms on the right of (25) cancel out. For the fourth term on the right of (25), the term inside the bracket is 0 in the NR and $O_p(||\hat{P}-\tilde{P}_{k-1}|| + ||\hat{\alpha}-\tilde{\alpha}_{k-1}|| + N^{-1/2})$ in the OPG from Lemma 11 (e) (f) and the information matrix equality. For the fifth term on the right of (25), it follows from the generalized Taylor's theorem and (15) that

$$\begin{aligned} D_P\overline{\psi}_\alpha(\tilde{P}_{k-1},\tilde{\alpha}_{k-1},\hat{\theta}_f) &= D_P\overline{\psi}_\alpha(P^0,\alpha^0,\theta^0) + O_p(||\tilde{P}_{k-1}-\hat{P}||) + O_p(||\tilde{\alpha}_{k-1}-\hat{\alpha}||) + O_p(N^{-1/2}) \\ &= O_p(||\tilde{P}_{k-1}-\hat{P}||) + O_p(||\tilde{\alpha}_{k-1}-\hat{\alpha}||) + O_p(N^{-1/2}). \end{aligned}$$

Therefore,

$$\begin{aligned} Q_N(\tilde{P}_{k-1},\tilde{\alpha}_{k-1})(\hat{\alpha}-\tilde{\alpha}_k) &= O_p(N^{-1/2}||\hat{P}-\tilde{P}_{k-1}||) + O_p(||\hat{\alpha}-\tilde{\alpha}_{k-1}||^2 + ||\hat{P}-\tilde{P}_{k-1}||^2) \\ &\quad (+O_p(N^{-1/2}||\hat{\alpha}-\tilde{\alpha}_{k-1}|| \text{ for } \text{ OPG}). \end{aligned}$$

The stated bound of $\tilde{\alpha}_k - \hat{\alpha}$ follows from $Q_N(\tilde{P}_{k-1},\tilde{\alpha}_{k-1}) \to_p E(\partial^2/\partial\alpha^2)\ln\Psi(P^0,\theta^0) < 0$.

The proof completes by showing the bound of $\tilde{P}_k - \hat{P}$. Similarly to the proof of Proposition 2, expanding $\Psi(\tilde{P}_{k-1},\tilde{\alpha}_k)$ around $(\hat{P},\hat{\alpha})$ and applying $D_P\Psi(\hat{P},\hat{\alpha}) = 0$ and Lemma 11 (b) give $\tilde{P}_k = \Psi(\tilde{P}_{k-1},\tilde{\alpha}_k) = \hat{P} + O_p(||\tilde{\alpha}_k-\hat{\alpha}|| + ||\tilde{P}_{k-1}-\hat{P}||^2) = \hat{P} + O_p(||\tilde{\alpha}_k-\hat{\alpha}||)$. The required result follows from induction. □

7.6 Proof of Proposition 6

We prove the result for only the NR and OPG. We suppress the superscript MPL from $\tilde{\alpha}_j^{MPL}$ and $\tilde{P}_j^{MPL}$, and we suppress $\hat{\theta}_f$ from $\overline{\psi}_\alpha^2(P, \alpha, \hat{\theta}_f)$ and $Q_N(P, \alpha, \hat{\theta}_f)$ when it does not lead to confusion.

The proof is similar to the proof of Proposition 5. Since the NFXP estimator satisfies the first order condition $\overline{\psi}_\alpha^2(\hat{P}, \hat{\alpha}) = 0$, applying the generalized Taylor's theorem to $\overline{\psi}_\alpha^2(\hat{P}, \hat{\alpha}) - \overline{\psi}_\alpha^2(\tilde{P}_{k-1}, \tilde{\alpha}_{k-1})$ gives

$$\begin{aligned} 0 &= Q_N(\tilde{P}_{k-1}, \tilde{\alpha}_{k-1})(\hat{\alpha} - \tilde{\alpha}_k) + \left[D_\alpha \overline{\psi}_\alpha^2(\tilde{P}_{k-1}, \tilde{\alpha}_{k-1}) - Q_N(\tilde{P}_{k-1}, \tilde{\alpha}_{k-1}) \right] (\hat{\alpha} - \tilde{\alpha}_{k-1}) \\ &\quad + D_P \overline{\psi}_\alpha^2(\tilde{P}_{k-1}, \tilde{\alpha}_{k-1})(\hat{P} - \tilde{P}_{k-1}) + R_{N,k}, \end{aligned} \tag{26}$$

where $R_{N,k} = O_p(||\hat{\alpha} - \tilde{\alpha}_{k-1}||^2 + ||\hat{P} - \tilde{P}_{k-1}||^2)$ from Lemma 11 (c). For the second term on the right of (26), the term inside the bracket is 0 in the NR and $O_p(||\hat{P} - \tilde{P}_{k-1}|| + ||\hat{\alpha} - \tilde{\alpha}_{k-1}|| + N^{-1/2})$ in the OPG from Lemma 11 (g) (h) and the information matrix equality. For the third term on the right of (26), it follows from the generalized Taylor's theorem, $D_P \overline{\psi}_\alpha^2(\hat{P}, \hat{\alpha}) = 0$ (see (21)), and Lemma 11 (c) that

$$D_P \overline{\psi}_\alpha^2(\tilde{P}_{k-1}, \tilde{\alpha}_{k-1}) = O_p(||\tilde{P}_{k-1} - \hat{P}||) + O_p(||\tilde{\alpha}_{k-1} - \hat{\alpha}||).$$

Therefore,

$$\begin{aligned} & Q_N(\tilde{P}_{k-1}, \tilde{\alpha}_{k-1})(\hat{\alpha} - \tilde{\alpha}_k) \\ = \; & O_p(||\hat{\alpha} - \tilde{\alpha}_{k-1}||^2 + ||\hat{P} - \tilde{P}_{k-1}||^2) + \begin{cases} 0, & \text{for NR}, \\ O_p(N^{-1/2}||\hat{\alpha} - \tilde{\alpha}_{k-1}||), & \text{for OPG}, \end{cases} \end{aligned}$$

and the stated bound of $\hat{\alpha} - \tilde{\alpha}_k$ follows.

For the bound of $\tilde{P}_k - \hat{P}$, expanding $\Psi(\tilde{P}_{k-1}, \tilde{\alpha}_k)$ around $(\hat{P}, \hat{\alpha})$ and applying $D_P \Psi(\hat{P}, \hat{\alpha}) = 0$ and Lemma 11 give $\tilde{P}_k = \Psi(\tilde{P}_{k-1}, \tilde{\alpha}_k) = \hat{P} + O_p(||\tilde{\alpha}_k - \hat{\alpha}|| + ||\tilde{P}_{k-1} - \hat{P}||^2) = \hat{P} + O_p(||\tilde{\alpha}_k - \hat{\alpha}||)$. The required result follows from induction. □

7.7 Proof of Lemma 7

First, the results of Lemmas 8 and 9 of Andrews (2001) (A01 hereafter) hold in our case, because we can replace Lemmas 5 and 7 of A01 in the proof of Lemmas 8 and 9 of A01 with 12 (a) and 14 and the proof carries through. Then the stated result follows from the proof of Theorem 1 of A01, because only Lemmas 7, 8, and 9 of A01 are used in the proof. □

7.8 Proof of Lemma 8

We drop the superscript PL and MPL from $\tilde{\alpha}_k$ and $\tilde{P}_k$. We show that, if $\tilde{\alpha}_0 = \alpha^0$ and $\tilde{P}_0 = P^0$, then for $k = 0, 1, \ldots$

$$\begin{aligned}
\sup_{\theta_0} \Pr_{\theta_0} \left(||\tilde{\alpha}_k - \hat{\alpha}|| > \mu_{N,k} \right) &= o(N^{-a}), \quad \sup_{\theta_0} \Pr_{\theta_0} \left(||\tilde{P}_k - \hat{P}|| > \mu_{N,k} \right) = o(N^{-a}), && (27)\\
\sup_{\theta_0} \Pr_{\theta_0} \left(||T_{N,k}(\theta_r^0) - T_N(\theta_r^0)|| > \mu_{N,k} \right) &= o(N^{-a}), && (28)\\
\sup_{\theta_0} \Pr_{\theta_0} \left(||\mathcal{W}_{N,k}(\theta_0) - \mathcal{W}_N(\theta_0)|| > \mu_{N,k} \right) &= o(N^{-a}), && (29)
\end{aligned}$$

Then the stated result follows from Lemma 2 of A01, because the condition on $\hat{\alpha}$ and $\hat{P} = P_{\hat{\theta}}$ (corresponding to $\hat{\theta}_N$ in A01) is satisfied by Lemmas 12 and 13.

