

Kasahara, Hiroyuki; Shimotsu, Katsumi

Working Paper

Nested Pseudo-likelihood Estimation and Bootstrap-based Inference for Structural Discrete Markov Decision Models

Queen's Economics Department Working Paper, No. 1063

Provided in Cooperation with:

Queen's University, Department of Economics (QED)

Suggested Citation: Kasahara, Hiroyuki; Shimotsu, Katsumi (2006) : Nested Pseudo-likelihood Estimation and Bootstrap-based Inference for Structural Discrete Markov Decision Models, Queen's Economics Department Working Paper, No. 1063, Queen's University, Department of Economics, Kingston (Ontario)

This Version is available at:

<https://hdl.handle.net/10419/189343>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Queen's Economics Department Working Paper No. 1063

Nested Pseudo-likelihood Estimation and Bootstrap-based Inference for Structural Discrete Markov Decision Models

Hiroyuki Kasahara
University of Western Ontario

Katsumi Shimotsu
Queen's University

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

2-2006

Nested Pseudo-likelihood Estimation and Bootstrap-based Inference for Structural Discrete Markov Decision Models*

Hiroyuki Kasahara
Department of Economics
University of Western Ontario
hkasahar@uwo.ca

Katsumi Shimotsu
Department of Economics
Queen's University
shimotsu@qed.econ.queensu.ca

February 4, 2006

Abstract

This paper analyzes the higher-order properties of nested pseudo-likelihood (NPL) estimators and their practical implementation for parametric discrete Markov decision models in which the probability distribution is defined as a fixed point. We propose a new NPL estimator that can achieve quadratic convergence without fully solving the fixed point problem in every iteration. We then extend the NPL estimators to develop one-step NPL bootstrap procedures for discrete Markov decision models and provide some Monte Carlo evidence based on a machine replacement model of Rust (1987). The proposed one-step bootstrap test statistics and confidence intervals improve upon the first order asymptotics even with a relatively small number of iterations. Improvements are particularly noticeable when analyzing the dynamic impacts of counterfactual policies.

Keywords: Edgeworth expansion, k-step bootstrap, maximum pseudo-likelihood estimators, nested fixed point algorithm, Newton-Raphson method, policy iteration.

JEL Classification Numbers: C12, C13, C14, C15, C44, C63.

*We are grateful to Victor Aguirregabiria, Chris Bennett, Christopher Ferrall, Silvia Gonçalves, Lance Lochner, James MacKinnon, John Rust, and seminar participants at Canadian Econometric Study Group, Indiana University, University of Maryland, and Queen's University for helpful comments. Shimotsu thanks the SSHRC for financial support.

1 Introduction

Understanding the dynamic response of individuals and firms is imperative for properly assessing various policy proposals. As numerous empirical studies have demonstrated, the estimation of dynamic structural models enhances our understanding of individual and firm behavior, especially when expectations play a major role in decision making.¹

The literature on estimating parametric discrete Markov decision models was pioneered by Rust (1987, 1988) who introduced the *nested fixed point algorithm* (NFXP). The NFXP requires repeatedly solving the fixed point problem during optimization and can be very costly when the dimensionality of state space is large. Hotz and Miller (1993) developed a simpler estimator, called the *conditional choice probabilities* (CCP) estimator, based on the *policy iteration* mapping—denoted by $\Psi(P, \theta)$ —which maps an arbitrary choice probability P and the model parameter θ to another choice probability. The true choice probability is characterized as a fixed point of the mapping, i.e., $P_\theta = \Psi(P_\theta, \theta)$. The CCP estimates the parameter θ by minimizing the discrepancy between the observed choice probabilities and $\Psi(\hat{P}^0, \theta)$, where \hat{P}^0 is an initial estimate. The CCP requires only one policy iteration to evaluate the objective function, leading to a significant computational gain over the NFXP.

Aguirregabiria and Mira (2002) [henceforth, AM] extended the CCP estimator and proposed the *nested pseudo-likelihood* (NPL) estimator. Upon obtaining $\hat{\theta}$ from the CCP, one can update the conditional choice probabilities estimate as $\hat{P}^1 = \Psi(\hat{P}^0, \hat{\theta})$, which provides a more accurate estimator of P_θ than \hat{P}^0 . Next, one can obtain another estimator of θ , $\hat{\theta}^1$, by using $\Psi(\hat{P}^1, \theta)$ instead of $\Psi(\hat{P}^0, \theta)$. Iterating this procedure generates a sequence of the NPL estimators, including the CCP as the initial element and the NFXP estimator as its limit. Somewhat surprisingly, AM showed that the NPL estimator for any number of iterations has the same limiting distribution as the NFXP estimator.

The NPL provides a menu of first-order equivalent estimators that empirical researchers can choose from, but little is known about their higher-order properties. Since the choice among these estimators involves a trade-off between efficiency and computational burden, understanding their higher-order properties is necessary for making an appropriate choice for a given situation.

¹Contributions include Miller (1984), Pakes (1986), Berkovec and Stern (1991), Rust (1987), Keane and Wolpin (1997), Rust and Phelan (1997), Gilleskie (1998), Eckstein and Wolpin (1999), Imai and Keane (2004).

In fact, the simulations by AM reveal that iterating the policy iteration mapping improves the accuracy of the parameter estimates, often by a substantial magnitude, suggesting that higher-order properties may be of practical importance.

We present the simulation results showing that tests based on first order asymptotics can be unreliable. While bootstrap tests are known to provide a better inferential tool than first-order asymptotic approximations, few studies have analyzed a bootstrap-based inference method for discrete Markov decision models. The main obstacle lies in the computational burden, because the bootstrap requires repeated parameter estimation under different simulated samples while it is not unusual for estimating one set of the parameters to take more than a day. This further increases the need for computationally attractive methods. Moreover, because the asymptotic improvement of the bootstrap relies on its higher-order properties, analyzing those properties is essential for practical applications.

The contributions of this paper are three-fold. First, we analyze the higher-order properties of the NPL estimator and derive the stochastic differences [c.f., Robinson (1988)] between the NFXP and the sequence of estimators generated by the NPL algorithm. We show the rate at which the sequence of the NPL estimators approaches the NFXP and provide a theoretical explanation for the simulation results in AM, in which iterating the NPL algorithm improves the accuracy of the NPL estimator.

Second, we propose two new estimators based on the NPL estimator. First, we develop a *nested modified pseudo-likelihood* (NMPL) estimator that uses a pseudo-likelihood defined in terms of *two policy iterations* as opposed to one policy iteration in the NPL. We show the convergence rate of the NMPL is faster than quadratic while that of the NPL is less than quadratic. Second, we propose a version of the NPL and NMPL estimators, called the *one-step* NPL and NMPL estimators, that use only one Newton-Raphson (NR) step to update the parameter θ during each iteration. By using only one NR step rather than fully solving the pseudo-likelihood problem for every iteration, we can reduce the computational cost significantly. The one-step NMPL estimator with the NR method achieves a quadratic convergence while the convergence rate of the one-step NPL estimator is less than quadratic.

Our one-step NPL and NMPL estimators are closely related to the k -step estimators analyzed by Pfanzagl (1974), Janssen, Jureckova, and Veraverbeke (1985), Robinson (1988), and Andrews (2002a), among others. Specifically, our one-step estimators may be viewed as a (semi-

parametric) k -step estimator in which an estimate of nuisance parameter P is updated between NR steps.

The key to understanding the convergence properties of the NPL and the NMPL algorithms is the orthogonality condition between the parameter of interest θ and the nuisance parameter P . When we define a pseudo-likelihood in terms of two policy iterations, θ and P become *orthogonal in any sample size*. This strengthens one of the key properties of the NPL that θ and P are *asymptotically orthogonal*. Consequently, the effect of the nuisance parameter P on the estimation of θ becomes negligible at a faster rate in the NMPL than in the NPL, leading to their different convergence rates.

The superior convergence properties of the NMPL over the NPL is not without cost. The computational cost for each NR step is larger in the NMPL, because its pseudo-likelihood is defined in terms of two policy iterations in contrast to one policy iteration in the NPL. Comparing the number of policy iterations required to achieve a particular level of convergence suggests that the overall computational cost of the one-step NMPL may be lower than that of the one-step NPL when the target level of convergence is high.

Third, we develop a computationally attractive bootstrap procedure for parametric discrete Markov decision models, applying the framework developed by Davidson and MacKinnon (1999a) and Andrews (2002b, 2005). Starting with an estimate from the original sample, a bootstrap estimator is obtained with the bootstrap sample by using the (one-step) NPL and NMPL, where taking a small number of iterations suffices to achieve higher-order improvements. Since their computational burden is substantially less than that of the NFXP, our proposed bootstrap is feasible for many discrete Markov decision models where the standard bootstrap procedure is too costly to implement. The computational burden is further reduced because the covariance matrix can be consistently estimated in the bootstrap sample using the derivatives of a *pseudo-likelihood function* instead of the likelihood function based on the fixed point solution. The proofs of higher-order properties of the proposed algorithm build on the results developed in Andrews (2002a,b, 2005).

We also consider two extensions of our bootstrap procedure: counterfactual experiments and models with unobserved heterogeneity. When estimated structural models are used to quantitatively assess the impact of counterfactual policies, the reliability of the estimated impact arises as an important issue. We develop a bootstrap procedure that allows us to construct

reliable CIs for the impact of counterfactual policies where asymptotic CIs may be unreliable. We also show that our bootstrap procedure can be applied to a finite mixture model, which is a popular approach when preferences are likely to be different across individuals.

In order to assess the performance of our bootstrap procedure, we provide Monte Carlo evidence based on a machine replacement model of Rust (1987) and Cooper, Haltiwanger, and Power (1999). We compare the performance of the bootstrap CIs for the impact of counterfactual policies with that of the asymptotic CIs. The bootstrap CIs perform better than the asymptotic CIs, and the one-step bootstrap CIs with a few iterations often achieve a similar performance to the bootstrap CIs based on the NFXP. The simulation results suggest that we may construct more reliable CIs by using our proposed one-step bootstrap procedure without facing a prohibitive computational burden.

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 NPL estimator. Section 4 describes our one-step estimation algorithm and proves its convergence properties. Section 5 analyzes the higher-order improvements from applying parametric bootstrapping to the one-step NPL estimators. Practical extensions are discussed in Section 6, and Section 7 reports some simulation results. Proofs and technical results are collected in Appendices A and B.

2 The Econometric Model

This section introduces the class of discrete Markov decision models considered in this paper. We closely follow the setup and the notations of Aguirregabiria and Mira (2002) [AM, hereafter]. 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 be chosen from the discrete and finite set $A = \{1, 2, \dots, J\}$. The transition probabilities are 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 the 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 -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 $g(\epsilon|x)$ has finite first moments and is twice differentiable in ϵ uniformly in $x \in X$; the support of $\epsilon(a)$ is the real line for all a .

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

Assumptions 1 and 2 are analogous 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:

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

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

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

where $v(x, a) = u(x, a) + \beta \int_X V(x')f(dx'|x, a)$ 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 the mapping $\Lambda(V) : B_V \rightarrow B_P$ as

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

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

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

where

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

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

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: $\tilde{v}_x = Q_x^{-1}(P_x)$. Invertibility allows us to express the conditional expectations of $\epsilon(a)$ in terms of the choice probabilities P_x as $e_x(a, P_x) \equiv E[\epsilon(a)|x, a; Q_x^{-1}(P_x), P(a|x)]$.

By substituting these functions into (4), we obtain

$$V(x) = u_P(x) + \beta E_P V(x), \quad (5)$$

where $u_P(x) = \sum_{a \in A} P(a|x)[u(x, a) + e_x(a, P_x)]$ and $E_P V(x) = \sum_{a \in A} P(a|x) \int_X V(x') f(dx'|x, a)$. Here, u_P is the expected utility function implied by the conditional choice probability P_x whereas 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): $V = (I - \beta E_P)^{-1} u_P$. The right-hand side of this equation can be viewed as a mapping from the choice probability space B_P to the value function space B_V . Define this mapping as $\varphi(P) \equiv (I - \beta E_P)^{-1} u_P$. Then we may define a policy iteration operator Ψ as a composite operator of $\varphi(\cdot)$ and $\Lambda(\cdot)$:

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

Given the fixed point of this policy iteration operator, P , the fixed point of the Bellman equation (1) can be expressed as $V = \varphi(P)$.

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

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

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 is a 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 analogously to the corresponding concepts defined for the functions in Euclidean spaces. For further details the reader is referred to Zeider (1986). Ichimura and Lee (2004) provide a concise summary on F-derivatives. Let $D^j g(x, y)$ denote the j th 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 $\partial g(x, y)/\partial x$.

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 proves this property in the case where the support of x_t is finite. The following proposition establishes that this zero-Jacobian property also holds even when the support of x_t is not finite and V does not belong to a Euclidean space.

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 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 an $L_\theta \times 1$ parameter vector $\theta \equiv (\theta_u, \theta_g, \theta_f)$, where θ_u, θ_g , and θ_f are the parameter vectors in the utility function u , the density of unobservable state variables 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 $\{w_i : i = 1, 2, \dots, N\}$ be a random sample of $w = (a, x', x)$ from the population, where x_i is drawn from the stationary distribution implied by P_θ and f_{θ_f} , a_i is drawn conditional on x_i from $P_\theta(\cdot|x_i)$, and x'_i is drawn from $f_{\theta_f}(\cdot|x_i, a_i)$. Under Assumption 2, the log-likelihood function

can be decomposed into conditional choice probability and transition probability terms as:

$$l_N(\theta) = l_{N,1}(\theta) + l_{N,2}(\theta_f) = \sum_{i=1}^N \ln P_\theta(a_i|x_i) + \sum_{i=1}^N \ln f_{\theta_f}(x'_i|x_i, a_i). \quad (6)$$

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 $l_{N,2}(\theta_f)$. Thus, $\Psi(P, \theta) = \Psi(P, \alpha, \theta_f)$, and we use both $\Psi(P, \theta)$ and $\Psi(P, \alpha, \theta_f)$ henceforth.

The maximum likelihood estimator solves the following constrained maximization problem:

$$\max_{\alpha} \frac{1}{N} \sum_{i=1}^N \ln P(a_i|x_i) \quad s.t. \quad P = \Psi(P, \alpha, \hat{\theta}_f). \quad (7)$$

Rust (1987) develops the celebrated Nested Fixed Point (NFXP) algorithm by formulating the parameter restriction in terms of Bellman's equation. The NFXP repeatedly solves the fixed point problem at each parameter value to maximize the likelihood with respect to α . Let $\hat{\alpha}$ denote the solution to the maximization problem (7), and let \hat{P} denote the associated conditional choice probability estimate characterized by the fixed point: $\hat{P} = \Psi(\hat{P}, \hat{\alpha}, \hat{\theta}_f)$.

3.1 Nested Pseudo-likelihood (NPL) Estimator

Assuming an initial consistent estimator \hat{P}_0 is available, the nested pseudo-likelihood (NPL) estimator developed by AM is recursively defined as follows.

Step 1: Given \hat{P}_{j-1}^{PL} , update α by

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

Step 2: Update P using the obtained estimate $\hat{\alpha}_j^{PL}$ by $\hat{P}_j^{PL} = \Psi(\hat{P}_{j-1}^{PL}, \hat{\alpha}_j^{PL}, \hat{\theta}_f)$.

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

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 $\Theta = \Theta_\alpha \times \Theta_f$. Following AM, consider the following regularity conditions:

Assumption 4. (a) Θ_α and Θ_f are compact. (b) $\Psi(P, \alpha, \theta_f)$ is three times continuously F-differentiable. (c) $\Psi(P, \alpha, \theta_f)(a|x) > 0$ for any $(a, x) \in A \times X$ and any $\{P, \alpha, \theta_f\} \in$

$B_P \times \Theta_\alpha \times \Theta_f$. (d) $w_i = \{a_i, x'_i, x_i\}$, for $i = 1, 2, \dots, N$, 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 . (e) There is a unique $\theta_f^0 \in \text{int}(\Theta_f)$ such that, for any $(a, x, x') \in A \times X \times X$, $f_{\theta_f^0}(x'|x, a) = f^0(x'|x, a)$. (f) There is a unique $\alpha^0 \in \text{int}(\Theta_\alpha)$ such that, for any $(a, x) \in A \times X$, $P_{\theta_f^0}(a|x) = P^0(a|x)$. For any $\alpha \neq \alpha^0$, $\Pr_{\theta_f^0}(\{(a, x) : \Psi(P^0, \alpha, \theta_f^0)(a|x) \neq P^0(a|x)\}) > 0$. (g) $E_{\theta_f^0} \sup_{(P, \alpha, \theta_f)} \|D^s \Psi(P, \alpha, \theta_f)(a|x)\|^2 < \infty$ for $s = 1, \dots, 4$. (h) $\hat{\theta}_f - \theta_f^0 = O_p(N^{-1/2})$, $\hat{P}_0^{PL} - P^0 = o_p(1)$, and the NFXP estimator $\hat{\alpha}$ satisfies $\sqrt{N}(\hat{\alpha} - \alpha^0) \rightarrow_d N(0, \Omega)$.

Assumptions 4(a)–4(f) are similar to the regularity conditions 4(a)–(f) in AM. The supremum in 4(g) may be taken in a neighborhood of $(P^0, \alpha^0, \theta_f^0)$.

Following Robinson (1988), for matrix/mapping 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 .

Our first main result shows that the NPL estimator converges to the MLE, $\hat{\alpha}$, at a superlinear, but less than quadratic, convergence rate.

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

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

This proposition provides a theoretical explanation for the result of the AM's Monte Carlo experiment. Their experiment illustrates that the finite sample properties of the NPL estimators improve monotonically with k and that the estimators with $k = 2$ or 3 substantially outperform the estimator with $k = 1$.

Note that $\hat{P}_0^{PL} - P^0 = O_p(N^{-b})$ with $b > 1/4$ suffices for $\sqrt{N}(\hat{\alpha}_k^{PL} - \alpha^0) \rightarrow_d N(0, \Omega)$ for all $k \geq 1$. This weakens assumption (g) of Proposition 4 of AM and also implies that the NPL estimator is valid even if x_t has an infinite support and a kernel-based estimator is used to estimate P^0 . The result suggests that the NPL algorithm may work even with relatively imprecise initial estimates of the conditional choice probabilities.