First, we prove the result for the k-step PL estimator, so $\mu_{N,k} = N^{-(k+1)/2} \ln^{k+1} N$. We use induction. For $k = 0$, (27) holds from Lemmas 12 and 13. Suppose (27) holds for $k = j - 1 \geq 0$. Then, from (25) in the proof of Proposition 5,

$$\begin{aligned}
\hat{\alpha} - \tilde{\alpha}_j &= Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})^{-1} \left[D_\alpha \overline{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) - Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) \right] (\hat{\alpha} - \tilde{\alpha}_{j-1})\\
&\quad + Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})^{-1} D_P \overline{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})(\hat{P} - \tilde{P}_{j-1}) + Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})^{-1} R_{N,j}, \quad (30)
\end{aligned}$$

where $||R_{N,j}|| \leq \sup_{(P,\alpha,\theta_f)} ||D^2 \overline{\psi}_\alpha(P, \alpha, \theta_f)|| (||\hat{\alpha} - \tilde{\alpha}_{j-1}||^2 + ||\hat{P} - \tilde{P}_{j-1}||^2)$. It follows from the generalized Taylor's theorem that

$$D_P \overline{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) = D_P \overline{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}, \hat{\theta}_f) = D_P \overline{\psi}_\alpha(P^0, \alpha^0, \theta_f^0) + R_{N,j}^b,$$

where $||R_{N,j}^b|| \leq \sup_{(P,\alpha,\theta_f)} ||D^2 \overline{\psi}_\alpha(P, \alpha, \theta_f)|| (||\tilde{\alpha}_{j-1} - \alpha^0|| + ||\tilde{P}_{j-1} - P^0|| + ||\hat{\theta}_f - \theta_f^0||)$. From Lemmas 3(b), 3(c) and 4 of A01 we have, for all $\varepsilon > 0$ and some $K < \infty$,

$$\begin{aligned}
&\sup_{\theta_0} \Pr_{\theta_0}(\sup_{(P,\alpha,\theta_f)} ||D^2 \overline{\psi}_\alpha(P, \alpha, \theta_f)|| > K) = o(N^{-a}),\\
&\sup_{\theta_0} \Pr_{\theta_0}(||D_P \overline{\psi}_\alpha(P^0, \alpha^0, \theta_f^0)|| > \varepsilon N^{-1/2} \ln N) = o(N^{-a}), \qquad (31)\\
&\sup_{\theta_0} \Pr_{\theta_0}(||Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})^{-1}|| > K) = o(N^{-a}).
\end{aligned}$$

In case of NR, because $D_\alpha \overline{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) - Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) = 0$, it follows that

$$\begin{aligned}
||\hat{\alpha} - \tilde{\alpha}_j|| \leq\ & \xi_{N,j}(||\hat{\alpha} - \tilde{\alpha}_{j-1}||^2 + ||\hat{P} - \tilde{P}_{j-1}||^2)\\
&+ \xi_{N,j} ||D_P \overline{\psi}_\alpha(P^0, \alpha^0, \theta_f^0)|| ||\hat{P} - \tilde{P}_{j-1}||\\
&+ \xi_{N,j} \left(||\tilde{\alpha}_{j-1} - \hat{\alpha}|| + ||\tilde{P}_{j-1} - \hat{P}|| \right) ||\hat{P} - \tilde{P}_{j-1}||\\
&+ \xi_{N,j} \left(||\hat{\alpha} - \alpha^0|| + ||\hat{P} - P^0|| + ||\hat{\theta}_f - \theta_f^0|| \right) ||\hat{P} - \tilde{P}_{j-1}|| \qquad (32)
\end{aligned}$$

where $\sup_{\theta_0}\Pr_{\theta_0}(||\xi_{N,j}|| > K) = o(N^{-a})$ for some $K < \infty$. Thus (27) holds for $k = j$ from (31) and Lemmas 12 and 13. In case of the default NR, line-search NR, and OPG, the right hand side of (32) has an additional term $\xi_{N,j}||D_\alpha\overline{\psi}_\alpha(\tilde{P}_{j-1},\tilde{\alpha}_{j-1}) - Q_N(\tilde{P}_{j-1},\tilde{\alpha}_{j-1})||||\hat{\alpha} - \tilde{\alpha}_{j-1}||$. By repeating the argument of the proof of Lemma 1 of A01, we can show

$$\sup_{\theta_0}\Pr_{\theta_0}(\sup_{(P,\alpha,\theta_f)}||D_\alpha\overline{\psi}_\alpha(\tilde{P}_{j-1},\tilde{\alpha}_{j-1}) - Q_N(\tilde{P}_{j-1},\tilde{\alpha}_{j-1})|| > N^{-1/2}\ln N) = o(N^{-a}),$$

and (27) follows, and the proof for general $k \geq 1$ follows from induction.

We proceed to prove (28) and (29). Let Σ_r denote $(\Sigma_N)_{rr}$. Let $\Sigma_{k,r}$ denote Σ_r with $D_N(\hat{\theta})$, $D_N^O(\hat{\theta})$, and $V_N(\hat{\theta})$ in its definition of (7) replaced with $D_N^{PL}(\tilde{P}_k,\tilde{\theta}_k)$, $D_N^{O,PL}(\tilde{P}_k,\tilde{\theta}_k)$ and $V_N^{PL}(\tilde{P}_k,\tilde{\theta}_k)$, where $\tilde{\theta}_k = ((\tilde{\alpha}_k)',\hat{\theta}_f')$. In view of the arguments in p. 26 of A01, (28) holds if there exists $K < \infty$ and a $\delta > 0$ such that

$$\sup_{\theta_0}\Pr_{\theta_0}(|\Sigma_r - \Sigma_{k,r}| > \mu_{N,k}) = o(N^{-a}), \tag{33}$$

$$\sup_{\theta_0}\Pr_{\theta_0}(\Sigma_{k,r} < \delta) = o(N^{-a}), \quad \sup_{\theta_0}\Pr_{\theta_0}(\Sigma_r < \delta) = o(N^{-a}). \tag{34}$$

Let $\bar{\theta}$ denote an estimator that satisfies: for all $\varepsilon > 0$, $\sup_{\theta_0}\Pr_{\theta_0}(||\bar{\theta} - \theta^0|| > \varepsilon) = o(N^{-a})$. Then, proceeding in the same way as the proof of Lemma 4 of A01, we obtain the following; for all $\varepsilon > 0$ and some $K < \infty$, $\sup_{\theta_0}\Pr_{\theta_0}(||V_N(\bar{\theta}) - V(\theta^0)|| > \varepsilon) = o(N^{-a})$, $\sup_{\theta_0}\Pr_{\theta_0}(||D_N(\bar{\theta}) - D(\theta^0)|| > \varepsilon) = o(N^{-a})$, and $\sup_{\theta_0}\Pr_{\theta_0}(||D_N^O(\bar{\theta}) - D(\theta^0)|| > \varepsilon) = o(N^{-a})$. Therefore, (34) holds. (33) holds if

$$\sup_{\theta_0}\Pr_{\theta_0}(||V_N^{PL}(\tilde{P}_k,\tilde{\theta}_k) - V_N(\hat{\theta})|| > \mu_{N,k}) = o(N^{-a}), \quad \sup_{\theta_0}\Pr_{\theta_0}(||D_N^{O,PL}(\tilde{P}_k,\tilde{\theta}_k) - D_N^O(\hat{\theta})|| > \mu_{N,k}) = o(N^{-a}).$$

For the first one, applying the generalized Taylor's theorem, Lemma 3(b) of A01, and (27) gives $\sup_{\theta_0}\Pr_{\theta_0}(||V_N^{PL}(\tilde{P}_k,\tilde{\theta}_k) - V_N^{PL}(\hat{P},\hat{\theta})|| > \mu_{N,k}) = o(N^{-a})$. Since $V_N^{PL}(\hat{P},\hat{\theta}) - V_N(\hat{\theta}) = 0$ from (17), the first result follows. The second result is proven in an analogous manner, and we complete the proof of (28). The corresponding result does not hold for $D_N^{PL}(\tilde{P}_k,\tilde{\theta}_k) - D_N(\hat{\theta})$, however, because $D_{\theta\theta}P_\theta \neq D_{\theta\theta}\Psi(P_\theta,\theta)$ in general from (18). (29) follows from (27) and the proof of (28) and because Lemma 9(a) of A01 holds in our case (see the proof of Lemma 7); see pp. 26-27 of A01.

The stated results for the k-step MPL estimator are proven in an analogous manner. Suppose (27) holds for $k = j - 1 \geq 0$ with $\mu_{N,k} = N^{-(k+1)/2}\ln^{k+1} N$ for the OPG and $\mu_{N,k} = N^{-2^{k-1}}\ln^{2^k}(N)$ in the other cases. From (26) in the proof of Proposition 6, the equality (30) holds for $\hat{\alpha} - \tilde{\alpha}_j$ with $\overline{\psi}_\alpha$ replaced by $\overline{\psi}^2_\alpha$ and $||R_{N,j}|| \leq \sup_{(P,\alpha,\theta_f)} ||D^2\overline{\psi}^2_\alpha(P,\alpha,\theta_f)||(||\hat{\alpha} - \tilde{\alpha}_{j-1}||^2 + ||\hat{P} - \tilde{P}_{j-1}||^2)$. It follows from $D_P\overline{\psi}^2_\alpha(\hat{P},\hat{\alpha}) = 0$ and the generalized Taylor's theorem that $||D_P\overline{\psi}^2_\alpha(\tilde{P}_{j-1},\tilde{\alpha}_{j-1})|| \leq \sup_{(P,\alpha,\theta_f)} ||D^2\overline{\psi}^2_\alpha(P,\alpha,\theta_f)||(||\tilde{P}_{j-1} - \hat{P}|| + ||\tilde{\alpha}_{j-1} - \hat{\alpha}||)$. Since the bound (31) holds if $\overline{\psi}_\alpha$ is replaced with $\overline{\psi}^2_\alpha$, we obtain

$$||\hat{\alpha}-\tilde{\alpha}_j|| \leq \xi_{N,j}(||\hat{\alpha}-\tilde{\alpha}_{j-1}||^2+||\hat{P}-\tilde{P}_{j-1}||^2)+\xi_{N,j}\left[D_\alpha\overline{\psi}^2_\alpha(\tilde{P}_{j-1},\tilde{\alpha}_{j-1}) - Q_N(\tilde{P}_{j-1},\tilde{\alpha}_{j-1})\right]||\hat{\alpha}-\tilde{\alpha}_{j-1}||,$$

with the same definition of $\xi_{N,j}$. For the NR, the second term on the right is 0 and hence (27) holds for $k = j$ with $\mu_{N,k} = N^{-2^{k-1}} \ln^{2^k}(N)$. Repeating the argument of the proof of Lemma 1 of A01 gives the required result for the default NR and line-search NR. In case of the OPG, repeating the argument of the proof of Lemma 1 of A01 gives

$$\sup_{\theta_0} \Pr_{\theta_0}(\sup_{(P,\alpha,\theta_f)} ||D_\alpha\overline{\psi}^2_\alpha(\tilde{P}_{j-1},\tilde{\alpha}_{j-1}) - Q_N(\tilde{P}_{j-1},\tilde{\alpha}_{j-1})|| > N^{-1/2}\ln N) = o(N^{-a}),$$

and (27) holds for $k = j$ with $\mu_{N,k} = N^{-(k+1)/2}\ln^{k+1} N$ for the OPG, and the proof for general $k \geq 1$ follows from induction.

(28) and (29) are proven using the same argument as the one for the k-step PL estimator. The only difference is it also holds that $\sup_{\theta_0} \Pr_{\theta_0}(||D_N^{MPL}(\tilde{P}_k,\tilde{\theta}_k) - D_N(\hat{\theta})|| > \mu_{N,k}) = o(N^{-a})$ by the virtue of (21). □

7.9 Proof of Lemma 9 and 10

This lemma is proven from applying the argument of p. 27 of A01. □

8 Appendix B: Auxiliary results

Lemma 11 collects the bounds that are used in the proof of Propositions 2-6. Lemma 12 is our version (i.e. for $\hat{\alpha}$ and $\hat{\theta}_f$) of Lemma 5 of A01. Lemma 14 is our version (i.e. for $\hat{\alpha}$ and $\hat{\theta}_f$) of Lemma 7 of A01.

Lemma 11 *Suppose Assumptions 1-5 hold. Let $\bar{P}$ and $\bar{\theta}$ denote estimators that converge to P^0*

and θ^0 *in probability, and let* $\psi_i(P,\theta) = \ln \Psi(P,\theta)(a_i|x_i)$ *and* $\psi_i^2(P,\theta) = \ln \Psi^2(P,\theta)(a_i|x_i)$.

$$\begin{array}{ll}
(a) & (\partial/\partial\theta)\Psi(\bar{P},\bar{\theta})(a_i|x_i) = O_p(1), \\
(b) & N^{-1}\sum_{i=1}^N \sup_{(P,\theta)\in\Theta_0} ||D^s\psi_i(P,\theta)|| = O_p(1) \quad \textit{for } s = 1,2,3, \\
(c) & N^{-1}\sum_{i=1}^N \sup_{(P,\theta)\in\Theta_0} ||D^s\psi_i^2(P,\theta)|| = O_p(1) \quad \textit{for } s = 1,2,3, \\
(d) & N^{-1}\sum_{i=1}^N D_{P\alpha}\psi_i(P^0,\theta^0) = O_p(N^{-1/2}), \\
(e) & N^{-1}\sum_{i=1}^N D_{\theta\theta}\psi_i(\bar{P},\bar{\theta}) = ED_{\theta\theta}\psi_i(P^0,\theta^0) + O_p(||\bar{P}-P^0|| + ||\bar{\theta}-\theta^0|| + N^{-1/2}), \\
(f) & \begin{array}{l} N^{-1}\sum_{i=1}^N D_\theta\psi_i(\bar{P},\bar{\theta})D_\theta\psi_i(\bar{P},\bar{\theta}) \\ = ED_\theta\psi_i(P^0,\theta^0)D_\theta\psi_i(P^0,\theta^0) + O_p(||\bar{P}-P^0|| + ||\bar{\theta}-\theta^0|| + N^{-1/2}), \end{array} \\
(g) & N^{-1}\sum_{i=1}^N D_{\theta\theta}\psi_i^2(\bar{P},\bar{\theta}) = ED_{\theta\theta}\psi_i^2(P^0,\theta^0) + O_p(||\bar{P}-P^0|| + ||\bar{\theta}-\theta^0|| + N^{-1/2}), \\
(h) & \begin{array}{l} N^{-1}\sum_{i=1}^N D_\theta\psi_i^2(\bar{P},\bar{\theta})D_\theta\psi_i^2(\bar{P},\bar{\theta}) \\ = ED_\theta\psi_i^2(P^0,\theta^0)D_\theta\psi_i^2(P^0,\theta^0) + O_p(||\bar{P}-P^0|| + ||\bar{\theta}-\theta^0|| + N^{-1/2}), \end{array}
\end{array}$$

If Assumptions 1-8 hold, then (b) and (c) hold for $(P,\theta) \in \Theta_1$.