If $\hat{P}_0^{PL} - P^0 = O_p(N^{-b})$ with $b \in (1/4, 1/2]$, repeated substitution gives

$$\hat{\alpha}_k^{PL} - \hat{\alpha} = O_p(N^{-(k-1)/2-2b}), \quad \hat{P}_k^{PL} - \hat{P} = O_p(N^{-(k-1)/2-2b}). \quad (8)$$

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

3.2 Nested Modified Pseudo-likelihood (NMPL) Estimator

We now introduce the nested modified pseudo-likelihood (NMPL) estimator that achieves a faster rate of convergence than the NPL estimator:

Step 1: Given \hat{P}_{j-1}^{MPL} , update α by

$$\hat{\alpha}_j^{MPL} = \arg \max_{\alpha} \frac{1}{N} \sum_{i=1}^N \ln \Psi_2(\hat{P}_{j-1}^{MPL}, \alpha, \hat{\theta}_f)(a_i | x_i),$$

where

$$\Psi_2(P, \alpha, \theta_f)(a_i | x_i) \equiv \Psi(\Psi(P, \alpha, \theta_f), \alpha, \theta_f)(a_i | x_i).$$

Step 2: Update P using the obtained estimate $\hat{\alpha}_j^{MPL}$ by $\hat{P}_j^{MPL} = \Psi(\hat{P}_{j-1}^{MPL}, \hat{\alpha}_j^{MPL}, \hat{\theta}_f)$.

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

Assumption 5. (a) For any $\alpha \neq \alpha^0$, $\Pr_{\theta^0}(\{(a, x) : \Psi_2(P^0, \alpha, \theta_f^0)(a|x) \neq P^0(a|x)\}) > 0$. (b) $E_{\theta^0} \sup_{(P, \alpha, \theta_f)} \|D^s \Psi_2(P, \alpha, \theta_f)(a_i | x_i)\|^2 < \infty$ for $s = 1, \dots, 4$. (c) $\hat{P}_0^{MPL} - P^0 = o_p(1)$.

The following proposition shows the NMPL estimator of α converges at a rate faster than quadratic while the NMPL estimator of P converges at a quadratic rate.

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

$$\hat{\alpha}_k^{MPL} - \hat{\alpha} = O_p(N^{-1/2} \|\hat{P}_{k-1}^{MPL} - \hat{P}\|^2 + \|\hat{P}_{k-1}^{MPL} - \hat{P}\|^3), \quad \hat{P}_k^{MPL} - \hat{P} = O_p(\|\hat{P}_{k-1}^{MPL} - \hat{P}\|^2).$$

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

$$\hat{\alpha}_k^{MPL} - \hat{\alpha} = O_p(N^{-1/2-b2^k} + N^{-3b2^{k-1}}), \quad \hat{P}_k^{MPL} - \hat{P} = O_p(N^{-b2^k}).$$

In particular, if $\hat{P}_0^{MPL} - P^0 = O_p(N^{-1/2})$, then we have $\hat{\alpha}_k^{MPL} - \hat{\alpha} = O_p(N^{-1/2-2^{k-1}})$. Note that $\hat{P}_0^{MPL} - P^0 = O_p(N^{-b})$ with $b > 1/6$ suffices for $\sqrt{N}(\hat{\alpha}_k^{MPL} - \alpha^0) \rightarrow_d N(0, \Omega)$ for all $k \geq 1$. Therefore, the NMPL estimator requires a weaker condition on the initial estimate of P^0 than the NPL estimator. The NMPL estimator may, therefore, be preferable to the NPL estimator

when we only have a poor initial estimate of P^0 , as is likely to be the case, for instance, in models with unobserved heterogeneity.

Using $\Psi_2(P, \alpha, \theta_f)$ instead of $\Psi(P, \alpha, \theta_f)$ achieves a faster rate of convergence. However, the NMPL algorithm requires more policy iterations than the NPL for computing each $\hat{\alpha}_j$, which implies that the overall computational cost for achieving a given rate of convergence may be higher with the NMPL.

The following two orthogonality conditions between $\hat{\alpha}$ and \hat{P} are the key to understanding the difference in the rates of convergence between the NPL and the NMPL estimators:²

$$\begin{aligned} N^{-1} \sum_{i=1}^N D_{P\alpha} \ln \Psi(P_{\hat{\theta}}, \hat{\theta})(a_i|x_i) &= O_p(N^{-1/2}), \\ N^{-1} \sum_{i=1}^N D_{P\alpha} \ln \Psi_2(P_{\hat{\theta}}, \hat{\theta})(a_i|x_i) &= 0. \end{aligned} \tag{9}$$

Thus, at the fixed point, $\hat{\alpha}$ and \hat{P} are *asymptotically* orthogonal in the NPL while they are orthogonal *in any sample size* in the NMPL. In case of the NPL, the asymptotic orthogonality in the first equation of (9) implies that the estimation error $\hat{P}_{k-1}^{PL} - \hat{P}$ has an asymptotically negligible effect on $\hat{\alpha}_k^{PL} - \hat{\alpha}$, diminishing at the rate of $N^{-1/2}$. Since the extent to which the impreciseness of \hat{P}_{k-1}^{PL} would be carried over to the estimate $\hat{\alpha}_k^{PL}$ is mitigated only at the rate of $N^{-1/2}$, the NPL converges at a superlinear, but less than quadratic, rate. In case of the NMPL, the second equation of (9) implies that $\hat{P}_{k-1}^{MPL} - \hat{P}$ has, at most, a second-order effect on $\hat{\alpha}_k^{MPL} - \hat{\alpha}$ for any sample size N and hence the NMPL converges, at least, at a quadratic rate. In the appendix, we also show that $N^{-1} \sum_{i=1}^N D_{PP\alpha} \ln \Psi_2(P_{\hat{\theta}}, \hat{\theta})(a_i|x_i) = O_p(N^{-1/2})$ [c.f., Lemma 9(b)], implying that the second-order effect is diminishing at the rate of $N^{-1/2}$, and thus the NMPL converges at a faster rate than quadratic.

3.3 Covariance Matrix Estimation and Test Statistics

Suppose $\hat{\theta}_f$ is obtained by maximizing $l_{N,2}(\theta_f)$. 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 of $\hat{\theta} = (\hat{\alpha}', \hat{\theta}_f)'$ as

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

²They follow from Lemma 8(c), the root- N consistency of $\hat{\theta}$, and Lemma 9(a).

where

$$D(\theta) = \begin{bmatrix} D_{11}(\theta) & D_{12}(\theta) \\ 0 & D_{22}(\theta) \end{bmatrix} = - \begin{bmatrix} E(\partial^2/\partial\alpha\partial\alpha') \ln P_\theta & E(\partial^2/\partial\alpha\partial\theta'_f) \ln P_\theta \\ 0 & E(\partial^2/\partial\theta_f\partial\theta'_f) \ln f_{\theta_f} \end{bmatrix},$$

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

The information matrix equality from the MLE based on $l_{N,2}(\theta)$ alone implies $D_{22}(\theta^0) = V_{22}(\theta^0)$, and the information matrix equality from the full MLE based on $l_N(\theta)$ implies $D_{11}(\theta^0) = V_{11}(\theta^0)$ and $-E(\partial^2/\partial\alpha\partial\theta'_f) \ln P_{\theta^0} = E(\partial/\partial\alpha) \ln P_{\theta^0} (\partial/\partial\theta'_f) (\ln P_{\theta^0} + \ln f_{\theta_f^0})$.

There are several ways to estimate $\Sigma(\theta^0)$ consistently. Let $D_N(\theta)$ and $V_N(\theta)$ be the sample analogue of $D(\theta)$ and $V(\theta)$, respectively, and define

$$D_N^O(\theta) = \frac{1}{N} \sum_{i=1}^N \begin{bmatrix} (\partial/\partial\alpha) \ln P_{\theta_i} (\partial/\partial\alpha') \ln P_{\theta_i} & (\partial/\partial\alpha) \ln P_{\theta_i} (\partial/\partial\theta'_f) (\ln P_{\theta_i} + \ln f_{\theta_f}) \\ 0 & (\partial/\partial\theta_f) \ln f_{\theta_f} (\partial/\partial\theta'_f) \ln f_{\theta_f} \end{bmatrix}.$$

$D_N^O(\theta)$ is an outer-product-of-the-gradient (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

$$\begin{aligned} \Sigma_N(\theta) &= D_N(\theta)^{-1} V_N(\theta) (D_N(\theta)^{-1})', \quad \text{or} \\ \Sigma_N(\theta) &= D_N^O(\theta)^{-1} V_N(\theta) (D_N^O(\theta)^{-1})'. \end{aligned} \tag{10}$$

The consistency of $\Sigma_N(\bar{\theta})$ follows from the standard argument. Notice, however, that computing $\Sigma_N(\bar{\theta})$ 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 NPL and NMPL 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 in the following Proposition, we can estimate $\Sigma(\theta^0)$ consistently using these estimates with the NPL and NMPL estimators of (P, α) and construct t - and Wald statistics with a limited number of policy iterations.

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

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 \mathbb{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 $T_N(\theta_r^0) \rightarrow_d N(0, 1)$ and $\mathcal{W}_N(\theta^0) \rightarrow_d \chi_{L_\eta}^2$ under the null hypotheses.

4 One-step NPL and NMPL Estimators

We propose one-step NPL and NMPL estimators which update the parameter α using one Newton step without fully solving the optimization problem. This reduces the computational cost of the corresponding estimators especially when the dimension of α is high. Let $L_N(P, \alpha, \theta_f)$ denote the objective function of the NPL estimator as

$$L_N(P, \alpha, \theta_f) = \frac{1}{N} \sum_{i=1}^N \ln \Psi(P, \alpha, \theta_f)(a_i | x_i). \quad (11)$$

The one-step NPL 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

$$\tilde{\alpha}_j^{PL} = \tilde{\alpha}_{j-1}^{PL} - (Q_{N,j-1})^{-1} \frac{\partial}{\partial \alpha'} L_N(\tilde{P}_{j-1}^{PL}, \tilde{\alpha}_{j-1}^{PL}, \hat{\theta}_f), \quad (12)$$

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 one-step NPL estimator uses the NR, default NR, line-search NR, or Gauss-Newton (GN) steps. The NR choice of $Q_{N,j-1}$ is $Q_{N,j-1}^{NR} = (\partial^2/\partial\alpha\partial\alpha')$ $L_N(\tilde{P}_{j-1}^{PL}, \tilde{\alpha}_{j-1}^{PL}, \hat{\theta}_f)$. The default NR choice of $Q_{N,j-1}$, denoted $Q_{N,j-1}^D$ equals $Q_{N,j-1}^{NR}$ if $\tilde{\alpha}_j^{PL}$ defined in (12) satisfies $L_N(\tilde{P}_{j-1}^{PL}, \tilde{\alpha}_j^{PL}, \hat{\theta}_f) \geq L_N(\tilde{P}_{j-1}^{PL}, \tilde{\alpha}_{j-1}^{PL}, \hat{\theta}_f)$, but equals some other matrix otherwise. Typically, $(1/\varepsilon)I_{\dim(\alpha)}$ for some small $\varepsilon > 0$ is used. The line-search NR choice, $Q_{N,j-1}^{LS}$, computes $\tilde{\alpha}_j^{PL,\lambda}$ for $\lambda \in (0, 1]$ using $(1/\lambda)Q_{N,j-1}^{NR}$ and chooses the one that maximizes the objective function. The GN choice, denoted $Q_{N,j-1}^{GN}$, uses a matrix that approximates the NR matrix $Q_{N,j-1}^{NR}$. A popular choice is the OPG estimator

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

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

The following proposition establishes that the one-step NPL estimator achieves a similar rate of convergence to the original NPL estimator. This is because taking one NR step brings the one-step NPL estimator sufficiently close to the NPL estimator. In fact, the distance between the one-step NPL estimator and the NPL estimator is at most of the same order of magnitude as the distance between the NFXP estimator and the NPL estimator.

Proposition 5 *Suppose the assumptions of Proposition 2 hold and the initial estimates $(\tilde{\alpha}_0^{PL}, \tilde{P}_0^{PL})$ are consistent. Then, for $k = 1, 2, \dots$,*

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

If the initial estimates satisfy $\tilde{\alpha}_0^{PL} - \alpha^0, \tilde{P}_0^{PL} - P^0 = O_p(N^{-b})$ with $b \in (1/4, 1/2]$, then repeated substitution gives³

$$\tilde{\alpha}_k^{PL} - \hat{\alpha} = O_p(N^{-(k-1)/2-2b}), \quad \tilde{P}_k^{PL} - \hat{P} = O_p(N^{-(k-1)/2-2b}), \quad (13)$$

³The initial root- N consistent estimate, $\tilde{\alpha}_0^{PL}$, can be obtained from applying the original NPL estimator with $k = 1$ or using Hotz and Miller's CCP estimator. Furthermore, when we apply the one-step NPL estimator to the bootstrap-based inference, we may use the estimate from the original sample as an initial root- N consistent estimate for the bootstrap sample.

and the one-step NPL estimator achieves the same convergence rate as the NPL estimator.

The one-step NMPL estimator $(\tilde{\alpha}_k^{MPL}, \tilde{P}_k^{MPL})$ 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 quadratic rate of convergence when the NR, default NR, or line-search NR is used. When the OPG is used, however, its convergence rate reduces to that of the one-step NPL estimator.

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

$$\begin{aligned} \tilde{\alpha}_k^{MPL} - \hat{\alpha} &= O_p(\|\tilde{\alpha}_{k-1}^{MPL} - \hat{\alpha}\|^2 + N^{-1/2}\|\tilde{P}_{k-1}^{MPL} - \hat{P}\|^2 + \|\tilde{P}_{k-1}^{MPL} - \hat{P}\|^3) \\ &\quad [+O_p(N^{-1/2}\|\tilde{\alpha}_{k-1}^{MPL} - \hat{\alpha}\| + \|\tilde{P}_{k-1}^{MPL} - \hat{P}\|^2) \text{ for OPG}], \\ \tilde{P}_k^{MPL} - \hat{P} &= O_p(\|\tilde{\alpha}_k^{MPL} - \hat{\alpha}\| + \|\tilde{P}_{k-1}^{MPL} - \hat{P}\|^2). \end{aligned}$$

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

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

For the NR, the default NR, and the line-search NR, the result follows from a quadratic convergence of NR iterations. For the OPG estimator, the convergence rate is less than quadratic because the matrix $Q_{N,j-1}^{OPG}$ approximates $(\partial^2/\partial\alpha\partial\alpha')L_N$, leading to an approximation error of the magnitude $O_p(N^{-1/2})$ in the NR search direction.

Comparing the number of policy iterations required to achieve a particular level of convergence with these estimators reveals that the one-step NMPL estimator requires fewer policy iterations than the one-step NPL estimator when the target level of convergence is high. We may also consider a hybrid algorithm that needs the fewest policy iterations by using the one-step NPL estimator for the first few steps and then switching to the one-step NMPL estimator.

5 Parametric Bootstrap and Higher-order Improvements

In this section, building upon Andrews (2005), we analyze the higher-order improvements from applying parametric bootstrapping to the parametric discrete Markov decision models.

5.1 The NFXP Parametric Bootstrap

First, consider bootstrapping the NFXP estimator. The parametric bootstrap sample $\{w_i^* : i = 1, \dots, n\}$ is generated using the parametric density at the (unrestricted) 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 $\theta^0 = (\alpha^{0'}, \theta_f^{0'})'$.⁴

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

$$\begin{aligned} \theta_f^* &= \arg \max_{\theta_f \in \Theta_f} l_{N,2}^*(\theta_f), \quad \text{where } l_{N,2}^*(\theta_f) = \frac{1}{N} \sum_{i=1}^N \ln f_{\theta_f}(x_i^* | x_i^*, a_i^*), \\ \alpha^* &= \arg \max_{\alpha \in \Theta_\alpha} \frac{1}{N} \sum_{i=1}^N \ln P(a_i^* | x_i^*) \quad \text{s.t. } P = \Psi(P, \alpha, \theta_f^*). \end{aligned} \quad (14)$$

The bootstrap covariance matrix estimator, Σ_N^* , is defined as $\Sigma_N^*(\theta^*)$ where $\Sigma_N^*(\theta)$ has the same definition as $\Sigma_N(\theta)$ in (10) but with the bootstrap sample in place of the original sample. The bootstrap t and Wald statistics are defined as

$$\begin{aligned} T_N^*(\hat{\theta}_r) &= N^{1/2}(\theta_r^* - \hat{\theta}_r) / (\Sigma_N^*)_{rr}^{1/2}, \\ \mathcal{W}_N^*(\hat{\theta}) &= H_N^*(\theta^*, \hat{\theta})' H_N^*(\theta^*, \hat{\theta}), \quad \text{where} \\ H_N^*(\theta, \hat{\theta}) &= \left(\frac{\partial}{\partial \theta'} \eta(\theta) \Sigma_N^*(\theta) \frac{\partial}{\partial \theta} \eta(\theta) \right)^{-1/2} N^{1/2}(\eta(\theta) - \eta(\hat{\theta})), \end{aligned} \quad (15)$$

where θ_r^* denotes the r -th element of θ^* , and $(\Sigma_N^*)_{rr}$ denotes the (r, r) -th element of Σ_N^* . Here, we use the bootstrap Wald statistics to test $H_0 : \eta(\theta^0) = 0$ versus $H_A : \eta(\theta^0) \neq 0$.

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

$$CI_{SYM}(\hat{\theta}_r) = [\hat{\theta}_r - z_{|T|, \alpha}^* (\Sigma_N(\hat{\theta}))_{rr}^{1/2} / N^{1/2}, \hat{\theta}_r + z_{|T|, \alpha}^* (\Sigma_N(\hat{\theta}))_{rr}^{1/2} / N^{1/2}]. \quad (16)$$

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

$$CI_{ET}(\hat{\theta}_r) = [\hat{\theta}_r - z_{T, \alpha/2}^* (\Sigma_N(\hat{\theta}))_{rr}^{1/2} / N^{1/2}, \hat{\theta}_r - z_{T, 1-\alpha/2}^* (\Sigma_N(\hat{\theta}))_{rr}^{1/2} / N^{1/2}]. \quad (17)$$

⁴If x_i is assumed to be exogenous, then $x_i^* = x_i$ needs to be used. If x_i is assumed to be drawn from its stationary distribution $\lambda(\theta)$ implied by P_θ and f_{θ_f} , then x_i^* is either equal to x_i or drawn from $\lambda(\hat{\theta})$.

The symmetric two-sided bootstrap t test of $H_0 : \theta_r = \theta_r^0$ versus $H_1 : \theta_r \neq \theta_r^0$ at significance level α rejects H_0 if $|T_N(\theta_r^0)| > z_{|T|, \alpha}^*$. The equal-tailed two-sided bootstrap t test at 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 rejects H_0 if $\mathcal{W}_N(\theta^0) > z_{\mathcal{W}, \alpha}^*$.

We introduce 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 4.1-4.3 in Andrews (2005). Let c be a non-negative constant such that $2c$ is an integer. Let $g(w_i, \theta) = ((\partial/\partial\theta') \ln P_\theta(a|x), (\partial/\partial\theta'_f) \ln f_{\theta_f}(x'|x, a))'$, and let $h(w_i, \theta) \in \mathbb{R}^{L_h}$ denote the vector containing the unique components of $g(w_i, \theta)$ and $g(w_i, \theta)g(w_i, \theta)'$ and their partial derivatives with respect to θ through order $d = \max\{2c + 2, 3\}$. Let $\lambda_{\min}(A)$ denote the smallest eigenvalue of the matrix A . Let $d(\theta, B)$ denote the distance between the point θ and the 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 reason why these sets need to be considered, see Andrews (2005).

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 $E_{\theta^0} \ln P_{\theta^0}(a_i|x_i) - E_{\theta^0} \ln P_\theta(a_i|x_i) > \eta$ and $E_{\theta^0} \ln f_{\theta_f}(x'_i|x_i, a_i) - E_{\theta^0} \ln f_{\theta_f}(x'_i|x_i, a_i) > \eta$ for all $\theta \in \Theta$ and $\theta^0 \in \Theta_1$. (c) $\sup_{\theta^0 \in \Theta_1} E_{\theta^0} \sup_{\theta \in \Theta} \|g(w_i, \theta)\|^{q_0} < \infty$, $\sup_{\theta^0 \in \Theta_1} E_{\theta^0} \sup_{\theta \in \Theta} \{|\ln P_\theta(a_i|x_i)|^{q_0} + |\ln f_{\theta_f}(x'_i|x_i, a_i)|^{q_0}\} < \infty$ for all $\theta \in \Theta$ for $q_0 = \max\{2c + 1, 2\}$.

Assumption 7. (a) $g(w, \theta)$ is $d = \max\{2c + 2, 3\}$ times partially differentiable with respect to θ on Θ_2 for all $w = (a, x', x) \in A \times X \times X$. (b) $\sup_{\theta^0 \in \Theta_1} E_{\theta^0} \|h(w_i, \theta^0)\|^{q_1} < \infty$ for some $q_1 > 2c + 2$. (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 $q_1 > 2c + 2$.

Assumption 8. (a) For all $\varepsilon > 0$, there exists a positive δ such that for all $t \in \mathbb{R}^{L_h}$ with $\|t\| > \varepsilon$, $|E_{\theta^0} \exp(it'h(w_i, \theta^0))| \leq 1 - \delta$ for all $\theta^0 \in \Theta_1$. (b) $\text{Var}_{\theta^0}(h(w_i, \theta^0))$ 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 , and the utility function, u_θ , are sufficiently smooth. Note that Assumption 4.1(b) of Andrews (2005) is satisfied by the definition of $\hat{\alpha}$ and $\hat{\theta}_f$. Assumption 4.1(c) of Andrews (2005) 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 4.1(d) of Andrews (2005) is satisfied by Assumption 6(b). Because w_i is iid, Assumption 4.3(a), (b), and (d) of Andrews (2005) are trivially satisfied, and his Assumption 4.3(c) reduces to the standard Cramér condition. Assumption 4.3(f) of Andrews (2005) follows from our Assumption 8(b) since w_i is iid. Assumption 8(a), however, is not satisfied when all elements of the observed state variable have a finite support.

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

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

- (a) $\sup_{\theta^0 \in \Theta_0} |\Pr_{\theta^0}(\theta_r^0 \in CI_{SYM}(\hat{\theta}_r)) - (1 - \alpha)| = O(N^{-2})$ for $c = 2$,
- (b) $\sup_{\theta^0 \in \Theta_0} |\Pr_{\theta^0}(\theta_r^0 \in CI_{ET}(\hat{\theta}_r)) - (1 - \alpha)| = o(N^{-1} \ln N)$ for $c = 1$,
- (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 $c = 3/2$.

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 rejection probability of a standard Wald test are $O(N^{-1})$. Davidson and MacKinnon (1999b) and Kim (2005) analyze an alternative parametric bootstrap procedure that draws the bootstrap sample using the restricted MLE where the null is imposed. The results in Davidson and MacKinnon and Kim indicate that the bootstrap equal-tailed t -test from the restricted parametric bootstrap have smaller errors in rejection probabilities than the unrestricted parametric bootstrap. In this paper, we mainly focus on CIs, but we conjecture that such a refinement from bootstrapping with the restricted MLE is also possible in our context.

5.2 One-step NPL and NMPL Parametric Bootstrap

Bootstrapping the NFXP estimator is computationally costly because one has to estimate the model repeatedly under different bootstrap samples, where each estimation requires the repeated full solution of the Bellman equation. For this reason, we propose the one-step boot-

strap NPL and NMPL 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 (14) 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 $\{w_i^* : i = 1, \dots, n\}$.

We estimate θ by the NFXP estimator in the original sample and use the fixed point at the NFXP estimator $P_{\hat{\theta}}$ as the initial estimate of P for the one-step estimation with the bootstrap samples. Using the NFXP and $P_{\hat{\theta}}$ does not increase the computational burden significantly, since we are required to estimate θ and compute $P_{\hat{\theta}}$ only once in the original sample.⁵

We use the derivatives of the pseudo-likelihood function defining the NPL or NMPL estimator to construct the covariance matrix estimate (c.f., Proposition 4). This is essential for developing computationally attractive bootstrap-based inference in this context. Evaluating the derivatives of the pseudo-likelihood functions involves a limited number of policy iterations and, under the assumption of extreme-value distributed unobserved state variables, the analytical expression for the first derivatives are available. The computational saving from using the pseudo-covariance matrix estimate can be substantial, since we need to compute the covariance matrix estimates as many times as the number of bootstraps.

With $(P_k^{*PL}, \theta_k^{*PL})$, we use the bootstrap covariance matrix estimator as

$$\Sigma_N^*(P, \theta) = D_N^{*O,PL}(P, \theta)^{-1} V_N^{*PL}(P, \theta) (D_N^{*O,PL}(P, \theta)^{-1})', \quad (18)$$

where $D_N^{*O,PL}(P, \theta)$ and $V_N^{*PL}(P, \theta)$ are the same as $D_N^{O,PL}(P, \theta)$ and $V_N^{PL}(P, \theta)$ but constructed with the bootstrap sample. Here, care must be exercised; using the bootstrap covariance matrix estimator defined as $D_N^{*PL}(P_k^{*PL}, \theta_k^{*PL})^{-1} V_N^{*PL}(P_k^{*PL}, \theta_k^{*PL}) (D_N^{*PL}(P_k^{*PL}, \theta_k^{*PL})^{-1})'$ *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 when evaluated at the fixed point.

With $(P_k^{*MPL}, \theta_k^{*MPL})$, we use either

$$\begin{aligned} \Sigma_N^*(P, \theta) &= D_N^{*MPL}(P, \theta)^{-1} V_N^{*MPL}(P, \theta) (D_N^{*MPL}(P, \theta)^{-1})', \quad \text{or} \\ \Sigma_N^*(P, \theta) &= D_N^{*O,MPL}(P, \theta)^{-1} V_N^{*MPL}(P, \theta) (D_N^{*O,MPL}(P, \theta)^{-1})', \end{aligned} \quad (19)$$

⁵Alternatively, we may estimate θ by the NPL or NMPL estimator in the original sample and use \hat{P}_k^{PL} or \hat{P}_k^{MPL} as the initial estimate for the bootstrap estimation. Here, we focus on the case of estimating θ by the NFXP estimator but the similar argument applies to the case of estimating θ by the NPL or NMPL estimator in the original sample.

with analogous definitions for $D_N^{*MPL}(P, \theta)$, $V_N^{*MPL}(P, \theta)$, and $D_N^{*O,MPL}(P, \theta)$. It is important to note that $D_N^{*MPL}(P, \theta)$ must be used if $D_N(\theta)$ is used in forming $\Sigma_N(\theta)$, and $D_N^{*O,MPL}(P, \theta)$ must be used if $D_N^O(\theta)$ is used in forming $\Sigma_N(\theta)$. For instance, using $D_N^{*MPL}(P, \theta)$ when $D_N^O(\theta)$ is used in forming $\Sigma_N(\theta)$ introduces an approximation error of magnitude $O_p(N^{-1/2})$ and, hence, does not yield the higher-order refinement.

The one-step bootstrap t - and Wald statistics, $T_{N,k}^*(\hat{\theta}_r)$ and $\mathcal{W}_{N,k}^*(\hat{\theta})$, are defined as in (15), but with (θ^*, Σ_N^*) replaced by $(\theta_k^{*PL}, \Sigma_N^*(P_k^{*PL}, \theta_k^{*PL}))$ or $(\theta_k^{*MPL}, \Sigma_N^*(P_k^{*MPL}, \theta_k^{*MPL}))$. The one-step bootstrap CIs, denoted $CI_{SYM,k}$, $CI_{ET,k}$, are defined analogously to (16) and (17) but using the $1 - \alpha$ quantiles of $|T_{N,k}^*(\hat{\theta}_r)|$ and $T_{N,k}^*(\hat{\theta}_r)$ instead of $|T_N^*(\hat{\theta}_r)|$ and $T_N^*(\hat{\theta}_r)$.

Define

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

Lemma 2 establishes the higher-order equivalence of the one-step NPL and NMPL bootstrap estimators and NFXP bootstrap estimator. Lemma 3 shows, under suitable conditions on c and k , the difference between the bootstrap test statistics constructed using the one-step NPL or NMPL estimator and the NFXP estimator is $o(N^{-c})$.

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

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

Lemma 3 *Suppose the assumptions of Lemma 2 hold and $\mu_{N,k} = o(N^{-(c+1/2)})$. Then, for all $\varepsilon > 0$,*

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

for $\Xi_k(z) = \Pr_{\hat{\theta}}^*(N^{1/2}(\theta_k^{*s} - \hat{\theta}) \leq z) - \Pr_{\hat{\theta}}^*(N^{1/2}(\theta^* - \hat{\theta}) \leq z)$ 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 $\Pr_{\hat{\theta}}^*(\mathcal{W}_{N,k}^*(\hat{\theta}) \leq z) - \Pr_{\hat{\theta}}^*(\mathcal{W}_N^*(\hat{\theta}) \leq z)$.

Admittedly, the additional finiteness assumptions on the derivatives of P and Ψ are strong. We conjecture they can be weakened to assumptions in terms of their moments, but doing so would require a longer proof. The following Lemma shows that the errors in coverage probability of the one-step NPL and NMPL bootstrap CIs are the same as those of the NFXP bootstrap CIs. Therefore, the one-step bootstrap estimators achieve the same level of higher-order refinement as the NFXP bootstrap estimator.

Lemma 4 *Suppose the assumptions of Lemma 2 hold.*

- (a) *If $c = 2$ and $\mu_{N,k} = o(N^{-5/2})$, then $\sup_{\theta^0 \in \Theta_0} |\Pr_{\theta^0}(\theta_r^0 \in CI_{SYM,k}(\hat{\theta}_r)) - (1 - \alpha)| = O(N^{-2})$.*
- (b) *If $c = 1$ and $\mu_{N,k} = o(N^{-3/2})$, then $\sup_{\theta^0 \in \Theta_0} |\Pr_{\theta^0}(\theta_r^0 \in CI_{ET,k}(\hat{\theta}_r)) - (1 - \alpha)| = o(N^{-1} \ln N)$.*
- (c) *If $c = 3/2$ and $\mu_{N,k} = o(N^{-3/2})$, then $\sup_{\theta^0 \in \Theta_0} |\Pr_{\theta^0}(\mathcal{W}_N(\theta^0) \leq z_{\mathcal{W},\alpha}^*) - (1 - \alpha)| = o(N^{-3/2} \ln N)$.*

The condition $\mu_{N,k} = o(N^{-5/2})$ requires $k \geq 3$ for the one-step NMPL estimator with the NR, default NR, and line-search NR, and requires $k \geq 5$ for the one-step NPL estimator and the one-step NMPL estimator with the OPG. Constructing a one-step NMPL bootstrap-t statistic requires 8 policy iterations. This is because the one-step bootstrap NMPL estimator with $k = 3$ requires 6 policy iterations and the pseudo-covariance matrix estimator based on the second equation of (19) requires 2 policy iterations.⁶ On the other hand, constructing a one-step NPL bootstrap-t statistic requires 6 policy iterations by using the one-step NPL estimator with $k = 5$ and using (18), and hence fewer computation. The fewest policy iterations with $\mu_{N,k} = o(N^{-5/2})$ are achieved if we use the one-step NPL estimator in the first and second iterations, the one-step NMPL estimator in the third iteration, and using the pseudo-covariance matrix estimator based on (18); this yields $\mu_{N,k} = O(N^{-3} \ln^6(N))$ with 5 policy iterations.

The NPL and NMPL estimators yield the same level of higher-order refinement as stated in Lemma 4 except that, reflecting the difference in their convergence rates, the definition of $\mu_{N,k}$ for the NMPL estimator is different from that for the one-step NMPL estimator. Specifically, we have $\mu_{N,k} = N^{-2^{k-1}-1/2} \ln^{2^k+1}(N)$ for the NMPL estimator with NR, default NR, and line search NR. We omit the proof because it is very similar to the proof of Lemmas 2-4.

⁶We may reduce the number of policy iterations from 8 to 7 by using the pseudo-covariance matrix estimator (18) instead of (19).

6 Practical Extensions

6.1 Bootstrapping Counterfactual Experiments

One important advantage of structural models over reduced-form models is that we can use them to quantitatively assess the dynamic impact of public policy proposals, often called counterfactual experiments. Thereby, the reliability of the estimated impact of policies arises as an important issue. Our proposed bootstrap method allows us to construct reliable CIs for the dynamic impact of counterfactual policies where asymptotic CIs may be unreliable.

Counterfactual policies are characterized by a counterfactual parameter which in turn depends on the true parameter. Given the true parameter θ , a counterfactual parameter is denoted by $\vartheta(\theta)$, where $\vartheta(\cdot)$ is a (non-random) smooth mapping from Θ to itself. The quantity of interest under a counterfactual policy often depends on the true parameter θ , a counterfactual parameter $\vartheta(\theta)$, as well as the conditional choice probabilities P_θ and $P_{\vartheta(\theta)}$; see the examples provided in Section 7. We assume that the quantity of interest takes a scalar value and denote it by $y(\theta) = g(\theta, \vartheta(\theta), P_\theta, P_{\vartheta(\theta)})$. Define $Y(\theta) = \partial y(\theta)/\partial \theta$. In practice, $Y(\theta)$ is evaluated by taking a numerical derivative of $y(\theta)$.

Denote the NFXP estimator by $\hat{\theta}$ and the covariance matrix estimator by $\Sigma_N(\hat{\theta})$. The asymptotic CI for $y(\hat{\theta})$ of confidence level $100(1 - \alpha)$ is $CI_{ASY} = [y(\hat{\theta}) - z_{\alpha/2}\hat{\sigma}_y/N^{1/2}, y(\hat{\theta}) + z_{\alpha/2}\hat{\sigma}_y/N^{1/2}]$, where $\hat{\sigma}_y^2 = Y(\hat{\theta})'\Sigma_N(\hat{\theta})Y(\hat{\theta})$ and z_α denotes the $1 - \alpha$ quantiles of the standard normal random variable. It is also straightforward to define the bootstrap CIs for $y(\hat{\theta})$. Define the bootstrap t -statistic as $T_y = N^{1/2}(y(\theta^*) - y(\hat{\theta}))/\sigma_y^*$, where $\sigma_y^{*2} = Y(\theta^*)'\Sigma_N^*(\theta^*)Y(\theta^*)$ and θ^* is the bootstrap NFXP estimator. Let $z_{T_y, \alpha}^*$ and $z_{|T_y|, \alpha}^*$ denote the $1 - \alpha$ quantiles of T_y and $|T_y|$. The symmetric and equal-tailed two-sided bootstrap CI for $y(\hat{\theta})$ of confidence level $100(1 - \alpha)$ are defined as $CI_{SYM}(y(\hat{\theta})) = [y(\hat{\theta}) - z_{|T_y|, \alpha}^*\hat{\sigma}_y/N^{1/2}, y(\hat{\theta}) + z_{|T_y|, \alpha}^*\hat{\sigma}_y/N^{1/2}]$ and $CI_{ET}(y(\hat{\theta})) = [y(\hat{\theta}) - z_{T_y, \alpha/2}^*\hat{\sigma}_y/N^{1/2}, y(\hat{\theta}) - z_{T_y, 1-\alpha/2}^*\hat{\sigma}_y/N^{1/2}]$, respectively.

Define $\theta_k^{*s} = (\alpha_k^{*s'}, \hat{\theta}_f')'$, where $s \in \{PL, MPL\}$. The one-step NPL or NMPL bootstrap CIs, denoted by $CI_{SYM, k}(y(\hat{\theta}))$ and $CI_{ET, k}(y(\hat{\theta}))$, are defined exactly as $CI_{SYM}(y(\hat{\theta}))$ and $CI_{ET}(y(\hat{\theta}))$ but with $(\theta^*, \Sigma_N^*(\theta^*))$ replaced with $(\theta_k^{*s}, \Sigma_N^*(P_k^{*s}, \theta_k^{*s}))$, where $s \in \{PL, MPL\}$ and $\Sigma_N^*(P, \theta)$ is defined by (18)-(19).