Proof Parts (a) and (b) follow from Assumption 4(7). Part (c) follows from Assumption 5(b), and part (d) follows Assumption 4(7) and $ED_{P\alpha}\ln\Psi(P^0,\theta^0) = 0$ (zero operator) from the information matrix equality and Proposition 1. Parts (e)-(h) follow from (b), (c), and the law of large numbers. □

Lemma 12 *Suppose Assumptions 1-8 hold. Then, for all* $\varepsilon > 0$,

$$\sup_{\theta^0\in\Theta^1} \Pr_{\theta^0}\left(N^{1/2}||\hat{\theta}_f - \theta_f^0|| + N^{1/2}||\hat{\alpha}-\alpha^0|| > \varepsilon\ln N\right) = o(N^{-a}),$$

Proof First, we can simply apply Lemma 5 of Andrews (2001) to $\hat{\theta}_f$ to show that, for all $\varepsilon > 0$,

$$\sup_{\theta_f^0\in\Theta_f^1} \Pr_{\theta_f^0}\left(N^{1/2}||\hat{\theta}_f - \theta_f^0|| > \varepsilon\ln N\right) = o(N^{-a}).$$

Define $Q_N(\alpha,\theta_f) = -N^{-1}\sum_{i=1}^N \ln P_{(\alpha,\theta_f)}(a_i|x_i)$ and $Q(\alpha,\theta_f) = -E\ln P_{(\alpha,\theta_f)}(a_i|x_i)$, so that $\hat{\alpha} = \arg\min Q_N(\alpha,\hat{\theta}_f)$. From the compactness of Θ, $\sup_{\theta^0\in\Theta^1} E_{\theta^0}\sup_{\theta\in\Theta} ||(\partial/\partial\theta_f)\ln P_\theta(a_i|x_i)||^{q_0} < \infty$, $\sup_{\theta^0\in\Theta^1} E_{\theta^0}|\ln P_\theta(a_i|x_i)|^{q_0} < \infty$ for all $\theta\in\Theta$, applying the proof of Lemma 5 of A01 gives, for all $\varepsilon > 0$,

$$\sup_{\theta^0\in\Theta^1} \Pr_{\theta^0}\left(\sup_{\alpha,\theta_f}|Q_N(\alpha,\theta_f) - Q(\alpha,\theta_f)| > \varepsilon\right) = o(N^{-a}).$$

Since $Q(\alpha, \theta_f)$ is minimized at (α^0, θ_f^0) and from some additional conditions, there exists $\delta > 0$ such that $||\alpha - \alpha^0|| > \varepsilon$ implies $Q(\alpha, \theta_f^0) - Q(\alpha^0, \theta_f^0) \geq \delta$. Therefore,

$$\begin{aligned}\sup_{\theta^0 \in \Theta^1} \Pr_{\theta^0}(||\hat{\alpha} - \alpha|| &> \varepsilon) \leq \sup_{\theta^0 \in \Theta^1} \Pr_{\theta^0}\left(Q(\hat{\alpha}, \theta_f^0) - Q(\alpha^0, \theta_f^0) \geq \delta\right) \\ &\leq \sup_{\theta^0 \in \Theta^1} \Pr_{\theta^0}\left(Q(\hat{\alpha}, \hat{\theta}_f) - Q(\alpha^0, \hat{\theta}_f) + 2 \sup_{\alpha \in \Theta_\alpha} |Q(\alpha, \theta_f^0) - Q(\alpha, \hat{\theta}_f)| \geq \delta\right).\end{aligned}$$

Since $\sup_{\theta_f^0 \in \Theta_f^1} \Pr_{\theta_f^0}(||\hat{\theta}_f - \theta_f^0|| > \varepsilon) = o(N^{-a})$ for any $\varepsilon > 0$ from the above and $Q(\alpha, \theta_f)$ is uniformly continuous, the above is no larger than

$$\begin{aligned}&\sup_{\theta^0 \in \Theta^1} \Pr_{\theta^0}\left(Q(\hat{\alpha}, \hat{\theta}_f) - Q_N(\hat{\alpha}, \hat{\theta}_f) + Q_N(\hat{\alpha}, \hat{\theta}_f) - Q(\alpha^0, \hat{\theta}_f) \geq \delta/2\right) + o(N^{-a}) \\ \leq\ &\sup_{\theta^0 \in \Theta^1} \Pr_{\theta^0}\left(Q(\hat{\alpha}, \hat{\theta}_f) - Q_N(\hat{\alpha}, \hat{\theta}_f) + Q_N(\alpha^0, \hat{\theta}_f) - Q(\alpha^0, \hat{\theta}_f) \geq \delta/2\right) + o(N^{-a}) \\ \leq\ &\sup_{\theta^0 \in \Theta^1} \Pr_{\theta^0}\left(2 \sup_{\alpha, \theta_f} |Q(\alpha, \theta_f) - Q_N(\alpha, \theta_f)| \geq \delta/2\right) + o(N^{-a}) = o(N^{-a}),\end{aligned}$$

giving the required result. □

Lemma 13 *Suppose Assumptions 1-8 hold. Suppose* $\sup_\theta ||(\partial/\partial\theta) P_\theta(a_i|x_i)|| < \infty$ *with probability one. Then*

$$\sup_{\theta^0 \in \Theta^1} \Pr_{\theta^0}\left(N^{1/2} ||P_{\hat{\theta}} - P^0|| > \varepsilon \ln N\right) = o(N^{-a}).$$

Proof Simply expand $P_{\hat{\theta}} = P_{(\hat{\alpha}, \hat{\theta}_f)}$ around (α^0, θ_f^0) to obtain

$$|P_{(\hat{\alpha}, \hat{\theta}_f)} - P^0| \leq ||(\partial/\partial\theta) P_{\bar{\theta}}||(||\hat{\alpha} - \alpha^0|| + ||\hat{\theta}_f - \hat{\theta}_f^0||), \quad \bar{\theta} \in [\hat{\theta}, \theta^0]$$

and the required result follows. □

Lemma 14 *Suppose Assumptions 1-8 hold. Define* $S_N(\theta) = N^{-1} \sum_{i=1}^N h(W_i, \theta)$ *and* $\hat{\theta} = (\hat{\alpha}', \hat{\theta}_f')'$. *Let* $\Delta_N(\theta^0)$ *denote* $N^{1/2}(\hat{\theta} - \theta^0)$, $T_N(\theta_r^0)$, *or* $H_N(\hat{\theta}, \beta^0)$, *where* $\theta^0 = (\beta^{0\prime}, \delta^{0\prime})$. *Let* L *denote the dimension of* $\Delta_N(\theta^0)$. *For each definition of* $\Delta_N(\theta^0)$, *there is an infinitely differentiable function* $G(\cdot)$ *that does not depend on* θ^0 *that satisfies* $G(E_{\theta^0} S_N(\theta^0)) = 0$ *for all* N *large and all* $\theta^0 \in \Theta_1$ *and*

$$\sup_{\theta^0 \in \Theta_1} \sup_{B \in \mathcal{B}_L} \left| \Pr_{\theta^0}(\Delta_N(\theta^0) \in B) - \Pr_{\theta^0}(N^{1/2} G(S_N(\theta^0)) \in B) \right| = o(N^{-a}).$$

Proof Suppose $\Delta_N(\theta^0) = N^{1/2}(\hat{\theta} - \theta^0)$. By Lemma 12, $\hat{\theta}$ is in the interior of Θ with probability $1 - o(N^{-a})$ and $\Pr_{\theta^0}((\partial/\partial\alpha')N^{-1}\sum_{i=1}^{N} \ln P_{(\hat{\alpha},\hat{\theta}_f)}(a_i|x_i) = 0)$ and $\Pr_{\theta^0}((\partial/\partial\theta_f')N^{-1}\sum_{i=1}^{N} \ln f_{\hat{\theta}_f}(x_i'|a_i, x_i) = 0)$ are $1 - o(N^{-a})$ uniformly in $\theta^0 \in \Theta_1$. Let

$$g(\theta) = \begin{bmatrix} (\partial/\partial\alpha')N^{-1}\sum_{i=1}^{N} \ln P_{(\alpha,\theta_f)}(a_i|x_i) \\ (\partial/\partial\theta_f')N^{-1}\sum_{i=1}^{N} \ln f_{\theta_f}(x_i'|a_i, x_i) \end{bmatrix}.$$