When $y(\theta)$ depends on $P_{\vartheta(\theta)}$, constructing the one-step bootstrap CIs often requires computing the numerical derivatives of $P_{\vartheta(\theta_k^{*s})}$ with respect to θ_k^{*s} . This is potentially expensive

because it requires solving the fixed point problem, $P = \Psi(P, \vartheta(\theta_k^{*s}))$, as many times as the number of bootstraps multiplied by the dimension of θ .⁷ Let θ_k^* denote either θ_k^{*PL} or θ_k^{*MPL} . We propose to reduce the computational burden in computing $y(\theta)$ by approximating the fixed point $P_{\vartheta(\theta_k^*)}$ by taking a finite number of policy iterations under $\vartheta(\theta_k^*)$ starting from the fixed point under $\hat{\theta}$. That is, starting from $P_{\vartheta,k}^{*0} = P_{\vartheta(\hat{\theta})}$, we repeat policy iterations under $\vartheta(\theta_k^*)$ as $P_{\vartheta,k}^{*j} = \Psi(P_{\vartheta,k}^{*j-1}, \vartheta(\theta_k^*))$ to obtain a sequence $\{P_{\vartheta,k}^{*j} : j \geq 0\}$. Since $P_{\vartheta,k}^{*0} - P_{\vartheta(\theta_k^*)} = O_p(N^{-1/2})$ and the policy iteration mapping $\Psi(\cdot, \vartheta(\theta_k^*))$ has the quadratic convergence property, we have $P_{\vartheta,k}^{*j} - P_{\vartheta(\theta_k^*)} = O_p(N^{-2^{j-1}})$. Under the assumption that $g(\theta, \vartheta(\theta), P_\theta, P_{\vartheta(\theta)})$ is a smooth functional of $P_{\vartheta(\theta)}$, it follows that $g(\theta_k^*, \vartheta(\theta_k^*), P_k^*, P_{\vartheta(\theta_k^*)}) - g(\theta_k^*, \vartheta(\theta_k^*), P_k^*, P_{\vartheta,k}^{*j}) = O_p(N^{-2^{j-1}})$. This suggests that a small value of j may suffice to achieve higher-order refinement in bootstrapping. Let $CI_{SYM,k}^j(y(\hat{\theta}))$ and $CI_{ET,k}^j(y(\hat{\theta}))$ be the approximated one-step bootstrap CIs that use the approximated conditional choice probabilities $P_{\vartheta,k}^{*j}$ in place of $P_{\vartheta(\theta_k^*)}$. Define $\mu_N^j = N^{-2^{j-1}} \ln^{2^j}(N)$ and $\mu_{N,k}^j = \max\{\mu_{N,k}, \mu_N^j\}$. The following Lemma shows choosing $j = k = 3$ ($j = k = 2$) suffices to achieve higher-order refinement in constructing the symmetric (equal-tailed) two-sided bootstrap CIs for $y(\hat{\theta})$.

Lemma 5 *Suppose the assumptions of Lemma 2 hold, $\vartheta(\theta)$ and $g(\theta, \vartheta, P_\theta, P_\vartheta)$ are continuously F -differentiable, and $\sup_\theta \|(\partial/\partial\theta)\vartheta(\theta)\|$, $\sup_{(\theta, \vartheta, P_\theta, P_\vartheta)} \|Dg(\theta, \vartheta, P_\theta, P_\vartheta)\| < \infty$ with probability one. Then*

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

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

6.2 Unobserved Heterogeneity

In the model of Section 2, it is assumed that individuals are homogenous in terms of the parameter θ representing their preferences and transition probabilities. However, in many empirical applications, preferences and transition probabilities are likely to be different across individuals.

⁷Note that numerically evaluating the derivative of $g(\theta_k^{*s}, \vartheta(\theta_k^{*s}), P_{\theta_k^{*s}}, P_{\vartheta(\theta_k^{*s})})$ with respect to θ_k^{*s} requires changing the value of an element of θ_k^{*s} slightly, computing $P_{\vartheta(\cdot)}$ for the new value θ_k^{*s} by solving the fixed point problem, and repeating it elementwise for all elements of θ_k^{*s} .

An approach often used in practice is to treat such heterogeneity as unobserved by econometricians and to allow for a finite mixture of types (c.f., Keane and Wolpin, 1997). This section discusses an extension of our bootstrap method to a finite mixture model.

Suppose there are M types of individuals, where type m is characterized by a type-specific parameter $\theta^m = (\alpha^{m'}, \theta_f^{m'})'$ and the probability of being type m in the population is π^m ($m = 1, \dots, M$).⁸ It is assumed that the number of types, M , is known and $\pi^m \in (0, 1)$. As often done in practice, we reparametrize the type probabilities as $\pi^m(\gamma) = \exp(\gamma^m)/(1 + \sum_{m=1}^{M-1} \exp(\gamma^i))$ for $m = 1, \dots, M-1$ and $\pi^M(\gamma) = 1 - \sum_{m=1}^{M-1} \pi^m(\gamma)$, where $\gamma = (\gamma^1, \dots, \gamma^{M-1})'$.

Let $\zeta = (\gamma', \theta^{1'}, \dots, \theta^{M'})'$ be the parameter to be estimated, and let Θ_ζ denote the set of possible values of ζ . Let $\{\{a_{it}, x_{it}, x_{i,t+1}\}_{t=1}^T\}_{i=1}^N$ be a panel data such that $w_i = \{a_{it}, x_{it}, x_{i,t+1}\}_{t=1}^T$ is randomly drawn across i 's from the population. In particular, the initial state x_{i1} is assumed to be randomly drawn from a type-specific stationary distribution implied by the conditional choice probability and the transition probability. We consider the asymptotics when T is fixed and $N \rightarrow \infty$.

Conditional on being type m , the likelihood of observing w_i is

$$L(w_i; \theta^m) = \lambda(x_{i1}; P_{\theta^m}, f_{\theta_f^m}) \prod_{t=1}^T f_{\theta_f^m}(x_{i,t+1} | x_{it}, a_{it}) P_{\theta^m}(a_{it} | x_{it}), \quad (20)$$

$$\lambda(x; P_{\theta^m}, f_{\theta_f^m}) = \int \sum_{a'=1}^J P_{\theta^m}(a' | x') f_{\theta_f^m}(x | x', a') d\lambda(x'; P_{\theta^m}, f_{\theta_f^m}), \quad (21)$$

where P_{θ^m} is the fixed point of $\Psi(\cdot, \theta^m)$. $\lambda(x; P_{\theta^m}, f_{\theta_f^m})$ is the stationary distribution of x for type m defined as the fixed point of the mapping defined by (21), and it is used to evaluate the (type-specific) likelihood contribution of the initial observation x_{i1} . Since solving (21) given $(P_{\theta^m}, f_{\theta_f^m})$ is often less computationally intensive than computing P_{θ^m} , we assume the full solution of (21) is available given $(P_{\theta^m}, f_{\theta_f^m})$.

The NFXP estimator of ζ is defined as

$$\hat{\zeta} = \arg \max_{\zeta \in \Theta_\zeta} \frac{1}{N} \sum_{i=1}^N l(w_i; \zeta), \quad \text{where } l(w_i; \zeta) = \ln \left(\sum_{m=1}^M \pi^m(\gamma) L(w_i; \theta^m) \right). \quad (22)$$

Let P^m be the conditional choice probability for type m . Stack P^m 's as $\mathbf{P} = (P^1, \dots, P^M)$, and let \mathbf{P}^0 denote its true value. Define $\Psi(\mathbf{P}, \zeta) = (\Psi(P^1, \theta^1), \dots, \Psi(P^M, \theta^M))$ and $\Psi_2(\mathbf{P}, \zeta) =$

⁸If the transition probabilities are common across types so that $\theta_f^m = \theta_f$ for $m = 1, \dots, M$, then we may use the 2-stage procedure analogous to that of Section 3.

$(\Psi_2(P^1, \theta^1), \dots, \Psi_2(P^M, \theta^M))$. The pseudo-likelihood function for the NPL estimator is

$$\mathcal{L}_N^{PL}(\mathbf{P}, \zeta) = \frac{1}{N} \sum_{i=1}^N l^{PL}(w_i; \mathbf{P}, \zeta), \quad \text{where } l^{PL}(w_i; \mathbf{P}, \zeta) = \ln \left(\sum_{m=1}^M \pi^m(\gamma) L^{PL}(w_i; P^m, \theta^m) \right),$$

and

$$L^{PL}(w_i; P^m, \theta^m) = \lambda \left(x_{i1}; \Psi(P^m, \theta^m), f_{\theta_f^m} \right) \prod_{t=1}^T f_{\theta_f^m}(x_{i,t+1} | x_{it}, a_{it}) \Psi(P^m, \theta^m)(a_{it} | x_{it}),$$

where λ is given by the fixed point of the mapping defined by (21). The pseudo-likelihood function for the NMPL estimator is defined by $\mathcal{L}_N^{MPL}(\mathbf{P}, \zeta) = N^{-1} \sum_{i=1}^N l^{MPL}(w_i; \mathbf{P}, \zeta)$, where $l^{MPL}(w_i; \mathbf{P}, \zeta) = l^{PL}(w_i; \Psi(\mathbf{P}, \zeta), \zeta)$, i.e., we replace P^m in the NPL pseudo-likelihood function $\mathcal{L}_N^{PL}(\mathbf{P}, \zeta)$ with $\Psi(P^m, \theta^m)$. Let $L^{MPL}(w_i; P^m, \theta^m) = L^{PL}(w_i; \Psi(P^m, \theta^m), \theta^m)$.

Let $\{\pi^{0,m}\}_{m=1}^M$ be the true set of type probabilities, and let $\{P^{0,m}, f^{0,m}\}_{m=1}^M$ be the true sets of type-specific conditional choice probabilities and transition probabilities. Let $P^0(w)$ denote the true set of probabilities for w defined as $P^0(w) \equiv \sum_{m=1}^M \pi^{0,m} \lambda(x_1; P^{0,m}, f^{0,m}) \times \prod_{t=1}^T f^{0,m}(x_{t+1} | x_t, a_t) P^{0,m}(a_t | x_t)$. Let $\hat{\mathbf{P}}_0^{PL}$ and $\hat{\mathbf{P}}_0^{MPL}$ be initial consistent estimators of \mathbf{P} . Consider the following regularity conditions that correspond to Assumptions 4 and 5.

Assumption 4UH. (a) Θ_ζ is compact. (b) $\lambda^m(x; P, f)$ is three times continuously F-differentiable.

(c) $\lambda(x; P, f_{\theta_f}) > 0$ for any $x \in X$ and any $\{P, \theta_f\} \in B_P \times \Theta_f$. (d) $w_i = \{(a_{it}, x_{it}, x_{i,t+1}) : t = 1, \dots, T\}$ for $i = 1, \dots, N$, are independently and identically distributed, and $dF(x) > 0$ for any $x \in X$, where $F(x)$ is the distribution function of x_i . (e) For any $\{P^m, \theta_f^m\} \in B_P \times \Theta_f$, there exists a unique solution to the fixed point problem of (21). (f) There is a unique $\zeta^0 \in \text{int}(\Theta_\zeta)$ such that, for any $w = \{(a_t, x_t, x_{t+1}) : t = 1, \dots, T\}$, $\sum_{m=1}^M \pi^m(\gamma^0) L(w; \theta^{0,m}) = P^0(w)$. For any $\zeta \neq \zeta^0$, $\Pr_{\zeta^0}(\{w : \sum_{m=1}^M \pi^m(\gamma) L^s(w; P^{0,m}, \theta^m) \neq P^0(w)\}) > 0$ for $s \in \{PL, MPL\}$. (g) $E_{\zeta^0} \sup_{(P,f)} \|D^s \lambda(x; P, f)\|^2 < \infty$ for $s = 0, \dots, 4$. (h) $\hat{\mathbf{P}}_0^{PL} - \mathbf{P}^0 = o_p(1)$, $\hat{\mathbf{P}}_0^{MPL} - \mathbf{P}^0 = o_p(1)$, and the NFXP estimator $\hat{\zeta}$ satisfies $\sqrt{N}(\hat{\zeta} - \zeta^0) \rightarrow_d N(0, \Omega_\zeta)$.

The following Lemma corresponds to Proposition 1 and equation (9) and establishes the key property of the pseudo-likelihood functions of the NPL and NMPL algorithm in the context of a finite mixture model. Define $\mathbf{P}_\zeta = (P_{\theta^1}, \dots, P_{\theta^M})$.

Lemma 6 *Suppose Assumptions 1-3 hold and $\Psi(\cdot)$ and $\lambda(\cdot; \cdot, \cdot)$ are F-differentiable. Then $D_{\mathbf{P}} l^{PL}(w_i; \mathbf{P}_\zeta, \zeta) = D_{\mathbf{P}} l^{MPL}(w_i; \mathbf{P}_\zeta, \zeta) = 0$. Suppose, in addition, Assumption 4(a)-(c), 4(e)-(g) and 4UH hold. Then $D_{\mathbf{P}_\zeta} \mathcal{L}_N^{PL}(\mathbf{P}_\zeta, \hat{\zeta}) = O_p(N^{-1/2})$ and $D_{\mathbf{P}_\zeta} \mathcal{L}_N^{MPL}(\mathbf{P}_\zeta, \hat{\zeta}) = 0$.*

Thus, at the fixed point, the parameter of interest ζ and the nuisance parameter \mathbf{P} are *asymptotically orthogonal* for the NPL estimator and are *orthogonal in any sample size* for the NMPL estimator. Given this result, we may develop the NPL and NMPL algorithms for a finite mixture model which have similar convergence properties to those in section 3.

The NPL and NMPL estimators are defined as follows. Let $s \in \{PL, MPL\}$.

Step 1: Given $\hat{\mathbf{P}}_{j-1}^s$, $\hat{\zeta}_j^s$ is computed by

$$\hat{\zeta}_j^{PL} = \arg \max_{\zeta \in \Theta_\zeta} \mathcal{L}_N^{PL}(\hat{\mathbf{P}}_{j-1}^{PL}, \zeta) \quad \text{or} \quad \hat{\zeta}_j^{MPL} = \arg \max_{\zeta \in \Theta_\zeta} \mathcal{L}_N^{MPL}(\hat{\mathbf{P}}_{j-1}^{MPL}, \zeta). \quad (23)$$

Step 2: For $m = 1, \dots, M$, update $\hat{P}_{j-1}^{s,m}$ using the obtained estimate $\hat{\theta}_j^{s,m}$ as $\hat{P}_j^{s,m} = \Psi(\hat{P}_{j-1}^{s,m}, \hat{\theta}_j^{s,m})$.

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

The following proposition corresponds to Propositions 2 and 3 and establishes the convergence rates of the NPL and the NMPL estimators for a finite mixture model. Define $\hat{\mathbf{P}} = \mathbf{P}_{\hat{\zeta}}$, the NFXP estimator of \mathbf{P} .

Proposition 7 *Suppose Assumptions 1-3, 4(a)-(c), 4(e)-(g), 5, and 4UH hold. Then, for $k = 1, 2, \dots$*

$$\begin{aligned} \hat{\zeta}_k^{PL} - \hat{\zeta} &= O_p(N^{-1/2} \|\hat{\mathbf{P}}_{k-1}^{PL} - \hat{\mathbf{P}}\| + \|\hat{\mathbf{P}}_{k-1}^{PL} - \hat{\mathbf{P}}\|^2), \quad \hat{\mathbf{P}}_k^{PL} - \hat{\mathbf{P}} = O_p(\|\hat{\zeta}_k^{PL} - \hat{\zeta}\|), \\ \hat{\zeta}_k^{MPL} - \hat{\zeta} &= O_p(N^{-1/2} \|\hat{\mathbf{P}}_{k-1}^{MPL} - \hat{\mathbf{P}}\|^2 + \|\hat{\mathbf{P}}_{k-1}^{MPL} - \hat{\mathbf{P}}\|^3), \quad \hat{\mathbf{P}}_k^{MPL} - \hat{\mathbf{P}} = O_p(\|\hat{\mathbf{P}}_{k-1}^{MPL} - \hat{\mathbf{P}}\|^2). \end{aligned}$$

The one-step NPL and NMPL estimators are analogously defined to the NPL and NMPL estimators except that they update the parameter ζ using one Newton step without fully solving the pseudo-maximization problem (23). Specifically, the one-step NPL estimator is updated as

$$\tilde{\zeta}_j^{PL} = \tilde{\zeta}_{j-1}^{PL} - Q_N^{PL}(\tilde{\mathbf{P}}_{j-1}^{PL}, \tilde{\zeta}_{j-1}^{PL})^{-1}(\partial/\partial\zeta)\mathcal{L}_N^{PL}(\tilde{\mathbf{P}}_{j-1}^{PL}, \tilde{\zeta}_{j-1}^{PL}).$$

Then, $\tilde{\mathbf{P}}_{j-1}^{PL}$ is updated as $\tilde{P}_j^{s,m} = \Psi(\tilde{P}_{j-1}^{s,m}, \tilde{\theta}_j^{s,m})$ for $m = 1, \dots, M$. This process is iterated for $j = 1, \dots, k$. The NR choice of Q_N^{PL} is $Q_N^{PL}(\mathbf{P}, \zeta) = (\partial^2/\partial\zeta\partial\zeta')\mathcal{L}_N^{PL}(\mathbf{P}, \zeta)$ whereas the OPG estimator is $Q_N^{PL}(\mathbf{P}, \zeta) = -N^{-1} \sum_{i=1}^N (\partial/\partial\zeta)l^{PL}(w_i; \mathbf{P}, \zeta)(\partial/\partial\zeta')l^{PL}(w_i; \mathbf{P}, \zeta)$. The one-step NMPL estimator is defined analogously.

The following proposition corresponds to Propositions 5 and 6 and shows that the one-step NPL/NMPL estimator achieves a similar rate of convergence as the original NPL/NMPL

estimator for a finite mixture model. The proof is omitted because it follows the proof of Propositions 5 and 6.

Proposition 8 *Suppose the assumptions of Proposition 7 hold and the initial estimates $(\tilde{\zeta}_0^{PL}, \tilde{\mathbf{P}}_0^{PL})$ and $(\tilde{\zeta}_0^{MPL}, \tilde{\mathbf{P}}_0^{MPL})$ are consistent. Then, for $k = 1, 2, \dots$,*

$$\begin{aligned}\tilde{\zeta}_k^{PL} - \hat{\zeta} &= O_p(\|\tilde{\zeta}_{k-1}^{PL} - \hat{\zeta}\|^2 + N^{-1/2}\|\tilde{\mathbf{P}}_{k-1}^{PL} - \hat{\mathbf{P}}\| + \|\tilde{\mathbf{P}}_{k-1}^{PL} - \hat{\mathbf{P}}\|^2) \\ &\quad [+ O_p(N^{-1/2}\|\tilde{\zeta}_{k-1}^{PL} - \hat{\zeta}\|) \text{ for OPG}], \\ \tilde{\mathbf{P}}_k^{PL} - \hat{\mathbf{P}} &= O_p(\|\tilde{\zeta}_k^{PL} - \hat{\zeta}\|). \\ \tilde{\zeta}_k^{MPL} - \hat{\zeta} &= O_p(\|\tilde{\zeta}_{k-1}^{MPL} - \hat{\zeta}\|^2 + N^{-1/2}\|\tilde{\mathbf{P}}_{k-1}^{MPL} - \hat{\mathbf{P}}\|^2 + \|\tilde{\mathbf{P}}_{k-1}^{MPL} - \hat{\mathbf{P}}\|^3) \\ &\quad [+ O_p(N^{-1/2}\|\tilde{\zeta}_{k-1}^{MPL} - \hat{\zeta}\| + \|\tilde{\mathbf{P}}_{k-1}^{MPL} - \hat{\mathbf{P}}\|^2) \text{ for OPG}], \\ \tilde{\mathbf{P}}_k^{MPL} - \hat{\mathbf{P}} &= O_p(\|\tilde{\zeta}_k^{MPL} - \hat{\zeta}\| + \|\tilde{\mathbf{P}}_{k-1}^{MPL} - \hat{\mathbf{P}}\|^2).\end{aligned}$$

The asymptotic covariance matrix of $\hat{\zeta}$ is given by $\Sigma(\zeta^0) = D(\zeta^0)^{-1}V(\zeta^0)(D(\zeta^0)^{-1})'$, where $D(\zeta) = -E(\partial^2/\partial\zeta\partial\zeta')l(w; \zeta)$ and $V(\zeta) = E(\partial/\partial\zeta)l(w; \zeta)(\partial/\partial\zeta')l(w; \zeta)$. As in Section 3.3, we may estimate the asymptotic covariance matrix either using the averages of the derivatives of $l(w_i; \hat{\zeta})$ or the derivatives of the summands of the pseudo-likelihood function.