Expanding these first order conditions around θ^0 gives

$$0 = g(\hat{\theta}) = g(\theta^0) + \sum_{j=1}^{d-1} \frac{1}{j!} D^j g(\theta^0)(\hat{\theta} - \theta^0, \ldots, \hat{\theta} - \theta^0) + \zeta_N(\theta^0),$$

where

$$\zeta_N(\theta^0) = \frac{1}{(d-1)!}\left(D^{d-1}g(\theta^+) - D^{d-1}g(\theta^0)\right)(\hat{\theta} - \theta^0, \ldots, \hat{\theta} - \theta^0),$$

where θ^+ lies between θ_N and θ^0, and $D^j g(\theta^0)(\hat{\theta} - \theta^0, \ldots, \hat{\theta} - \theta^0)$ denotes $D^j g(\theta^0)$ as a j-linear map, whose coefficients are partial derivatives of $g(\theta^0)$ of order j, applied to the j-tuple $(\hat{\theta} - \theta^0, \ldots, \hat{\theta} - \theta^0)$. Let $R_N(\theta^0)$ denote the column vector whose elements are the unique components of $g(\theta^0), D^1 g(\theta^0), \ldots, D^{d-1} g(\theta^0)$. Each element of $R_N(\theta^0)$ is an element of $S_N(\theta^0)$. Let $e_N(\theta^0) = (\zeta_n(\theta^0)', 0, \ldots, 0)'$ be conformable to $R_N(\theta^0)$. The first equation can be written as $\nu(R_N(\theta^0) + e_N(\theta^0), \hat{\theta} - \theta^0) = 0$, where $\nu(\cdot,\cdot)$ is infinitely differentiable, $\nu(E_{\theta^0} R_N(\theta^0), 0) = 0$, and

$$\begin{aligned} \frac{\partial}{\partial x}\nu(E_{\theta^0} R_N(\theta^0), x)|_{x=0} &= E\begin{bmatrix} D^1(\partial/\partial\alpha') \ln P_{\theta^0}(a_i|x_i) \\ D^1(\partial/\partial\theta_f') \ln f_{\theta_f^0}(x_i'|a_i, x_i) \end{bmatrix} \\ &= E\begin{bmatrix} (\partial/\partial\alpha\partial\alpha') \ln P_{\theta^0}(a_i|x_i) & (\partial/\partial\theta_f\partial\alpha') \ln P_{\theta^0}(a_i|x_i) \\ 0 & (\partial/\partial\theta_f\partial\theta_f') \ln f_{\theta_f^0}(x_i'|a_i, x_i) \end{bmatrix} \end{aligned}$$

is negative definite. Hence, the implicit function theorem can be applied to $\nu(\cdot,\cdot)$ at the point $(E_{\theta^0} R_N(\theta^0), 0)$ to obtain

$$\inf_{\theta^0\in\Theta_1} \Pr_{\theta^0}\left(\hat{\theta} - \theta^0 = \Lambda(R_N(\theta^0) + e_N(\theta^0))\right) = 1 - o(N^{-a}),$$

where Λ is a function that does not depend on N or θ^0, is infinitely differentiable in a neighborhood of $E_{\theta^0} R_N(\theta^0)$ and satisfies $\Lambda(E R_N(\theta^0)) = 0$.

Therefore, we can apply Lemma 6 of A01 to obtain the required result. □

References

[1] Aguirregabiria, V. and P. Mira (2002). "Swapping the nested fixed point algorithm: a class of estimators for discrete Markov decision models." *Econometrica* 70(4): 1519-1543.

[2] Andrews, D. W. K. (2001). Higher-order improvements of the parametric bootstrap for Markov processes. Cowles Foundation Discussion Paper No. 1334, Yale University. Forthcoming in "Identification and Inference for Econometric Models," Cambridge University Press.

[3] Andrews, D. W. K. (2002a). "Higher-order improvements of a computationally attractive k-step bootstrap for extremum estimators." *Econometrica* 70(1): 119-162.

[4] Andrews, D. W. K. (2002b). "Equivalence of the higher order asymptotic efficiency of k-step and extremum statistics." *Econometric Theory* 18: 1040-1085.

[5] Atkinson, E. K. (1997) *The Numerical Solution of Integral Equation of the Second Kind.* Cambridge University Press.

[6] Davidson, R. and J. G. MacKinnon (1999a). "Bootstrap testing in nonlinear models." *International Economic Review* 40(2): 487-508.

[7] Davidson, R. and J. G. MacKinnon (2000). "Bootstrap tests: how many bootstraps?" *Econometric Reviews*, 19: 55-68.

[8] Efron, B. and R. J. Tibshirani (1993) *An Introduction ot the Bootstrap.* New York, Chapman & Hall/CRC.

[9] Hall, P. (1992) *The Bootstrap and Edgeworth Expansion.* New York, Springer-Verlag.

[10] Hotz, J. and R. A. Miller (1993). "Conditional choice probabilities and the estimation of dynamic models." *Review of Economic Studies* 60: 497-529.

[11] Ichimura, H. and S. Lee (2004). "Calculation of the asymptotic distribution of semiparametric M-estimators." Mimeographed, University College of London.

[12] Newey, W. K., and McFadden, D. (1994) "Large Sample Estimation and Hypothesis Testing." in *Handbook of Econometrics, Volume 4*, eds. R. F. Engle and D. L. McFadden, Elsevier.

[13] Robinson, P. M. (1988) "The Stochastic difference between econometric statistics." *Econometrica* 56(3): 531-548.

[14] Puterman, M. (1994). *Markov Decision Processes.* New York, Wiley.

[15] Rust, J. (1987). "Optimal replacement of GMC bus engines: an empirical model of Harold Zurcher." *Econometrica* 55(5): 999-1033.

[16] Rust, J. (1988). "Maximum likelihood estimation of discrete control processes." *SIAM Journal of Control and Optimization* 26(5): 1006-1024.

[17] Rust, J. (1994a). "Estimation of dynamic structural models, problems and prospects: discrete decision processes," in *Advances in Econometrics, Proceedings of the Sixth World Congress of the Econometric Society* ed. by C. Sims. Cambridge, Cambridge University Press: 119-170.

[18] Rust, J. (1994b). "Structural estimation of Markov decision processes." in *Handbook of Econometrics, Vol. IV.* D. McFadden and R. F. Engle. New York, North-Holland: 3082-3139.

[19] Rust, J. (1996). "Numerical Dynamic Programming in Economics." in *Handbook of Computational Economics* H. Amman, D. Kendrick and J. Rust. Elsevier, North Holland.

[20] Zeidle, E. (1986). *Nonlinear Functional Analysis and its Applications I: Fixed-Point Theorems.* New York, Sprinver-Verlag.

Table 1: Monte Carlo Distribution of MLE

	θ_0	θ_1	θ_2
Mean Absolute Error	0.2262	0.1250	0.0382
Median Absolute Error	0.1852	0.0689	0.0282
Standard Error	0.2865	0.2245	0.0530

Notes: "Mean Absolute Error" is the mean of |ML estimate - parameter value|. "Median Absolute Error" is the median of |ML estimate - parameter value|. "Standard Error" is the square root of the mean squared deviation of ML estimate from parameter value.