Applying our bootstrap-based inference method to a finite mixture model is straightforward. We estimate ζ by the NFXP estimator as (22) in the original sample and use $\hat{\zeta}$ and $P_{\hat{\theta}_m}$'s as the initial estimates for the bootstrap samples. The one-step bootstrap NPL and NMPL estimators $(\mathbf{P}_k^{*PL}, \zeta_k^{*PL}, \mathbf{P}_k^{*MPL}, \zeta_k^{*MPL})$ are defined exactly as $(\tilde{\mathbf{P}}_k^{PL}, \tilde{\zeta}_k^{PL}, \tilde{\mathbf{P}}_k^{MPL}, \tilde{\zeta}_k^{MPL})$ but computing from the bootstrap sample. The bootstrap covariance matrix estimator, $\Sigma_N^{*PL}(\mathbf{P}_k^{*PL}, \zeta_k^{*PL})$ (or $\Sigma_N^{*MPL}(\mathbf{P}_k^{*MPL}, \zeta_k^{*MPL})$), is defined analogously to the covariance matrix estimator, $\Sigma_N(\hat{\zeta})$, except that we use the bootstrap sample and the corresponding pseudo-likelihood function. The one-step bootstrap t - and Wald statistics, $T_{N,k}^*(\hat{\zeta}_r)$ and $\mathcal{W}_{N,k}^*(\hat{\zeta})$, are then defined as in (15), but with (θ^*, Σ_N^*) replaced by $(\zeta_k^{*PL}, \Sigma_N^{*PL}(\mathbf{P}_k^{*PL}, \zeta_k^{*PL}))$ or $(\zeta_k^{*MPL}, \Sigma_N^{*MPL}(\mathbf{P}_k^{*MPL}, \zeta_k^{*MPL}))$. The one-step bootstrap CIs are defined similarly to (16) and (17).

Before presenting the final lemma, we define some notation. Let $h_\zeta(w_i, \zeta) \in \mathbb{R}^{Lh_\zeta}$ denote the vector containing the unique components of $(\partial/\partial\zeta)l(w; \zeta)$ and $(\partial/\partial\zeta)l(w; \zeta)(\partial/\partial\zeta')l(w; \zeta)$ and their partial derivatives with respect to ζ through order $d = \max\{2c+2, 3\}$. We assume the true parameter ζ^0 lies in a subset $\Theta_{\zeta,0}$ of Θ_ζ . For some $\delta > 0$, let $\Theta_{\zeta,1} = \{\zeta \in \Theta_\zeta : d(\theta, \Theta_{\zeta,0}) < \delta/2\}$ and $\Theta_{\zeta,2} = \{\zeta \in \Theta_\zeta : d(\theta, \Theta_{\zeta,0}) < \delta\}$. The following lemma establishes the higher-order

improvements of the one-step bootstrap NPL and NMPL algorithms for a finite mixture model. The proof follows the proof of Lemmas 1-4 and is therefore omitted.

Lemma 7 *Suppose Assumptions 1-3, 4(a)-(c), 4(e)-(g), 5, and 4UH hold. Suppose Assumptions 6-8 hold with $\theta, \Theta, \Theta_1, \Theta_2, P_\theta(a|x), h(w_i, \theta)$ replaced by $\zeta, \Theta_\zeta, \Theta_{\zeta,1}, \Theta_{\zeta,2}, l(w; \zeta), h_\zeta(w_i, \zeta)$, respectively, for some $c > 0$ with $2c$ an integer. Suppose $\sup_{\theta \in \Theta} \|(\partial/\partial\theta)P_\theta(a|x)\|, \sup_{(P,\theta)} \|D\Psi(P, \theta)(a|x)\|, \sup_{(P,\theta)} \|D^2\Psi(P, \theta)(a|x)\| < \infty$ with probability one. Then the errors in coverage probability of $CI_{SYM,k}(\tilde{\zeta}_r)$ and $CI_{ET,k}(\tilde{\zeta}_r)$ and the errors in rejection probability of the one-step bootstrap Wald test are given by Lemma 4(a)-(c), respectively.*

7 Monte Carlo Experiments

This section compares the performance of our proposed bootstrap-based inference method with that of the standard inference method based on first-order asymptotics.

7.1 Experimental Design

The model we consider is a version of the machine replacement models of Rust (1987) and Cooper, Haltiwanger, and Power (1999). There are two observable state variables in the model: machine age $s_t \in \mathbb{N}$ and productivity shock $\omega_t \in \mathbb{R}$. We denote the vector of observed state variables by $x_t = (s_t, \omega_t)'$ and let the variable $a_t \in \{0, 1\}$ represent the machine replacement decision. The profit function is given by $u(x_t, a_t) + \epsilon(a_t)$, where

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

with

$$\begin{aligned} rc(a_t) &= \theta_0 a_t, \\ y(s_t, \omega_t, a_t) &= \exp(\theta_1 s_t(1 - a_t) + \omega_t), \\ mc(s_t, a_t) &= \theta_2 s_t(1 - a_t). \end{aligned}$$

Here, $y(s_t, \omega_t, a_t)$ is a revenue function; $c(s_t)$ is a machine maintenance cost; $rc(s_t)$ is a replacement cost; and $\epsilon(a_t)$ is an unobserved state variable which follows an extreme value distribution independently across alternatives. The transition function of s_t is given by $s_t =$

$a_{t-1} + (1 - a_{t-1})(s_{t-1} + 1)$ and productivity shock ω_t follows an AR(1) process $\omega_t = \rho\omega_{t-1} + \eta_t$ with $\eta_t \sim N(0, \sigma_\eta^2)$. The model requires estimation of the three structural parameters whose true value is given by $\theta \equiv (\theta_0, \theta_1, \theta_2)' = (2.0, -0.2, 0.1)'$. We assume that the other parameters in the model, $(\beta, \rho, \sigma_\eta)$, are known and fixed at $(\beta, \rho, \sigma_\eta) = (0.96, 0.8, 0.2)$.

We generate a cross-sectional data set of sample size N from a parametric model by first randomly drawing the initial states $\{(s_i, \omega_i) : i = 1, \dots, N\}$ from the stationary distribution of (s, ω) under θ and then simulating a_i 's using the conditional choice probabilities $P_\theta(a|s_i, \omega_i)$. The data set consists of $\{(s_i, \omega_i, a_i) : i = 1, \dots, N\}$.

To simulate the data from the model with a 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 points outside of the grids. This allows us to generate a sample with continuously distributed ω from the approximated model and to evaluate a likelihood function at points outside of the grids. Finally, we approximate the state space of ω by 10 grid points using Gauss-Hermit quadrature points while the state space of s_t is given by $\{1, \dots, 10\}$.⁹

7.2 Parametric Bootstrapping

We conduct parametric bootstraps with 1000 simulated samples consisting of $N = 1000$ observations. For each simulated sample, we estimate the parameters by Maximum Likelihood (ML) using the NFXP algorithm and draw $B=599$ bootstrap samples from the parametric model evaluated at the ML estimates.¹⁰ Then we estimate parameters for each bootstrap sample using ML, NPL, NMPL, one-step NPL, and one-step NMPL estimators starting from the ML estimates and the corresponding conditional choice probabilities in the original sample. The covariance

⁹The choice of approximation methods can potentially affect estimation and inference. We checked the robustness of the results by repeating the same bootstrapping exercise with the NFXP using alternative approximation methods. First, using 15 instead of 10 grid points in approximating the state space of ω does not substantially change the results. Second, using the method of Tauchen (1986) instead of Gauss-Hermit quadrature method to approximate the state space of ω and their transition probabilities produces similar results.

¹⁰We draw the bootstrap samples of $\{(s_i^*, \omega_i^*) : i = 1, \dots, N\}$ from the stationary distribution under the ML estimate $\hat{\theta}$. We examine the alternative case in which (s_i^*, ω_i^*) is set to the original observation (s_i, ω_i) and find that the results are similar. We also experiment with $B = 999$ in some cases and find that the results do not change substantially.

matrices of the ML estimates are constructed by the OPG estimator using the derivatives of the likelihood function while those of the NPL, NMPL, one-step NPL, and one-step NMPL are constructed by the OPG estimator using the derivatives of their pseudo-likelihood functions.

We first compare the performance of the bootstrap Wald test and the asymptotic Wald test. The null hypothesis we test is $H_0 : (\theta_1, \theta_2) = (-0.2, 0.1)$. Table 1 reports the rejection frequencies of the asymptotic Wald test at .10, .05, and .01 levels for different sample sizes: $N = 500, 1000, \text{ and } 2000$. The asymptotic Wald test overrejects the null hypothesis at all three levels. While the severity of overrejection decreases with the sample size, it is substantial at all levels even with the sample size of 1000.

Table 2 reports the rejection frequencies of the bootstrap Wald test at .10, .05, and .01 levels for ML, NPL, NMPL, one-step NPL, and one-step NMPL estimators for a sample size $N = 1000$. In the table, “1-NPL” and “1-NMPL” represent one-step NPL and one-step NMPL estimators, respectively. The bootstrap Wald tests using ML slightly underreject at .10 and .05 levels but its overall performance is substantially better than that of the asymptotic Wald test. We also conduct the bootstrap Wald test based on the restricted ML estimator where the null is imposed. Its performance is reported in the row “MLE-NULL” and is similar to the one based on the unrestricted ML estimator. The results from the bootstrap Wald tests using NPL and NMPL with one iteration (i.e., $k = 1$) are similar to those using ML and are better than that of asymptotic Wald test at all three levels. Furthermore, the bootstrap Wald tests using one-step NPL and one-step NMPL perform well; 1-NPL and 1-NMPL with five iterations (i.e., $k = 5$) perform better than the asymptotic Wald test at all three levels.

Next, we compare the performance of the bootstrap CIs and the asymptotic CIs for the parameters $\theta_1, \theta_2, \text{ and } \theta_3$. Table 3 reports the coverage performance of the asymptotic 90% and 95% CIs, indicating the frequencies that the confidence intervals missed the true values on the left and right sides. In the case of the 90% CI, for instance, 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% CIs severely overcover on the right while they undercover on the left, suggesting that the center of these CIs is substantially larger than the true parameter value. The *asymmetry* of miscoverage for θ_1 is still substantial even at $N = 2000$. On the other hand, the asymmetry of miscoverage for θ_0 and θ_2 is not as severe as that for θ_1 .¹¹ In terms of the overall coverage

¹¹This may be due to the difference in the degree of nonlinearity. The parameter θ_1 enters into the profit

probabilities, the asymptotic CIs for θ_0 and θ_1 overcover for sizes of 500 and 1000 while the asymptotic CIs for θ_2 undercover.

Table 4 reports the coverage performance of bootstrap 90% and 95% CIs with $N = 1000$. The performance of symmetric bootstrap CIs from ML are similar to that of the asymptotic CIs in Table 3; in particular, both symmetric bootstrap CIs and asymptotic CIs for θ_1 severely overcover on the right while they undercover on the left. On the other hand, equal-tailed bootstrap CIs cover more equally on the right and on the left and thus are better centered around the true parameter value although they slightly undercover overall. The bootstrap CIs from NPL and NMPL with $k = 1$ and the bootstrap CIs from 1-NPL and 1-NMPL with $k = 3$ performs similar to ML.

7.3 Counterfactual Policy Experiments

Our proposed bootstrap method may allow us to construct reliable CIs for the impact of counterfactual policies where asymptotic CIs may be unreliable. We examine the finite sample properties of the bootstrap CIs for the impact of the following counterfactual policy experiments:

1. A government introduces a policy that permanently increases (or decreases) replacement cost by 30 percent. The agents in the economy know that the new policy is permanent.
2. Starting from the steady state, a government unexpectedly introduces a policy that temporarily increases (or decreases) replacement cost by 30 percent for a duration of one period. The agents in the economy know that the new policy only lasts one period.

We focus on the impact of these counterfactual policies on average revenue and revenue dispersion, where the latter is measured as the standard deviation of the logarithm of revenues. In particular, we examine these statistics (i) at the steady state under the new policy in the first experiment and (ii) at the initial period when the new policy is unexpectedly introduced in the second experiment. Table 5 compares the values of these statistics as well as average replacement rate across different experiments.

Given the estimated parameter $\hat{\theta}$, the estimate of a counterfactual parameter is denoted by $\vartheta(\hat{\theta}) = (1.3\hat{\theta}_0, \hat{\theta}_1, \hat{\theta}_2)'$ in the case of a 30% increase or $\vartheta(\hat{\theta}) = (0.7\hat{\theta}_0, \hat{\theta}_1, \hat{\theta}_2)'$ in the case of a 30% decrease.

function through exponential function while θ_0 and θ_1 are linearly related to the profit function; consequently, the degree of nonlinearity in θ_1 is larger than those in θ_0 and θ_2 .

decrease. Average revenue and revenue dispersion for (i)-(ii) above depend on $\hat{\theta}$ and $\vartheta(\hat{\theta})$ as well as the conditional choice probabilities $P_{\hat{\theta}}$ and $P_{\vartheta(\hat{\theta})}$. For instance, average revenue at the steady state under $\vartheta(\hat{\theta})$ may be written as

$$\bar{y}^{(i)}(\hat{\theta}) = \int \sum_{a'=0,1} y_{\vartheta(\hat{\theta})}(s', \omega', a') P_{\vartheta(\hat{\theta})}(a'|s', \omega') d\pi_{\vartheta(\hat{\theta})}^*(s', \omega'),$$

where $\pi_{\vartheta(\hat{\theta})}^*$ is the stationary distribution of (s, ω) under the parameter $\vartheta(\hat{\theta})$ defined as a fixed point of $\pi_{\vartheta(\hat{\theta})}^*(s, \omega) = \int \sum_{a=0,1} P_{\vartheta(\hat{\theta})}(a|s, \omega) f_s(s', \omega'|s, \omega, a) d\pi_{\vartheta(\hat{\theta})}^*(s', \omega')$.

As discussed in Section 6.1, constructing the bootstrap CIs for average revenue and revenue dispersion under counterfactual policies requires repeatedly solving the fixed point problem under counterfactual parameter evaluated at different bootstrap estimates. To construct the bootstrap CIs using NPL, NMPL, one-step NPL, and one-step NMPL, we apply the result of Lemma 5 and approximate the policy function under counterfactual bootstrap estimates by taking 3 policy iterations starting from the fixed point under the counterfactual parameter evaluated at the original estimates, $\vartheta(\hat{\theta})$. On the other hand, for the bootstrap CIs using ML, we use the full solution of the fixed point problem under counterfactual parameter estimates. The asymptotic CIs are constructed by the standard delta method.

Table 6 reports the coverage performance of the asymptotic and the bootstrap 95% CIs for counterfactual average revenues. The asymptotic CIs undercover both on the left and on the right across all counterfactual policies. Both the symmetric and the equal-tailed bootstrap CIs constructed from ML perform slightly better than the asymptotic CIs in terms of coverage probabilities. The average lengths of the asymptotic CIs for average revenues are shorter than those of the bootstrap CIs for all cases (not reported). The bootstrap CIs from NPL and NMPL with one iteration ($k = 1$) perform as well as those from ML while the bootstrap CIs from 1-NPL and 1-NMPL with three iterations ($k = 3$) achieve performance similar to those from ML.

The results are more striking in Table 7, which reports the coverage performance of the asymptotic and the bootstrap 95% CIs for counterfactual revenue dispersions. In terms of coverage probabilities, both the symmetric and the equal-tailed bootstrap CIs constructed from ML perform substantially better than the asymptotic CIs while the symmetric bootstrap CIs perform better than the equal-tailed bootstrap CIs. For instance, for the counterfactual experiment with a permanent 30% decrease in replacement cost, the coverage probabilities of nominal 95% asymptotic, symmetric bootstrap, and equal-tailed bootstrap CIs are .86, .96, and .91, respec-

tively. Furthermore, the asymptotic CIs and the symmetric bootstrap CIs severely overcover on the left and undercover on the right. On the other hand, the equal-tailed bootstrap CIs are better centered around the true parameter values.

The bootstrap CIs from NPL and NMPL with one iteration ($k = 1$) and 1-NPL and 1-NMPL with three iterations ($k = 3$) perform as well as the bootstrap CIs from ML. The results indicate that we can reduce the cost of constructing bootstrap CIs by considering computationally attractive one-step bootstrap procedures, such as one-step NPL and one-step NMPL, instead of the standard bootstrap procedure which is often infeasible in the context of structural discrete Markov decision models.

We acknowledge that the experiment provided in this section has a limited scope and that these results can be different in other applications. Nonetheless, the Monte Carlo evidence suggests that our one-step bootstrap procedure can be used to construct more reliable confidence intervals for the dynamic impact of counterfactual policies where asymptotic confidence intervals may be unreliable and yet the standard bootstrap procedure is too costly to implement.

8 Appendix A: proofs

For an n -linear operator $M(x_1, \dots, x_n)$ such as an n -th F-derivative, the operator norm of M is defined as $\|M\| = \sup_{\|x_1\|=\dots=\|x_n\|=1} \|M(x_1, \dots, x_n)\|$. To simplify the notation, let $\bar{\psi}_\alpha(P, \alpha, \theta_f) = N^{-1} \sum_{i=1}^N (\partial/\partial\alpha') \ln \Psi(P, \alpha, \theta_f)(a_i|x_i)$ and $\bar{\psi}_{2\alpha}(P, \alpha, \theta_f) = N^{-1} \sum_{i=1}^N (\partial/\partial\alpha') \ln \Psi_2(P, \alpha, \theta_f)(a_i|x_i)$.