Table 2: Monte Carlo Distribution of PL and MPL Estimators relative to MLE

				Steps		
	Statistics	k=1	k=2	k=3	k=4	k=5
PL Estimator						
θ_0	Mean AE	2.4475	0.7073	0.6389	0.6358	0.6354
	Median AE	1.1164	0.0276	0.0051	0.0045	0.0044
	Top 5% AE	6.5374	0.3665	0.0249	0.0148	0.0136
θ_1	Mean AE	6.0626	4.5582	4.4541	4.4468	4.4462
	Median AE	1.1366	0.0412	0.0025	0.0016	0.0015
	Top 5% AE	7.4709	0.4293	0.0440	0.0143	0.0130
θ_2	Mean AE	5.7233	0.6233	0.3759	0.3540	0.3514
	Median AE	2.3286	0.0674	0.0047	0.0028	0.0027
	Top 5% AE	22.6752	1.4575	0.1272	0.0183	0.0136
MPL Estimator						
θ_0	Mean AE	0.0251	0.0050	0.0048	0.0041	0.0053
	Median AE	0.0072	0.0038	0.0036	0.0027	0.0042
	Top 5% AE	0.0956	0.0136	0.0132	0.0124	0.0135
θ_1	Mean AE	0.0304	0.0030	0.0029	0.0022	0.0029
	Median AE	0.0056	0.0015	0.0012	0.0009	0.0013
	Top 5% AE	0.1444	0.0113	0.0133	0.0089	0.0105
θ_2	Mean AE	0.1172	0.0042	0.0037	0.0032	0.0041
	Median AE	0.0106	0.0027	0.0022	0.0019	0.0024
	Top 5% AE	0.5499	0.0132	0.0128	0.0109	0.0141
k-PL Estimator						
θ_0	Mean AE	16.1415	24.8390	29.0127	32.4305	36.5971
	Median AE	8.0803	2.0326	0.4774	0.1314	0.0378
	Top 5% AE	63.4902	145.5991	138.0032	109.8606	120.8148
θ_1	Mean AE	32.4990	25.1312	24.4792	23.5439	23.6059
	Median AE	7.6545	2.4350	0.6622	0.1862	0.0666
	Top 5% AE	197.8664	136.2302	172.5427	156.0622	152.8299
θ_2	Mean AE	31.9763	16.1400	13.9012	13.8815	14.6365
	Median AE	15.3367	3.2848	0.8816	0.2379	0.0783
	Top 5% AE	112.6538	78.4998	88.1842	95.4993	92.0382
k-MPL Estimator						
θ_0	Mean AE	16.1415	25.8171	30.4228	33.3208	36.8865
	Median AE	8.0803	2.0068	0.4608	0.1411	0.0377
	Top 5% AE	63.4902	164.6399	165.5518	123.9891	132.2554
θ_1	Mean AE	32.4990	25.2531	24.4088	23.2165	23.3835
	Median AE	7.6545	2.4782	0.6706	0.2122	0.0671
	Top 5% AE	197.8664	136.8179	162.9697	155.9485	170.6729
θ_2	Mean AE	31.9763	16.6148	12.8602	12.4662	13.7243
	Median AE	15.3367	3.5368	0.9198	0.2437	0.0799
	Top 5% AE	112.6538	78.3703	80.6983	86.9023	92.0255

Notes: All entries are percentage differences relative to the standard error of ML estimate. For instance, "Mean AE" for "PL Estimator" is computed by averaging the absolute percentage difference between PL estimate and ML estimate and then dividing by the standard deviation of ML estimate [i.e., the mean of 100×|PL estimate - ML estimate| /(S.E. of ML esitimates)].

Table 3: Monte Carlo Distribution of t-statistics of PL and MPL relative to MLE

		Steps				
	Statistics	k=1	k=2	k=3	k=4	k=5
PL Estimator						
θ_0	Mean AE	0.0253	0.0061	0.0053	0.0053	0.0053
	Median AE	0.0118	0.0003	0.0001	0.0000	0.0000
	Top 5% AE	0.0816	0.0041	0.0003	0.0002	0.0001
θ_1	Mean AE	0.0863	0.0233	0.0193	0.0189	0.0188
	Median AE	0.0163	0.0006	0.0000	0.0000	0.0000
	Top 5% AE	0.3275	0.0257	0.0022	0.0003	0.0002
θ_2	Mean AE	0.0621	0.0086	0.0061	0.0059	0.0059
	Median AE	0.0276	0.0008	0.0001	0.0000	0.0000
	Top 5% AE	0.2341	0.0130	0.0011	0.0002	0.0001
MPL Estimator						
θ_0	Mean AE	0.0003	0.0001	0.0001	0.0000	0.0001
	Median AE	0.0001	0.0000	0.0000	0.0000	0.0000
	Top 5% AE	0.0012	0.0001	0.0001	0.0001	0.0001
θ_1	Mean AE	0.0017	0.0001	0.0000	0.0000	0.0000
	Median AE	0.0001	0.0000	0.0000	0.0000	0.0000
	Top 5% AE	0.0076	0.0002	0.0002	0.0001	0.0002
θ_2	Mean AE	0.0010	0.0000	0.0000	0.0000	0.0000
	Median AE	0.0001	0.0000	0.0000	0.0000	0.0000
	Top 5% AE	0.0048	0.0001	0.0001	0.0001	0.0001
k-PL Estimator						
θ_0	Mean AE	0.1903	0.3223	0.3736	0.4160	0.4677
	Median AE	0.0849	0.0191	0.0047	0.0012	0.0004
	Top 5% AE	0.8436	1.7149	1.5378	1.1962	1.3031
θ_1	Mean AE	3.6032	3.5311	69.4189	1372.93	1400.63
	Median AE	0.1064	0.0367	0.0112	0.0036	0.0012
	Top 5% AE	14.5354	18.6584	18.1188	19.0707	17.3141
θ_2	Mean AE	0.5008	0.2840	0.2628	0.2616	0.2711
	Median AE	0.1800	0.0404	0.0100	0.0028	0.0009
	Top 5% AE	1.8934	1.5735	1.8417	2.0554	2.0272
k-MPL Estimator						
θ_0	Mean AE	0.1900	0.3359	0.3942	0.4302	0.4715
	Median AE	0.0849	0.0187	0.0047	0.0013	0.0004
	Top 5% AE	0.8256	2.0135	2.0030	1.5392	1.3450
θ_1	Mean AE	3.9070	3.5743	74.077	465.47	465.33
	Median AE	0.1084	0.0375	0.0115	0.0038	0.0012
	Top 5% AE	19.3222	18.3311	17.5195	18.2301	15.6829
θ_2	Mean AE	0.5431	0.2936	0.2606	0.2564	0.2760
	Median AE	0.1905	0.0422	0.0098	0.0030	0.0010
	Top 5% AE	2.0742	1.5167	1.7347	1.9661	2.0205

Notes: All entries are differences in t-statistics between MLE and alternative estimators. For instance, to compute "Mean AE" for "PL Estimator," we first compute MLE and PL's t-statistics under the null of the true values of parameters as t^{ML} =(ML estimate - ture parameter)/(S.E. of ML estimate) and t^{PL} =(PL estimate - ture parameter)/(S.E. of PL estimate), and then compute the average of absolute differences between t^{ML} and t^{PL} across 1000 samples.

Table 4: Monte Carlo Distribution of t-statistics of MLE

	θ_0	θ_1	θ_2
Mean	0.1183	0.3917	0.2433
Median	0.1448	0.0212	0.3046
5 percentile	-1.6059	-0.8710	-1.2613
95 percentile	1.6851	2.7422	1.5592

Notes: All entries are computed from t-statistics under the null of the true value of parameters. For instance, to compute "5 percentile," we first compute t-statistics as t^{ML} =(ML estimate - ture parameter)/(S.E. of ML estimate) for every sample estimate, then compute the 5 percentile of t^{ML} over 1000 samples.

Table 5: Monte Carlo Distribution of t-statistics of Various Esimators at $k = 3$

		θ_0	θ_1	θ_2
Mean	MLE	0.1183	0.3917	0.2433
	PL	0.1209	0.3910	0.2447
	MPL	0.1183	0.3917	0.2433
	k-PL	-0.1941	69.7565	0.3017
	k-MPL	-0.2141	74.4163	0.2984
	k-PL (Default)	0.0809	0.1004	0.0246
	k-MPL (Default)	0.0925	0.0688	0.0086
Median	MLE	0.1447	0.0211	0.3046
	PL	0.1471	0.0211	0.3057
	MPL	0.1447	0.0211	0.3046
	k-PL	0.0611	0.0382	0.2467
	k-MPL	0.0561	0.0353	0.2542
	k-PL (Default)	0.0031	-0.0084	0.0685
	k-MPL (Default)	0.0030	-0.0084	0.0400
5 percentile	MLE	-1.6058	-0.8710	-1.2613
	PL	-1.6058	-0.8802	-1.2612
	MPL	-1.6058	-0.8710	-1.2613
	k-PL	-2.5172	-0.5838	-1.4871
	k-MPL	-2.6066	-0.5813	-1.4890
	k-PL (Default)	-1.5074	-0.5597	-1.3175
	k-MPL (Default)	-1.4839	-0.5545	-1.2821
95 percentile	MLE	1.6852	2.7422	1.5593
	PL	1.6845	2.7424	1.5264
	MPL	1.6852	2.7423	1.5593
	k-PL	1.6731	20.5810	2.7319
	k-MPL	1.6697	19.8682	2.5367
	k-PL (Default)	1.5388	1.3051	1.2256
	k-MPL (Default)	1.5329	1.1607	1.1989

Notes: All entries are computed from t-statistics under the null of the true value of parameters. For instance, to compute "5 percentile," we first compute t-statistics as t^{ML} =(ML estimate - ture parameter)/(S.E. of ML estimate) for every sample estimate, then compute the 5 percentile of t^{ML} over 1000 samples. $k = 3$ steps are taken for all estimators.

Table 6: Rejection Frequencies for Asymptotic Wald test at .10, .05, and .01 Levels for $H_0 : (\theta_1, \theta_2) = (-0.1, 0.05)$

	.10	.05	.01
$N = 200$	0.198	0.160	0.124
$N = 500$	0.190	0.144	0.092
$N = 1000$	0.176	0.132	0.078

Notes: Based on 500 simulated samples with the NFXP algorithm. N is the number of observations for each sample.

Table 7: Coverage Performance of Asymptotic 90% and 95% Confidence Intervals

		θ_0		θ_1		θ_2	
		Miss Left	Miss Right	Miss Left	Miss Right	Miss Left	Miss Right
90% CI	$N = 200$	0.040	0.028	0.138	0.000	0.038	0.012
	$N = 500$	0.050	0.040	0.108	0.000	0.042	0.030
	$N = 1000$	0.048	0.050	0.110	0.020	0.040	0.046
95% CI	$N = 200$	0.010	0.010	0.116	0.000	0.014	0.006
	$N = 500$	0.020	0.022	0.088	0.000	0.022	0.020
	$N = 1000$	0.020	0.032	0.070	0.004	0.016	0.042

Notes: Based on 500 simulated samples. N represents the number of observations for each sample. The table shows the frequencies that the confidence intervals missed the true values of $\theta_0 = 2.0$, $\theta_1 = -0.1$, and $\theta_2 = 0.05$ on the left or right side. For example, "Miss Left" for θ_0 means that the left endpoint was larger than 2.0. The true coverage is 0.9 so that the ideal values of "Miss Left" and "Miss Right" are 0.05. The variance covariance matrix is estimated by OPG with the derivative of an actual likelihood function.

Table 8: Rejection Frequencies for Bootstrap Wald test at .10, .05, and .01 Levels for $H_0: (\theta_1, \theta_2) = (-0.1, 0.05)$

		Significance Levels		
		.10	.05	.01
ML Estimator		0.118	0.056	0.010
PL Estimator	$k = 1$	0.078	0.030	0.004
	$k = 2$	0.094	0.058	0.018
	$k = 3$	0.092	0.058	0.018
MPL Estimator	$k = 1$	0.092	0.058	0.016
	$k = 2$	0.092	0.058	0.018
	$k = 3$	0.092	0.058	0.018
k-PL Estimator	$k = 1$	0.016	0.008	0.002
	$k = 2$	0.042	0.016	0.002
	$k = 3$	0.060	0.032	0.002
	$k = 4$	0.076	0.042	0.004
	$k = 5$	0.082	0.042	0.006
k-MPL Estimator	$k = 1$	0.034	0.006	0.000
	$k = 2$	0.074	0.024	0.000
	$k = 3$	0.088	0.052	0.002
	$k = 4$	0.102	0.066	0.010
	$k = 5$	0.098	0.072	0.014

Notes: Based on 500 simulated samples, each with the sample size of 500. The number of bootstrap samples is 199. The null hypothesis is $(\theta_1, \theta_2) = (-0.1, 0.05)$. For PL and MPL, we only report up to $k = 3$ because the results of $k = 4, 5$ are identical to that of $k = 3$. The Quasi-Newton method with the outer products of gradient (OPG) is used to obtain k-step estimates. The variance-covariance matrix is also computed using the OPG, where the derivatives of an actual likelihood are used for MLE while the derivatives of a pseudo likelihood defined in terms of 2 policy iterations are used for PI, MPI, k-PL, and k-MPL.

Table 9: Coverage Performance of Bootstrap 90% CIs for ML, PL and MPL Estimators

			Symmetric CI		Equal-tailed CI		Percentile CI	
			Miss Left	Miss Right	Miss Left	Miss Right	Miss Left	Miss Right
ML Estimator		θ_0	0.054	0.056	0.060	0.064	0.092	0.022
		θ_1	0.128	0.014	0.088	0.034	0.058	0.064
		θ_2	0.026	0.056	0.038	0.090	0.098	0.006
PL Estimator	θ_0	$k=1$	0.084	0.040	0.116	0.042	0.100	0.024
		$k=2$	0.088	0.040	0.112	0.042	0.104	0.022
		$k=3$	0.088	0.040	0.112	0.042	0.104	0.022
	θ_1	$k=1$	0.090	0.010	0.056	0.066	0.048	0.112
		$k=2$	0.102	0.010	0.064	0.064	0.040	0.108
		$k=3$	0.102	0.010	0.064	0.064	0.040	0.108
	θ_2	$k=1$	0.026	0.048	0.030	0.066	0.110	0.004
		$k=2$	0.026	0.046	0.028	0.072	0.132	0.002
		$k=3$	0.026	0.046	0.028	0.072	0.132	0.002
MPL Estimator	θ_0	$k=1$	0.090	0.040	0.110	0.042	0.104	0.022
		$k=2$	0.090	0.040	0.110	0.042	0.104	0.022
		$k=3$	0.090	0.040	0.110	0.042	0.104	0.022
	θ_1	$k=1$	0.102	0.020	0.064	0.064	0.040	0.108
		$k=2$	0.102	0.018	0.064	0.064	0.040	0.108
		$k=3$	0.102	0.018	0.064	0.064	0.040	0.108
	θ_2	$k=1$	0.026	0.048	0.028	0.070	0.128	0.002
		$k=2$	0.026	0.048	0.028	0.070	0.128	0.002
		$k=3$	0.026	0.048	0.028	0.070	0.128	0.002

Notes: Based on 500 simulated samples, each with the sample size of 500. The number of bootstrap samples is 199. We report only up to $k = 3$ because the results of $k = 4, 5$ are identical to that of $k = 3$. The table shows the frequencies that the confidence intervals missed the true values of $\theta_0 = 2.0$, $\theta_1 = -0.1$, and $\theta_2 = 0.05$ on the left or right side. For example, "Miss Left" for θ_0 means that the left endpoint was larger than 2.0. The true coverage is 0.9 so that the ideal values of "Miss Left" and "Miss Right" are 0.05. The variance-covariance matrix is computed using the OPG, where the derivatives of an actual likelihood are used for MLE while the derivatives of a pseudo likelihood defined in terms of 2 policy iterations are used for PI and MPI.

Table 10: Coverage Performance of Bootstrap 90% CIs for k-PL and k-MPL Estimators

			Symmetric CI		Equal-tailed CI		Percentile CI	
			Miss Left	Miss Right	Miss Left	Miss Right	Miss Left	Miss Right
k-PL Estimator	θ_0	$k = 1$	0.052	0.046	0.040	0.048	0.104	0.028
		$k = 2$	0.030	0.036	0.054	0.034	0.034	0.032
		$k = 3$	0.020	0.036	0.074	0.036	0.024	0.028
		$k = 4$	0.018	0.036	0.078	0.038	0.024	0.026
		$k = 5$	0.018	0.038	0.078	0.038	0.022	0.026
	θ_1	$k = 1$	0.034	0.000	0.010	0.076	0.108	0.000
		$k = 2$	0.076	0.000	0.020	0.046	0.056	0.000
		$k = 3$	0.100	0.000	0.044	0.046	0.048	0.000
		$k = 4$	0.106	0.000	0.056	0.036	0.046	0.000
		$k = 5$	0.108	0.000	0.064	0.036	0.044	0.000
	θ_2	$k = 1$	0.016	0.040	0.014	0.028	0.042	0.028
		$k = 2$	0.014	0.024	0.016	0.024	0.074	0.014
		$k = 3$	0.018	0.020	0.022	0.026	0.074	0.008
		$k = 4$	0.020	0.022	0.028	0.026	0.076	0.004
		$k = 5$	0.018	0.022	0.028	0.024	0.080	0.002
k-MPL Estimator	θ_0	$k = 1$	0.054	0.058	0.046	0.058	0.098	0.036
		$k = 2$	0.020	0.048	0.052	0.038	0.024	0.030
		$k = 3$	0.010	0.050	0.074	0.044	0.016	0.024
		$k = 4$	0.010	0.046	0.076	0.046	0.014	0.022
		$k = 5$	0.008	0.048	0.076	0.050	0.016	0.026
	θ_1	$k = 1$	0.068	0.000	0.018	0.092	0.116	0.000
		$k = 2$	0.088	0.000	0.042	0.062	0.080	0.000
		$k = 3$	0.118	0.000	0.068	0.064	0.074	0.000
		$k = 4$	0.120	0.000	0.078	0.052	0.072	0.000
		$k = 5$	0.120	0.000	0.086	0.058	0.072	0.000
	θ_2	$k = 1$	0.018	0.040	0.022	0.032	0.056	0.028
		$k = 2$	0.014	0.024	0.022	0.022	0.080	0.012
		$k = 3$	0.024	0.024	0.030	0.022	0.092	0.006
		$k = 4$	0.030	0.018	0.032	0.024	0.094	0.004
		$k = 5$	0.030	0.018	0.036	0.020	0.096	0.000