8.1 Proof of Proposition 1

Let \bar{P} be an arbitrary set of conditional choice probabilities, and let $h = h(a|x)$ be a mapping such that $\bar{P} + h \in B_P$. From the relation $\varphi(\bar{P})(x) = u_{\bar{P}}(x) + \beta E_{\bar{P}}\varphi(\bar{P})(x)$, we obtain

$$\varphi(\bar{P} + h)(x) - \varphi(\bar{P})(x) = (1 - \beta E_{\bar{P}+h})^{-1} [u_{\bar{P}+h}(x) - u_{\bar{P}}(x) + \beta (E_{\bar{P}+h} - E_{\bar{P}}) \varphi(\bar{P})(x)].$$

Recall $u_{\bar{P}}(x) = \sum_{a \in A} \bar{P}(a|x)u(x, a) + \sum_{a \in A} \bar{P}(a|x)e_x(a, \bar{P}_x)$, and note that $\sum_{a \in A} h(a|x) = 0$ because $\bar{P}, \bar{P} + h \in B_P$. Furthermore, Lemmas 1 and 2 of AM hold uniformly in $x \in X$ by Assumptions 1 and 2. Consequently, applying Lemma 2 of AM to $u_{\bar{P}+h}(x) - u_{\bar{P}}(x)$ gives $u_{\bar{P}+h}(x) - u_{\bar{P}}(x) = \sum_{a \in A} h(a|x)u(x, a) - Q_x^{-1}(\bar{P}_x)' \bar{h}_x + o(\|h\|)$, where $\bar{h}_x = (h(2|x), \dots, h(J|x))'$, $\bar{P}_x = (\bar{P}_x(2), \dots, \bar{P}_x(J))'$, and $o(\|h\|)$ term is uniform in $x \in X$.

Let P be the fixed point of Ψ , so that $\varphi(P)(x) = V(x)$. Then

$$\begin{aligned} \beta(E_{P+h} - E_P)\varphi(P)(x) &= \beta \sum_{a \in A} h(a|x) \int_X V(x') f(dx'|x, a) \\ &= \sum_{a \in A} h(a|x)v(x, a) - \sum_{a \in A} h(a|x)u(x, a) \\ &= \tilde{v}'_x \bar{h}_x - \sum_{a \in A} h(a|x)u(x, a). \end{aligned}$$

Because $\tilde{v}_x = Q_x^{-1}(P_x)$ when P is the fixed point of Ψ , it follows that $\varphi(P+h) - \varphi(P) = o(\|h\|)$ for any h and hence $D\varphi(P) = 0$. Since $\Psi = \Lambda \circ \varphi$, application of the chain rule in B-spaces gives $D\Psi(P) = D\Lambda(\varphi(P))D\varphi(P) = 0$. \square

8.2 Proof of Proposition 2

Because the NFXP estimator maximizes the objective function of the NPL estimator if $P = \hat{P}$ (c.f. equation (Ap.3) of AM p. 1540), it follows that

$$\bar{\psi}_\alpha(\hat{P}_{k-1}^{PL}, \hat{\alpha}_k^{PL}, \hat{\theta}_f) = \bar{\psi}_\alpha(\hat{P}, \hat{\alpha}, \hat{\theta}_f) = 0. \quad (24)$$

We use induction. First, assume $\hat{P}_{k-1}^{PL} - P^0 = o_p(1)$. Then $\hat{\alpha}_k^{PL}$ is consistent, because the consistency proof in the proof of Proposition 4 of AM does not depend on the finiteness of X . Applying the generalized Taylor's theorem [c.f., pp.148-149 of Zeidler (1986)] to $\bar{\psi}_\alpha(\hat{P}_{k-1}^{PL}, \hat{\alpha}_k^{PL}, \hat{\theta}_f) - \bar{\psi}_\alpha(\hat{P}, \hat{\alpha}, \hat{\theta}_f)$ gives

$$\int_0^1 (\partial/\partial\alpha)\bar{\psi}_\alpha(P_\tau, \alpha_\tau, \hat{\theta}_f)(\hat{\alpha}_k^{PL} - \hat{\alpha})d\tau + \int_0^1 D_P\bar{\psi}_\alpha(P_\tau, \alpha_\tau, \hat{\theta}_f)(\hat{P}_{k-1}^{PL} - \hat{P})d\tau = 0 \quad (25)$$

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 $\hat{P} - P^0 = P_{\hat{\theta}} - P_{\theta^0} = O_p(N^{-1/2})$ because $\hat{\theta} - \theta^0 = O_p(N^{-1/2})$ and $\partial P_\theta/\partial\theta = \partial\Psi(P_\theta, \theta)/\partial\theta = O_p(1)$ from Lemma 8(a). For the first term on the left of (25), $\int_0^1 (\partial/\partial\alpha)\bar{\psi}_\alpha(P_\tau, \alpha_\tau, \hat{\theta}_f)d\tau \rightarrow_p E(\partial^2/\partial\alpha\partial\alpha') \ln \Psi(P^0, \theta^0)$ follows from Lemma 8(d) and the consistency of $P_\tau, \hat{\theta}_f$, and α_τ . For the second term on the left of (25), expanding $D_P\bar{\psi}_\alpha(P_\tau, \alpha_\tau, \hat{\theta}_f)$ around $(\hat{P}, \hat{\alpha}, \hat{\theta}_f)$ and using $\|P_\tau - \hat{P}\| \leq \|\hat{P}_{k-1}^{PL} - \hat{P}\|$, $\|\alpha_\tau - \hat{\alpha}\| \leq \|\hat{\alpha}_k^{PL} - \hat{\alpha}\|$, Lemma 8(b)(c), and root- N consistency of $(\hat{\alpha}, \hat{\theta}_f, \hat{P})$, we obtain

$$D_P\bar{\psi}_\alpha(P_\tau, \alpha_\tau, \hat{\theta}_f) = O_p(N^{-1/2}) + O_p(\|\hat{P}_{k-1}^{PL} - \hat{P}\|) + O_p(\|\hat{\alpha}_k^{PL} - \hat{\alpha}\|),$$

uniformly in τ . Therefore, rearranging the terms in (25) gives

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

and $\hat{\alpha}_k^{PL} - \hat{\alpha} = O_p(N^{-1/2} \|\hat{P}_{k-1}^{PL} - \hat{P}\| + \|\hat{P}_{k-1}^{PL} - \hat{P}\|^2)$ follows because $E(\partial^2/\partial\alpha\partial\alpha') \ln \Psi(P^0, \theta^0)$ is a nonsingular negative definite matrix (see AM p.1541).

For the convergence rate of \hat{P}_k^{PL} , expand $\hat{P}_k^{PL} = \Psi(\hat{P}_{k-1}^{PL}, \hat{\alpha}_k^{PL}, \hat{\theta}_f)$ around $(\hat{P}, \hat{\alpha}, \hat{\theta}_f)$, apply $\Psi(\hat{P}, \hat{\alpha}, \hat{\theta}_f) = \hat{P}$ and $D_P \Psi(\hat{P}, \hat{\alpha}, \hat{\theta}_f) = 0$, and use Lemma 8(a), to obtain

$$\hat{P}_k^{PL} = \Psi(\hat{P}_{k-1}^{PL}, \hat{\alpha}_k^{PL}, \hat{\theta}_f) = \hat{P} + O_p(\|\hat{\alpha}_k^{PL} - \hat{\alpha}\|) + O_p(\|\hat{P}_{k-1}^{PL} - \hat{P}\|^2). \quad (26)$$

The required result for all k follows from induction because $\hat{P}_0^{PL} - P^0 = o_p(1)$ by Assumption 4(g). \square

8.3 Proof of Proposition 3

We use induction. Assume $\hat{P}_{k-1}^{MPL} - P^0 = o_p(1)$. The consistency of $\hat{\alpha}_k^{MPL}$ follows from an argument similar to the proof of consistency of $\hat{\alpha}_k^{PL}$ by AM. From the first order conditions for the NMPL and NFXP estimator and Lemma 9(a), we have

$$\bar{\psi}_{2\alpha}(\hat{P}_{k-1}^{MPL}, \hat{\alpha}_k^{MPL}, \hat{\theta}_f) = \bar{\psi}_{2\alpha}(\hat{P}, \hat{\alpha}, \hat{\theta}_f) = 0. \quad (27)$$

Applying the generalized Taylor's theorem to (27) gives

$$\int_0^1 (\partial/\partial\alpha) \bar{\psi}_{2\alpha}(P_\tau, \alpha_\tau, \hat{\theta}_f) (\hat{\alpha}_k^{MPL} - \hat{\alpha}) d\tau + \int_0^1 D_P \bar{\psi}_{2\alpha}(P_\tau, \alpha_\tau, \hat{\theta}_f) (\hat{P}_{k-1}^{MPL} - \hat{P}) d\tau = 0, \quad (28)$$

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 (28), $\int_0^1 (\partial/\partial\alpha) \bar{\psi}_{2\alpha}(P_\tau, \alpha_\tau, \hat{\theta}_f) d\tau \rightarrow_p E(\partial^2/\partial\alpha\partial\alpha') \ln \Psi_2(P^0, \alpha^0, \theta_f^0) = E(\partial^2/\partial\alpha\partial\alpha') \ln P_{\theta^0}$ from Lemma 8(d) and the consistency of P_τ , α_τ , and $\hat{\theta}_f$. For the second term on the left of (28), recall $D_P \bar{\psi}_{2\alpha}(\hat{P}, \hat{\alpha}, \hat{\theta}_f) = 0$ from Lemma 9(a) because \hat{P} is the fixed point of $\Psi(\cdot, \hat{\alpha}, \hat{\theta}_f)$. Thus, applying the generalized Taylor's theorem to $D_P \bar{\psi}_{2\alpha}(P_\tau, \alpha_\tau, \hat{\theta}_f) - D_P \bar{\psi}_{2\alpha}(\hat{P}, \hat{\alpha}, \hat{\theta}_f)$ yields

$$D_P \bar{\psi}_{2\alpha}(P_\tau, \alpha_\tau, \hat{\theta}_f) = \int_0^1 D_{PP} \bar{\psi}_{2\alpha}(P_b, \alpha_b, \hat{\theta}_f) (P_\tau - \hat{P}) db + \int_0^1 D_{\alpha P} \bar{\psi}_{2\alpha}(P_b, \alpha_b, \hat{\theta}_f) (\alpha_\tau - \hat{\alpha}) db, \quad (29)$$

where $P_b = bP_\tau + (1 - b)\hat{P}$ and $\alpha_b = b\alpha_\tau + (1 - b)\hat{\alpha}$. For the right hand side of (29), first note that $D_{PP} \bar{\psi}_{2\alpha}(P^0, \alpha^0, \theta_f^0)$ and $D_{\alpha P} \bar{\psi}_{2\alpha}(P^0, \alpha^0, \theta_f^0)$ are $O_p(N^{-1/2})$ from Lemma 9(b) and $w_i \sim$ iid. Consequently, we obtain, uniformly in b ,

$$D_{PP} \bar{\psi}_{2\alpha}(P_b, \alpha_b, \hat{\theta}_f), D_{\alpha P} \bar{\psi}_{2\alpha}(P_b, \alpha_b, \hat{\theta}_f) = O_p(N^{-1/2} + \|\alpha_\tau - \hat{\alpha}\| + \|P_\tau - \hat{P}\|), \quad (30)$$

by expanding the left hand side around $(P^0, \alpha^0, \theta_f^0)$, applying the triangle inequality to $\|P_b - P^0\|$ and $\|\alpha_b - \alpha^0\|$, and using Lemma 8(b) and the root- N consistency of $(\hat{\alpha}, \hat{P}, \hat{\theta}_f)$. Substituting (30) into (29) gives, uniformly in τ ,

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

Consequently, rearranging the terms in (28) gives $[E(\partial^2/\partial\alpha\partial\alpha') \ln P_{\theta^0} + o_p(1)](\hat{\alpha}_k^{MPL} - \hat{\alpha}) = O_p(N^{-1/2} \|\hat{P}_{k-1}^{MPL} - \hat{P}\|^2 + \|\hat{P}_{k-1}^{MPL} - \hat{P}\|^3)$, and the stated bound on $\hat{\alpha}_k^{MPL} - \hat{\alpha}$ follows because $E(\partial^2/\partial\alpha\partial\alpha') \ln P_{\theta^0}$ is a nonsingular negative definite matrix.

For \hat{P}_k^{MPL} , we have $\hat{P}_k^{MPL} = \Psi(\hat{P}_{k-1}^{MPL}, \hat{\alpha}_k^{MPL}, \hat{\theta}_f) = \hat{P} + O_p(\|\hat{\alpha}_k^{MPL} - \hat{\alpha}\|) + O_p(\|\hat{P}_{k-1}^{MPL} - \hat{P}\|^2)$ from the same argument as (26). The required result for all k follows from induction because $\hat{P}_0^{MPL} - P^0 = o_p(1)$ by Assumption 5(c). \square

8.4 Proof of Proposition 4

First, consider a MLE based on $l_3(\theta) = N^{-1} \sum_{i=1}^N [\ln \Psi(P^0, \theta) + \ln f_{\theta_f}]$. The information matrix equality associated with it implies $-E(\partial^2/\partial\alpha\partial\theta'_f) \ln \Psi(P^0, \theta^0) = E(\partial/\partial\alpha) \ln \Psi(P^0, \theta^0) \times (\partial/\partial\theta'_f)(\ln \Psi(P^0, \theta^0) + \ln f_{\theta_f^0})$. Then, the required result for the (1,2)-th block of $D_N^{PL}(\bar{P}, \bar{\theta})$ follows from Lemma 8, the information matrix equality and (47) 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)-th 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 Lemma 9(a) instead of (47). An analogous argument gives the proof for $D_N^{Q,s}(\bar{P}, \bar{\theta})$ and $V_N^s(\bar{P}, \bar{\theta})$. \square

8.5 Proof of Proposition 5

We prove the result for only the NR and OPG methods. The proof for the default NR and line-search NR is essentially the same except for showing $\Pr(Q_N^D \neq Q_N^{NR}) \rightarrow 0$ and $\Pr(Q_N^{LS} \neq Q_N^{NR}) \rightarrow 0$; see the proof of Lemma 7.1 of Andrews (2005) (A05 hereafter). We suppress the

superscript PL from $\tilde{\alpha}_j^{PL}$ and \tilde{P}_j^{PL} , and we suppress $\hat{\theta}_f$ from $\bar{\psi}_\alpha(P, \alpha, \hat{\theta}_f)$ and $Q_N(P, \alpha, \hat{\theta}_f)$ when it does not lead to confusion.

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

$$\begin{aligned}
0 &= \bar{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) + D_\alpha \bar{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})(\hat{\alpha} - \tilde{\alpha}_{j-1}) \\
&\quad + D_P \bar{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})(\hat{P} - \tilde{P}_{j-1}) + R_{N,j} \\
&= \bar{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) + Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})(\hat{\alpha} - \tilde{\alpha}_{j-1}) + Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})(\hat{P} - \tilde{P}_{j-1}) \\
&\quad + \left[D_\alpha \bar{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) - Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) \right] (\hat{\alpha} - \tilde{\alpha}_{j-1}) \\
&\quad + D_P \bar{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})(\hat{P} - \tilde{P}_{j-1}) + R_{N,j}, \tag{31}
\end{aligned}$$

where $R_{N,j} = O_p(\|\hat{P} - \tilde{P}_{j-1}\|^2 + \|\hat{\alpha} - \tilde{\alpha}_{j-1}\|^2)$ from Lemma 8(b). The first two terms on the right of (31) cancel out. For the fourth term on the right of (31), the term inside the bracket is zero in the NR and $O_p(\|\hat{P} - \tilde{P}_{j-1}\| + \|\hat{\alpha} - \tilde{\alpha}_{j-1}\| + N^{-1/2})$ in the OPG from Lemma 8(d), (e) and the information matrix equality. For the fifth term on the right of (31), it follows from the generalized Taylor's theorem, Lemma 8(c), and $\hat{P} - P^0, \hat{\theta} - \theta^0 = O_p(N^{-1/2})$ that

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

Therefore,

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

The stated bound of $\tilde{\alpha}_j - \hat{\alpha}$ follows from $Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) \rightarrow_p E(\partial^2 / \partial \alpha \partial \alpha') \ln \Psi(P^0, \theta^0)$, which is negative definite.

We complete the proof by showing the bound of $\tilde{P}_j - \hat{P}$. Similarly to the proof of Proposition 2, expanding $\tilde{P}_j = \Psi(\tilde{P}_{j-1}, \tilde{\alpha}_j)$ around $(\hat{P}, \hat{\alpha})$ and applying $D_P \Psi(\hat{P}, \hat{\alpha}) = 0$ and Assumption 4(g) gives $\tilde{P}_j = \hat{P} + O_p(\|\tilde{\alpha}_j - \hat{\alpha}\| + \|\tilde{P}_{j-1} - \hat{P}\|^2) = \hat{P} + O_p(\|\tilde{\alpha}_j - \hat{\alpha}\|)$. The required result follows by induction. \square

8.6 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 $\bar{\psi}_{2\alpha}(P, \alpha, \hat{\theta}_f)$ and $Q_N(P, \alpha, \hat{\theta}_f)$.

The proof is similar to the proof of Proposition 5. Since the NFXP estimator satisfies the first order condition $\bar{\psi}_{2\alpha}(\hat{P}, \hat{\alpha}) = 0$, applying the generalized Taylor's theorem to $\bar{\psi}_{2\alpha}(\hat{P}, \hat{\alpha}) - \bar{\psi}_{2\alpha}(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})$ and proceeding similarly to (31) gives

$$\begin{aligned} 0 &= Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})(\hat{\alpha} - \tilde{\alpha}_{j-1}) + \left[D_{\alpha} \bar{\psi}_{2\alpha}(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) - Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) \right] (\hat{\alpha} - \tilde{\alpha}_{j-1}) \\ &\quad + D_P \bar{\psi}_{2\alpha}(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})(\hat{P} - \tilde{P}_{j-1}) + R_{N,j}, \end{aligned} \quad (32)$$

where $\|R_{n,j}\| \leq 2 \sup_{(P,\alpha)} (\|D_{PP} \bar{\psi}_{2\alpha}(P, \alpha)\| + \|D_{\alpha P} \bar{\psi}_{2\alpha}(P, \alpha)\|) (\|\hat{P} - \tilde{P}_{j-1}\|^2 + \|\hat{\alpha} - \tilde{\alpha}_{j-1}\|^2) + \sup_{(P,\alpha)} (\|D_{\alpha\alpha} \bar{\psi}_{2\alpha}(P, \alpha)\|) (\|\hat{\alpha} - \tilde{\alpha}_{j-1}\|^2)$, where the supremum is taken for all the pairs of (P, α) that lie between $(\hat{P}, \hat{\alpha})$ and $(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})$.