Notes: Based on 500 simulated samples, each with the sample size of 500. The number of bootstrap samples is 199. The table shows the frequencies that the confidence intervals missed the true values of $\theta_0 = 2.0$, $\theta_1 = -0.1$, and $\theta_2 = 0.05$ on the left or right side. For example, "Miss Left" for θ_0 means that the left endpoint was larger than 2.0. The true coverage is 0.9 so that the ideal values of "Miss Left" and "Miss Right" are 0.05. The variance-covariance matrix is computed using the OPG, where the derivatives of an actual likelihood are used for MLE while the derivatives of a pseudo likelihood defined in terms of 2 policy iterations are used for k-PL and k-MPL.

Table 11: Coverage Performance of Bootstrap 95% CIs for ML, PL and MPL Estimators

			Symmetric CI		Equal-tailed CI		Percentile CI	
			Miss Left	Miss Right	Miss Left	Miss Right	Miss Left	Miss Right
ML Estimator		θ_0	0.038	0.028	0.042	0.030	0.040	0.010
		θ_1	0.086	0.000	0.040	0.032	0.020	0.048
		θ_2	0.010	0.042	0.018	0.048	0.042	0.004
PL Estimator		$k = 1$	0.034	0.026	0.078	0.020	0.058	0.012
	θ_0	$k = 2$	0.034	0.026	0.072	0.024	0.064	0.012
		$k = 3$	0.034	0.026	0.072	0.024	0.064	0.012
		$k = 1$	0.056	0.000	0.026	0.058	0.016	0.096
	θ_1	$k = 2$	0.064	0.000	0.040	0.058	0.012	0.096
		$k = 3$	0.064	0.000	0.042	0.058	0.012	0.096
		$k = 1$	0.014	0.032	0.014	0.030	0.058	0.002
	θ_2	$k = 2$	0.014	0.028	0.014	0.038	0.072	0.002
		$k = 3$	0.014	0.028	0.014	0.038	0.072	0.002
MPL Estimator		$k = 1$	0.034	0.026	0.072	0.022	0.064	0.012
	θ_0	$k = 2$	0.034	0.026	0.072	0.022	0.064	0.012
		$k = 3$	0.034	0.026	0.072	0.022	0.064	0.012
		$k = 1$	0.064	0.000	0.038	0.058	0.014	0.096
	θ_1	$k = 2$	0.064	0.000	0.042	0.058	0.014	0.096
		$k = 3$	0.064	0.000	0.042	0.058	0.014	0.096
		$k = 1$	0.014	0.030	0.014	0.038	0.072	0.002
	θ_2	$k = 2$	0.014	0.030	0.014	0.038	0.072	0.002
		$k = 3$	0.014	0.030	0.014	0.038	0.072	0.002

Notes: Based on 500 simulated samples, each with the sample size of 500. The number of bootstrap samples is 199. We report only up to $k = 3$ because the results of $k = 4, 5$ are identical to that of $k = 3$. The table shows the frequencies that the confidence intervals missed the true values of $\theta_0 = 2.0$, $\theta_1 = -0.1$, and $\theta_2 = 0.05$ on the left or right side. For example, "Miss Left" for θ_0 means that the left endpoint was larger than 2.0. The true coverage is 0.9 so that the ideal values of "Miss Left" and "Miss Right" are 0.05. The variance-covariance matrix is computed using the OPG, where the derivatives of an actual likelihood are used for MLE while the derivatives of a pseudo likelihood defined in terms of 2 policy iterations are used for PI and MPI.

Table 12: Coverage Performance of Bootstrap 95% CIs for k-PL and k-MPL Estimators

			Symmetric CI		Equal-tailed CI		Percentile CI	
			Miss Left	Miss Right	Miss Left	Miss Right	Miss Left	Miss Right
k-PL Estimator		$k = 1$	0.052	0.024	0.040	0.028	0.094	0.012
		$k = 2$	0.046	0.022	0.056	0.024	0.078	0.012
	θ_0	$k = 3$	0.046	0.022	0.058	0.024	0.074	0.012
		$k = 4$	0.044	0.022	0.058	0.024	0.072	0.012
		$k = 5$	0.042	0.022	0.056	0.024	0.070	0.012
		$k = 1$	0.030	0.044	0.012	0.176	0.096	0.182
		$k = 2$	0.038	0.040	0.014	0.076	0.040	0.176
	θ_1	$k = 3$	0.050	0.032	0.028	0.048	0.020	0.160
		$k = 4$	0.062	0.040	0.036	0.042	0.018	0.156
		$k = 5$	0.066	0.030	0.042	0.042	0.016	0.148
		$k = 1$	0.016	0.048	0.052	0.030	0.054	0.036
		$k = 2$	0.016	0.042	0.020	0.046	0.064	0.018
	θ_2	$k = 3$	0.016	0.042	0.024	0.046	0.066	0.010
		$k = 4$	0.016	0.042	0.024	0.052	0.074	0.006
		$k = 5$	0.018	0.042	0.028	0.046	0.070	0.006
k-MPL Estimator		$k = 1$	0.030	0.042	0.024	0.036	0.066	0.014
		$k = 2$	0.008	0.032	0.036	0.020	0.004	0.010
	θ_0	$k = 3$	0.004	0.032	0.042	0.022	0.004	0.008
		$k = 4$	0.004	0.028	0.040	0.030	0.004	0.008
		$k = 5$	0.004	0.030	0.048	0.030	0.006	0.008
		$k = 1$	0.018	0.000	0.000	0.060	0.096	0.000
		$k = 2$	0.042	0.000	0.008	0.052	0.054	0.000
	θ_0	$k = 3$	0.066	0.000	0.030	0.052	0.044	0.000
		$k = 4$	0.076	0.000	0.042	0.052	0.032	0.000
		$k = 5$	0.084	0.000	0.056	0.052	0.034	0.000
		$k = 1$	0.004	0.028	0.002	0.022	0.026	0.018
		$k = 2$	0.004	0.014	0.002	0.012	0.046	0.002
	θ_0	$k = 3$	0.006	0.012	0.004	0.010	0.050	0.002
		$k = 4$	0.006	0.012	0.012	0.012	0.054	0.000
		$k = 5$	0.008	0.010	0.018	0.014	0.056	0.000

Notes: Based on 500 simulated samples, each with the sample size of 500. The number of bootstrap samples is 199. The table shows the frequencies that the confidence intervals missed the true values of $\theta_0 = 2.0$, $\theta_1 = -0.1$, and $\theta_2 = 0.05$ on the left or right side. For example, "Miss Left" for θ_0 means that the left endpoint was larger than 2.0. The true coverage is 0.9 so that the ideal values of "Miss Left" and "Miss Right" are 0.05. The variance-covariance matrix is computed using the OPG, where the derivatives of an actual likelihood are used for MLE while the derivatives of a pseudo likelihood defined in terms of 2 policy iterations are used for k-PL and k-MPL.