For the second term on the right of (32), the term inside the bracket is 0 in the NR and $O_p(\|\hat{P} - \tilde{P}_{j-1}\| + \|\hat{\alpha} - \tilde{\alpha}_{j-1}\| + N^{-1/2})$ in the OPG from Lemma 8(d)(e) and the information matrix equality. For the third term on the right of (32), we obtain

$$D_P \bar{\psi}_{2\alpha}(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) = O_p(N^{-1/2} \|\tilde{\alpha}_{j-1} - \hat{\alpha}\| + N^{-1/2} \|\tilde{P}_{j-1} - \hat{P}\| + \|\tilde{\alpha}_{j-1} - \hat{\alpha}\|^2 + \|\tilde{P}_{j-1} - \hat{P}\|^2) \quad (33)$$

by expanding $D_P \bar{\psi}_{2\alpha}(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})$ around $(\hat{P}, \hat{\alpha})$ and applying $D_P \bar{\psi}_{2\alpha}(\hat{P}, \hat{\alpha}) = 0$ and $D_{PP} \bar{\psi}_{2\alpha}(\hat{P}, \hat{\alpha})$, $D_{\alpha P} \bar{\psi}_{2\alpha}(\hat{P}, \hat{\alpha}) = O_p(N^{-1/2})$, which follows from Lemma 9, the root- N consistency of $(\hat{P}, \hat{\theta})$, and Lemma 8(b). Finally, for the bound of $R_{n,j}$, applying the argument that is used to show (30) gives $\sup_{(P,\alpha)} D_{PP} \bar{\psi}_{2\alpha}(P, \alpha)$, $\sup_{(P,\alpha)} D_{\alpha P} \bar{\psi}_{2\alpha}(P, \alpha) = O_p(N^{-1/2} + \|\tilde{\alpha}_{j-1} - \hat{\alpha}\| + \|\tilde{P}_{j-1} - \hat{P}\|)$ with the range of the supremum stated above. Lemma 8(b) gives $\sup_{(P,\alpha)} \|D_{\alpha\alpha} \bar{\psi}_{2\alpha}(P, \alpha)\| = O_p(1)$.

Combining all the bounds in conjunction with $Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) \rightarrow_p E(\partial^2 / \partial \alpha \partial \alpha') \ln P_{\theta_0}$ gives $\hat{\alpha} - \tilde{\alpha}_j = O_p(\|\tilde{\alpha}_{j-1} - \hat{\alpha}\|^2 + N^{-1/2} \|\tilde{P}_{j-1} - \hat{P}\|^2 + \|\tilde{P}_{j-1} - \hat{P}\|^3 (+ O_p(N^{-1/2} \|\tilde{\alpha}_{j-1} - \hat{\alpha}\| + \|\tilde{P}_{j-1} - \hat{P}\|^2))$ for OPG). The bound of $\tilde{P}_j - \hat{P}$ follows from the same argument as the proof of Proposition 5, and induction gives the required result. \square

8.7 Proof of Lemma 1

The stated result follows from applying the proof of Theorem 6.1 of A05. Note that only Lemmas A.6, A.7, and A.8 of A05 are used in his proof. Our Lemma 11 corresponds to Lemma A.6 of

A05. The results of Lemmas A.7 and A.8 of A05 hold in our case, because we can replace Lemmas A.4 and A.6 of A05 in the proof of Lemmas A.7 and A.8 of A05 with our Lemmas 10 and 11 and the proof carries through. \square

8.8 Proof of Lemma 2

The proof follows the same line of approach as the proof of Theorem 7.1 of Andrews (2005). 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 $k = 0, 1, \dots$ (this corresponds to (A.9) of A05)

$$\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0} (|\tilde{\alpha}_k - \hat{\alpha}| > \mu_{N,k}) = o(N^{-c}), \quad \sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0} (|\tilde{P}_k - \hat{P}| > \mu_{N,k}) = o(N^{-c}) \quad (34)$$

$$\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0} (|T_{N,k}(\theta_r^0) - T_N(\theta_r^0)| > N^{-1/2} \mu_{N,k}) = o(N^{-c}), \quad (35)$$

$$\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0} (|\mathcal{W}_{N,k}(\theta^0) - \mathcal{W}_N(\theta^0)| > N^{-1/2} \mu_{N,k}) = o(N^{-c}). \quad (36)$$

Then, as in the proof of Theorem 7.1 of A05 (p. 203), the stated result follows from applying Lemma A.1 of A05 three times, because the condition on $\hat{\theta}$ (corresponding to $\hat{\theta}_N$ in A05) in Lemma A.1 of A05 is satisfied by our Lemma 10.

First, using an induction argument, we prove the result for the one-step NPL estimator. Let $\mu_{N,k} = N^{-(k+1)/2} \ln^{k+1} N$. For $k = 0$, (34) holds from Lemma 10 and $\sup_{\theta \in \Theta} \|(\partial/\partial\theta)P_\theta\| < \infty$. Suppose (34) holds for $k = j - 1 \geq 0$. Then, from (31) in the proof of Proposition 5, we have

$$\begin{aligned} \tilde{\alpha}_j - \hat{\alpha} &= Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})^{-1} \left[D_\alpha \bar{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) - Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) \right] (\hat{\alpha} - \tilde{\alpha}_{j-1}) \\ &\quad + Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})^{-1} D_P \bar{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})(\hat{P} - \tilde{P}_{j-1}) + Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})^{-1} R_{N,j}, \end{aligned} \quad (37)$$

where $\|R_{N,j}\| \leq (\sup_{(P,\alpha,\theta_f)} \|D^2 \bar{\psi}_\alpha(P, \alpha, \theta_f)\|) (\|\hat{\alpha} - \tilde{\alpha}_{j-1}\|^2 + \|\hat{P} - \tilde{P}_{j-1}\|^2)$.

From Lemmas A.2(b), A.2(c), and A.3 of A05 and Assumption 7(c), we have, for all $\varepsilon > 0$ and some $K < \infty$,

$$\begin{aligned} \sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0} (\sup_{(P,\alpha,\theta_f)} \|D^2 \bar{\psi}_\alpha(P, \alpha, \theta_f)\| > K) &= o(N^{-c}), \\ \sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0} (\|D_P \bar{\psi}_\alpha(P^0, \alpha^0, \theta_f^0)\| > \varepsilon N^{-1/2} \ln N) &= o(N^{-c}), \\ \sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0} (\|Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})^{-1}\| > K) &= o(N^{-c}). \end{aligned} \quad (38)$$

Thus, expanding $D_P \bar{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) = D_P \bar{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}, \hat{\theta}_f)$ in (37) around $(P^0, \alpha^0, \theta_f^0)$ gives

$$\begin{aligned} \|D_P \bar{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})\| &\leq \xi_{N,j} \left(N^{-1/2} \ln N + \|\tilde{P}_{j-1} - \hat{P}\| + \|\tilde{\alpha}_{j-1} - \hat{\alpha}\| \right) \\ &\quad + \xi_{N,j} \left(\|\hat{P} - P^0\| + \|\hat{\alpha} - \alpha^0\| + \|\hat{\theta}_f - \theta_f^0\| \right), \end{aligned} \quad (39)$$

where $\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0}(\|\xi_{N,j}\| > K) = o(N^{-c})$ for some $K < \infty$.

In case of NR, the first term on the right of (37) is zero. Hence, the first equation of (34) for $k = j$ follows from (37)-(39) and Lemma 10. In case of the default NR, line-search NR, and OPG, we can show

$$\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0}(\|D_\alpha \bar{\psi}_\alpha(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1}) - Q_N(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})\| > N^{-1/2} \ln N) = o(N^{-c}), \quad (40)$$

by repeating the argument of the proof of Lemma 1 of Andrews (2001), to which the proof of Lemma 7.1 of A05 refers. Using (40) to bound the first term on the right of (37), we establish that the first equation of (34) holds for $k = j$.

To show that the second equation of (34) holds for $k = j$, expanding $\Psi(\tilde{P}_{j-1}, \tilde{\alpha}_j)$ around $(\hat{P}, \hat{\alpha})$ and applying $D_P \Psi(\hat{P}, \hat{\alpha}) = 0$ give

$$\|\tilde{P}_j - \hat{P}\| \leq \|D_\alpha \Psi(\hat{P}, \hat{\alpha})\| \|\tilde{\alpha}_j - \hat{\alpha}\| + \left(\sup_{(P, \alpha)} \|D^2 \Psi(P, \alpha, \hat{\theta}_f)\| \right) (\|\tilde{\alpha}_j - \hat{\alpha}\|^2 + \|\tilde{P}_{j-1} - \hat{P}\|^2).$$

Then the required result follows from $\sup_{(P, \theta)} \|D \Psi(P, \theta)\| < \infty$ and $\sup_{(P, \theta)} \|D^2 \Psi(P, \theta)\| < \infty$.

We proceed to prove (35) and (36). Let Σ_r denote $(\Sigma_N(\hat{\theta}))_{rr}$. Also, 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 (10) 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 pp. 205-6 of A05, (35) holds if there exists $K < \infty$ and $\delta > 0$ such that

$$\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0}(|\Sigma_r - \Sigma_{k,r}| > \mu_{N,k}) = o(N^{-c}), \quad (41)$$

$$\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0}(\Sigma_{k,r} < \delta) = o(N^{-c}), \quad \sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0}(\Sigma_r < \delta) = o(N^{-c}). \quad (42)$$

Let $\bar{\theta}$ denote an estimator that satisfies: for all $\varepsilon > 0$, $\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0}(\|\bar{\theta} - \theta^0\| > \varepsilon) = o(N^{-c})$. Then, proceeding in the same way as the proof of Lemma A.3 of A05, we obtain the following; for all $\varepsilon > 0$ and some $K < \infty$, $\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0}(\|V_N(\bar{\theta}) - V(\theta^0)\| > \varepsilon) = o(N^{-c})$, $\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0}(\|D_N(\bar{\theta}) - D(\theta^0)\| > \varepsilon) = o(N^{-c})$, and $\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0}(\|D_N^O(\bar{\theta}) - D(\theta^0)\| > \varepsilon) = o(N^{-c})$. Thus, (42) holds. Equation (41) holds if

$$\begin{aligned} \sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0}(\|V_N^{PL}(\tilde{P}_k, \tilde{\theta}_k) - V_N(\hat{\theta})\| > \mu_{N,k}) &= o(N^{-c}), \quad \text{and} \\ \sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0}(\|D_N^{O,PL}(\tilde{P}_k, \tilde{\theta}_k) - D_N^O(\hat{\theta})\| > \mu_{N,k}) &= o(N^{-c}). \end{aligned}$$

Note that $V_N(\hat{\theta}) = V_N^{PL}(\hat{P}, \hat{\theta})$ from (47). Therefore, the first result follows from applying the generalized Taylor's theorem to $V_N^{PL}(\tilde{P}_k, \tilde{\theta}_k) - V_N^{PL}(\hat{P}, \hat{\theta})$ in conjunction with Lemma A.2(b) of

A05 and (34). The second result is proven in an analogous manner, and we complete the proof of (35). 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 (48). Finally, in view of the argument in p. 206 of A05, (36) follows from (34) and the proof of (35), because Lemma A.8(a) of A05 holds in our case (see the proof of Lemma 1). The proof for the one-step NPL for general $k \geq 1$ follows by induction.

The proof for the one-step NMPL estimator follows an analogous argument. Suppose (34) 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 all other cases. From (32) in the proof of Proposition 6 and the bounds analogous to (38), equation (37) holds for the one-step NMPL estimator with $\bar{\psi}_\alpha$ replaced by $\bar{\psi}_{2\alpha}$, where the reminder term satisfies $\|R_{N,j}\| \leq \xi_{N,j}(\|\hat{\alpha} - \tilde{\alpha}_{j-1}\|^2 + N^{-1/2}\|\hat{P} - \tilde{P}_{j-1}\|^2 + \|\hat{P} - \tilde{P}_{j-1}\|^3)$ with the same definition of $\xi_{N,j}$. Applying the argument used to show (33), we have, in place of (39),

$$\|D_P \bar{\psi}_{2\alpha}(\tilde{P}_{j-1}, \tilde{\alpha}_{j-1})\| \leq \xi_{N,j}(N^{-1/2}\|\hat{\alpha} - \tilde{\alpha}_{j-1}\| + N^{-1/2}\|\hat{P} - \tilde{P}_{j-1}\| + \|\hat{\alpha} - \tilde{\alpha}_{j-1}\|^2 + \|\hat{P} - \tilde{P}_{j-1}\|^2).$$

Therefore, repeating the argument of the proof for the one-step NPL estimator following equation (39) shows that the first equation of (34) holds for $k = j$ with $\mu_{N,k} = N^{-2^{k-1}} \ln^{2^k}(N)$ ($\mu_{N,k} = N^{-(k+1)/2} \ln^{k+1} N$ for the OPG). Note that (40) holds with $D_\alpha \bar{\psi}_\alpha$ replaced by $D_\alpha \bar{\psi}_{2\alpha}$. The second equation of (34) follows from $\sup_{(P,\theta)} \|D\Psi_2(P, \theta)\| < \infty$ and $\sup_{(P,\theta)} \|D^2\Psi_2(P, \theta)\| < \infty$, and the proof for general $k \geq 1$ follows from induction.

Equations (35) and (36) are proven using the same argument as the one for the one-step NPL estimator. The only difference is that $\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0}(\|D_N^{MPL}(\tilde{P}_k, \tilde{\theta}_k) - D_N(\hat{\theta})\| > \mu_{N,k}) = o(N^{-c})$ holds by virtue of Lemma 9. \square

8.9 Proof of Lemma 3 and 4

These lemmas correspond to Theorems 7.1(b) and 7.2 of A05. They are proven by applying the argument of pp. 206-7 of A05. \square

8.10 Proof of Lemma 5

We drop the superscript PL and MPL from $\tilde{\alpha}_k$ and \tilde{P}_k . Define $\tilde{P}_{\vartheta,k}^j$ exactly as $\tilde{P}_{\vartheta,k}^{*j}$ but using the original sample in place of the bootstrap samples. In view of the proof of Lemmas 2-4, the required result follows if we show that, if $\tilde{\alpha}_0 = \alpha^0$, $\tilde{P}_0 = P^0$, and $\tilde{P}_{\vartheta,k}^0 = P_{\vartheta}(\theta^0)$, then for

$k = 0, 1, \dots$ and $j = 0, 1, \dots$

$$\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0} \left(\|\vartheta(\tilde{\theta}_k) - \vartheta(\hat{\theta})\| > \mu_{N,k} \right) = o(N^{-c}), \quad (43)$$

$$\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0} \left(\|\tilde{P}_{\vartheta,k}^j - P_{\vartheta(\hat{\theta})}\| > \mu_{N,k}^j \right) = o(N^{-c}), \quad (44)$$

$$\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0} \left(\|g(\tilde{\theta}_k, \vartheta(\tilde{\theta}_k), \tilde{P}_k, \tilde{P}_{\vartheta,k}^j) - g(\hat{\theta}, \vartheta(\hat{\theta}), \hat{P}, P_{\vartheta(\hat{\theta})})\| > \mu_{N,k}^j \right) = o(N^{-c}). \quad (45)$$

Equation (43) follows from applying the mean value expansion to $\vartheta(\tilde{\theta}_k) - \vartheta(\hat{\theta})$ and using (34) and the finiteness of $\partial\vartheta(\theta)/\partial\theta$. To prove (44), note that applying the mean value expansion to $P_{\vartheta(\tilde{\theta}_k)} - P_{\vartheta(\hat{\theta})}$ with (43) gives $\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0} (\|P_{\vartheta(\tilde{\theta}_k)} - P_{\vartheta(\hat{\theta})}\| > \mu_{N,k}) = o(N^{-c})$. Therefore, it suffices to show

$$\sup_{\theta^0 \in \Theta_1} \Pr_{\theta^0} \left(\|\tilde{P}_{\vartheta,k}^j - P_{\vartheta(\tilde{\theta}_k)}\| > \mu_{N,k}^j \right) = o(N^{-c}). \quad (46)$$

For $j = 0$, (46) holds because $\tilde{P}_{\vartheta,k}^0 = P_{\vartheta(\theta^0)}$. Suppose (46) holds for $j = r - 1 \geq 0$. Expanding $\Psi(\tilde{P}_{\vartheta,k}^{r-1}, \vartheta(\tilde{\theta}_k))$ around $(P_{\vartheta(\tilde{\theta}_k)}, \vartheta(\tilde{\theta}_k))$ and applying $\Psi(P_{\vartheta(\tilde{\theta}_k)}, \vartheta(\tilde{\theta}_k)) = P_{\vartheta(\tilde{\theta}_k)}$ and $D_P\Psi(P_{\vartheta(\tilde{\theta}_k)}, \vartheta(\tilde{\theta}_k)) = 0$ give

$$\tilde{P}_{\vartheta,k}^r = \Psi(\tilde{P}_{\vartheta,k}^{r-1}, \vartheta(\tilde{\theta}_k)) = P_{\vartheta(\tilde{\theta}_k)} + R_{N,r},$$

where $\|R_{N,r}\| \leq K\|\tilde{P}_{\vartheta,k}^{r-1} - P_{\vartheta(\tilde{\theta}_k)}\|^2$ for a finite constant K . Thus (46) holds for $j = r$, and the proof for general $j \geq 1$ follows from induction. This proves (44). Finally, (45) follows from the finiteness of $Dg(\theta, \vartheta, P_\theta, P_\vartheta)$ and (34), (43), and (44). \square

8.11 Proof of Lemma 6

First, $D_{\mathbf{P}}l^{PL}(w_i; \mathbf{P}_\zeta, \zeta) = D_{\mathbf{P}}l^{MPL}(w_i; \mathbf{P}_\zeta, \zeta) = 0$ follows from the chain rule and Proposition 1. We proceed to prove the orthogonality results. $D_{\mathbf{P}}l^{PL}(w_i; \mathbf{P}_\zeta, \zeta) = 0$ and the information matrix equality imply that $E_{\zeta^0} D_{\mathbf{P}_\zeta} l^{PL}(w_i; \mathbf{P}^0, \zeta^0) = 0$. It follows that $D_{\mathbf{P}_\zeta} \mathcal{L}_N^{PL}(\mathbf{P}^0, \zeta^0) = O_p(N^{-1/2})$ since w_i is iid. Then, $D_{\mathbf{P}_\zeta} \mathcal{L}_N^{PL}(\hat{\mathbf{P}}_\zeta, \hat{\zeta}) = O_p(N^{-1/2})$ follows from expanding $D_{\mathbf{P}_\zeta} \mathcal{L}_N^{PL}(\mathbf{P}_\zeta, \hat{\zeta})$ around (\mathbf{P}^0, ζ^0) and using $\hat{\mathbf{P}} - \mathbf{P}^0, \hat{\zeta} - \zeta^0 = O_p(N^{-1/2})$ and Assumptions 4(g) and 4UH(g).

For the NMPL estimator, $D_{\mathbf{P}_\zeta} l^{MPL}(w_i; \mathbf{P}_\zeta, \zeta) = 0$ follows from the chain rule, $D_{\mathbf{P}} l^{PL}(w_i; \mathbf{P}_\zeta, \zeta) = 0$, and $D_{\mathbf{P}} \Psi(\mathbf{P}_\zeta, \zeta) = 0$. \square

8.12 Proof of Proposition 7

The proof follows the proofs of Proposition 2 and 3. Because the NFXP estimator maximizes the objective function of the NPL estimator if $\mathbf{P} = \hat{\mathbf{P}}$, we have, in place of (24), $D_\zeta \mathcal{L}^{PL}(\hat{\mathbf{P}}_{k-1}^{PL}, \hat{\zeta}_k^{PL}) =$

$D_\zeta \mathcal{L}^{PL}(\hat{\mathbf{P}}, \hat{\zeta}) = 0$. Assume $\hat{\mathbf{P}}_{k-1}^{PL} - \mathbf{P}^0 = o_p(1)$, then applying the generalized Taylor's theorem and following the argument used to prove Proposition 2 in conjunction with Lemma 6 gives $[E_{\zeta^0} D_{\zeta \zeta} l^{PL}(w_i; \mathbf{P}^0, \zeta^0) + o_p(1)](\hat{\zeta}_k^{PL} - \hat{\zeta}) = O_p(N^{-1/2} \|\hat{\mathbf{P}}_{k-1}^{PL} - \hat{\mathbf{P}}\| + \|\hat{\mathbf{P}}_{k-1}^{PL} - \hat{\mathbf{P}}\|^2)$. The stated result follows because $E_{\zeta^0} D_{\zeta \zeta} l^{PL}(w_i; \mathbf{P}^0, \zeta^0)$ is negative definite. The bound of $\hat{\mathbf{P}}_k^{PL} - \hat{\mathbf{P}}$ can be shown by expanding $\Psi(P_{k-1}^{PL, m}, \hat{\theta}_k^{PL, m})$ around $(P_{\hat{\theta}^m}, \hat{\theta}^m)$ and applying $D_P \Psi(P_{\hat{\theta}^m}, \hat{\theta}^m) = 0$. The required result follows by induction.

In case of the NMPL estimator, we have $D_\zeta \mathcal{L}^{MPL}(\hat{\mathbf{P}}_{k-1}^{MPL}, \hat{\zeta}_k^{MPL}) = D_\zeta \mathcal{L}^{MPL}(\hat{\mathbf{P}}, \hat{\zeta}) = 0$ in place of (27). The required result follows from repeating the argument of the proof of Proposition 3 in conjunction with Lemma 6 and $D_{\mathbf{P} \mathbf{P} \zeta} \mathcal{L}_N^{MPL}(\mathbf{P}^0, \zeta^0), D_{\zeta \mathbf{P} \zeta} \mathcal{L}_N^{MPL}(\mathbf{P}^0, \zeta^0) = O_p(N^{-1/2})$, which holds because $E_{\zeta^0} D_{\mathbf{P} \mathbf{P} \zeta} l^{MPL}(w_i; \mathbf{P}^0, \zeta^0) = 0$ and $E_{\zeta^0} D_{\zeta \mathbf{P} \zeta} l^{MPL}(w_i; \mathbf{P}^0, \zeta^0) = 0$ from the chain rule, $D_{\mathbf{P}} \Psi(\mathbf{P}^0, \zeta^0) = 0, E_{\zeta^0} D_{\mathbf{P} \mathbf{P} \zeta} l^{PL}(w_i; \mathbf{P}^0, \zeta^0) = 0$ and $E_{\zeta^0} D_{\mathbf{P} \zeta} l^{PL}(w_i; \mathbf{P}^0, \zeta^0) = 0$. \square

9 Appendix B: Auxiliary results

Lemma 8 collects the bounds that are used in the proof of Propositions 2-6. Lemma 9 collects the results on the derivatives of $\ln \Psi_2(P, \theta)$. Lemma 10 is our version (i.e., for $\hat{\alpha}$ and $\hat{\theta}_f$) of Lemma A.4 of A05. Lemma 11 is our version (i.e., for $\hat{\alpha}$ and $\hat{\theta}_f$) of Lemma A.6 of A05.

Lemma 8 *Suppose Assumptions 1-5 hold, $\bar{P} \rightarrow_p P^0$, and $\bar{\theta} \rightarrow_p \theta^0$. Let $\psi_i(P, \theta)$ denote either $\ln \Psi(P, \theta)(a_i|x_i)$ or $\ln \Psi_2(P, \theta)(a_i|x_i)$. Then*

- (a) $D^s \Psi(\bar{P}, \bar{\theta})(a_i|x_i) = O_p(1)$ for $s = 1, 2$,
- (b) $N^{-1} \sum_{i=1}^N \sup_{(P, \theta) \in B_P \times \Theta_0} \|D^s \psi_i(P, \theta)\|^q = O_p(1)$ for $q = 1, 2$ and $s = 1, \dots, 4$,
- (c) $N^{-1} \sum_{i=1}^N D_{P\alpha} \ln \Psi(P^0, \theta^0)(a_i|x_i) = O_p(N^{-1/2})$,
- (d) $N^{-1} \sum_{i=1}^N D^2 \psi_i(\bar{P}, \bar{\theta}) = E_{\theta^0} D^2 \psi_i(P^0, \theta^0) + O_p(\|\bar{P} - P^0\| + \|\bar{\theta} - \theta^0\| + N^{-1/2})$,
- (e) $\begin{cases} N^{-1} \sum_{i=1}^N D_\theta \psi_i(\bar{P}, \bar{\theta}) D_\theta \psi_i(\bar{P}, \bar{\theta}) \\ = E_{\theta^0} D_\theta \psi_i(P^0, \theta^0) D_\theta \psi_i(P^0, \theta^0) + O_p(\|\bar{P} - P^0\| + \|\bar{\theta} - \theta^0\| + N^{-1/2}). \end{cases}$

If Assumptions 1-8 hold, then (b) holds for $(P, \theta) \in B_P \times \Theta_1$.

Proof Parts (a) and (b) follow from Assumptions 4(c), 4(g), and 5(b). Part (c) follows because $E_{\theta^0} D_{P\alpha} \ln \Psi(P^0, \theta^0) = 0$ (zero operator) from the information matrix equality and Proposition 1 and w_i is iid. Parts (d) and (e) follow from part (b) and the law of large numbers. \square

Lemma 9 *Suppose Assumptions 1-4 hold. Then*

$$(a) \begin{cases} D_P \ln \Psi_2(P_\theta, \theta)(a_i|x_i) = 0, & D_\theta \ln \Psi_2(P_\theta, \theta)(a_i|x_i) = D \ln P_\theta(a_i|x_i), \\ D_{\theta\theta} \ln \Psi_2(P_\theta, \theta)(a_i|x_i) = D^2 \ln P_\theta(a_i|x_i), & D_{P\theta} \ln \Psi_2(P_\theta, \theta)(a_i|x_i) = 0. \\ \text{The same results hold for the derivatives of } \Psi_2(P_\theta, \theta)(a_i|x_i) \text{ and } P_\theta(a_i|x_i). \end{cases}$$

$$(b) \quad E_{\theta^0} D_{PP\theta} \ln \Psi_2(P^0, \theta^0)(a_i|x_i) = 0, \quad E_{\theta^0} D_{\theta P\theta} \ln \Psi_2(P^0, \theta^0)(a_i|x_i) = 0.$$

Proof The first result of part (a) is a simple consequence of Proposition 1 and the chain rule. For the other results of part (a), recall $P_\theta(a_i|x_i)$ is defined implicitly as a function of θ as $P_\theta(a_i|x_i) = \Psi(P_\theta, \theta)(a_i|x_i)$. Taking the derivative of $\ln P_\theta(a_i|x_i) = \ln \Psi(P_\theta, \theta)(a_i|x_i)$ and using Proposition 1 gives

$$D \ln P_\theta(a_i|x_i) = D_P \ln \Psi(P_\theta, \theta)(a_i|x_i) D P_\theta + D_\theta \ln \Psi(P_\theta, \theta)(a_i|x_i) = D_\theta \ln \Psi(P_\theta, \theta)(a_i|x_i). \quad (47)$$

It follows from the chain rule and $D_P \Psi(P_\theta, \theta) = 0$ that, for all $h \in \Theta$

$$\begin{aligned} D^2 \ln P_\theta(a_i|x_i)h &= D_{PP} \ln \Psi(P_\theta, \theta)(a_i|x_i) D P_\theta h \cdot D P_\theta + D_{\theta P} \ln \Psi(P_\theta, \theta)(a_i|x_i)h \cdot D P_\theta \\ &\quad + D_{P\theta} \ln \Psi(P_\theta, \theta)(a_i|x_i) \cdot D P_\theta h + D_{\theta\theta} \ln \Psi(P_\theta, \theta)(a_i|x_i)h. \end{aligned} \quad (48)$$

Now collect the derivatives of $\ln \Psi_2(P, \theta) = \ln \Psi(\Psi(P, \theta), \theta)$, where P is not necessarily the fixed point of $\Psi(\cdot, \theta)$.

$$D_\theta \ln \Psi_2(P, \theta)(a_i|x_i) = D_P \ln \Psi(\Psi(P, \theta), \theta)(a_i|x_i) D_\theta \Psi(P, \theta) + D_\theta \ln \Psi(\Psi(P, \theta), \theta)(a_i|x_i), \quad (49)$$

where $D_P \ln \Psi(\Psi(P, \theta), \theta)$ is the F-derivative of $\ln \Psi(P, \theta)$ with respect to P evaluated at $(\Psi(P, \theta), \theta)$, and similarly for $D_{PP} \ln \Psi(\Psi(P, \theta), \theta)$ etc. Furthermore, for all $h \in \Theta$

$$\begin{aligned} D_{\theta\theta} \ln \Psi_2(P, \theta)(a_i|x_i)h &= D_{PP} \ln \Psi(\Psi(P, \theta), \theta)(a_i|x_i) D_\theta \Psi(P, \theta)h \cdot D_\theta \Psi(P, \theta) \\ &\quad + D_{\theta P} \ln \Psi(\Psi(P, \theta), \theta)(a_i|x_i)h \cdot D_\theta \Psi(P, \theta) + D_P \ln \Psi(\Psi(P, \theta), \theta)(a_i|x_i) D_{\theta\theta} \Psi(P, \theta)h \\ &\quad + D_{P\theta} \ln \Psi(\Psi(P, \theta), \theta)(a_i|x_i) D_\theta \Psi(P, \theta)h + D_{\theta\theta} \ln \Psi(\Psi(P, \theta), \theta)(a_i|x_i)h. \end{aligned} \quad (50)$$

The cross derivative of $\Psi_2(P, \theta)$ takes the form, for all $h \in B_P$

$$\begin{aligned} D_{P\theta} \ln \Psi_2(P, \theta)(a_i|x_i)h &= D_{PP} \ln \Psi(\Psi(P, \theta), \theta)(a_i|x_i) D_P \Psi(P, \theta)h \cdot D_\theta \Psi(P, \theta) \\ &\quad + D_P \ln \Psi(\Psi(P, \theta), \theta)(a_i|x_i) D_{P\theta} \Psi(P, \theta)h + D_{P\theta} \ln \Psi(\Psi(P, \theta), \theta)(a_i|x_i) D_P \Psi(P, \theta)h. \end{aligned} \quad (51)$$

Evaluating (49)-(51) at $P = P_\theta$ with $D_P\Psi(P_\theta, \theta) = 0$ gives the first set of the results in part (a). The required results for the derivatives of $\Psi_2(P_\theta, \theta)(a_i|x_i)$ and $P_\theta(a_i|x_i)$ follow from the same argument.

To show part (b), taking the F-derivative of (51) and evaluating it at $P = P_\theta$ gives, for all $h_1, h_2 \in B_P$

$$\begin{aligned} D_{PP\theta} \ln \Psi_2(P_\theta, \theta)(a_i|x_i)h_1h_2 &= D_{PP} \ln \Psi(P_\theta, \theta)(a_i|x_i)D_{PP}\Psi(P_\theta, \theta)h_1h_2 \cdot D_\theta\Psi(P_\theta, \theta) \\ &\quad + D_{P\theta} \ln \Psi(P_\theta, \theta)(a_i|x_i)D_{PP}\Psi(P_\theta, \theta)h_1h_2, \\ D_{\theta P\theta} \ln \Psi_2(P_\theta, \theta)(a_i|x_i)h_1h_2 &= D_{PP} \ln \Psi(P_\theta, \theta)(a_i|x_i)D_{\theta P}\Psi(P_\theta, \theta)h_1h_2 \cdot D_\theta\Psi(P_\theta, \theta) \\ &\quad + D_{P\theta} \ln \Psi(P_\theta, \theta)(a_i|x_i)D_{\theta P}\Psi(P_\theta, \theta)h_1h_2. \end{aligned}$$

Part (b) follows because $E_{\theta^0}D_{PP} \ln \Psi(P^0, \theta^0)(a_i|x_i) = 0$ and $E_{\theta^0}D_{P\theta} \ln \Psi(P^0, \theta^0)(a_i|x_i) = 0$ from Proposition 1, the information matrix equality, and $w_i \sim \text{iid}$. \square

Lemma 10 *Suppose Assumptions 1-8 hold. Then, for all $\varepsilon > 0$,*

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

Proof From Lemma 5 of Andrews (2001), we have $\sup_{\theta_f^0 \in \Theta_f^1} \Pr_{\theta_f^0} (N^{1/2} \|\hat{\theta}_f - \theta_f^0\| > \varepsilon \ln N) = o(N^{-c})$ for all $\varepsilon > 0$.

Define $\rho_N(\alpha, \theta_f) = -N^{-1} \sum_{i=1}^N \ln P_{(\alpha, \theta_f)}(a_i|x_i)$ and $\rho(\alpha, \theta_f) = -E_{\theta^0} \ln P_{(\alpha, \theta_f)}(a_i|x_i)$, so that $\hat{\alpha} = \arg \min_{\alpha \in \Theta_\alpha} \rho_N(\alpha, \hat{\theta}_f)$. By Assumption 6(b), given any $\varepsilon > 0$, there exists $\delta > 0$ such that $\|\alpha - \alpha^0\| > \varepsilon$ implies $\rho(\alpha, \theta_f^0) - \rho(\alpha^0, \theta_f^0) \geq \delta$. Therefore, $\sup_{\theta^0 \in \Theta^1} \Pr_{\theta^0} (\|\hat{\alpha} - \alpha^0\| > \varepsilon) \leq \sup_{\theta^0 \in \Theta^1} \Pr_{\theta^0} (\rho(\hat{\alpha}, \theta_f^0) - \rho(\alpha^0, \theta_f^0) \geq \delta)$. Since $\rho(\alpha, \theta_f)$ is uniformly continuous, the right hand is no larger than

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

where the first inequality follows from $\rho_N(\hat{\alpha}, \hat{\theta}_f) - \rho_N(\alpha^0, \hat{\theta}_f) \leq 0$ and the last equality follows from $\sup_{\theta^0 \in \Theta^1} \Pr_{\theta^0} (\sup_{(\alpha, \theta_f) \in \Theta} |\rho_N(\alpha, \theta_f) - \rho(\alpha, \theta_f)| > \eta) = o(N^{-c})$ for all $\eta > 0$, which follows from (8.49) in Andrews (2001).

Therefore, we can use the argument in p. 34 of Andrews (2001) following his equation (8.51) to obtain $\inf_{\theta^0 \in \Theta^1} \Pr_{\theta^0} ((\partial/\partial\alpha)\rho_N(\hat{\alpha}, \hat{\theta}_f) = 0) = 1 - o(N^{-c})$. Then, the stated result for $\hat{\alpha}$ follows

from expanding $(\partial/\partial\alpha)\rho_N(\hat{\alpha}, \hat{\theta}_f)$ around (α^0, θ_f^0) and applying an argument similar to (8.52) in Andrews (2001). \square

Lemma 11 *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 $H_N(\hat{\theta}, \theta^0)$. 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 and that satisfies $G(E_{\theta^0} S_N(\theta^0)) = 0$ for all N large and all $\theta^0 \in \Theta_1$, and*

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

where \mathcal{B}_L denotes the class of all convex sets in \mathbb{R}^L .

Proof The proof follows the proof of Lemma A.6 of A05. Suppose $\Delta_N(\theta^0) = N^{1/2}(\hat{\theta} - \theta^0)$.

Define

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

From Lemma 10, $\hat{\theta}$ is in the interior of Θ with probability $1 - o(N^{-c})$, and we have $\inf_{\theta^0 \in \Theta_1} \Pr_{\theta^0}(s(\hat{\theta}) = 0) = 1 - o(N^{-c})$. Consequently, the proof of Lemma A.6 of A05 carries through if we replace $(\partial/\partial\theta)\rho_N(\theta)$ and $\hat{\theta}_N$ in A05 with our $s(\theta)$ and $\hat{\theta}$. The only difference is $(\partial/\partial x)\nu(E_{\theta^0} R_N(\theta_0), x)|_{x=0} = N^{-1} \sum_{i=1}^N E_{\theta^0} g(\tilde{W}_i, \theta_0)g(\tilde{W}_i, \theta_0)'$ in line 20, p. 210 of A05 needs to be replaced with

$$\frac{\partial}{\partial x} \nu(E_{\theta^0} R_N(\theta^0), x)|_{x=0} = E \begin{bmatrix} (\partial^2/\partial\alpha\partial\alpha') \ln P_{\theta^0}(a_i|x_i) & (\partial^2/\partial\alpha\partial\theta_f') \ln P_{\theta^0}(a_i|x_i) \\ 0 & (\partial^2/\partial\theta_f\partial\theta_f') \ln f_{\theta_f^0}(x'_i|a_i, x_i) \end{bmatrix}.$$

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

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

This equation corresponds to (A.35) of A05, where $R_N(\theta^0)$ and $e_N(\theta^0)$ are defined in the same manner as in A05 but his $(\partial/\partial\theta)\rho_N(\theta_0)$ replaced with our $s(\theta^0)$. The remaining part of his proof carries through, because Lemmas A.5 and A.8 of A05 holds in our context by our Assumptions 1-8, and our Lemma 10 plays the role of Lemma A.4 of A05. \square

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.
- [3] Andrews, D. W. K. (2002a). “Equivalence of the higher order asymptotic efficiency of k -step and extremum statistics.” *Econometric Theory* 18: 1040-1085.
- [4] Andrews, D. W. K. (2002b). “Higher-order improvements of a computationally attractive k -step bootstrap for extremum estimators.” *Econometrica* 70(1): 119-162.
- [5] Andrews, D. W. K. (2005). “Higher-order improvements of the parametric bootstrap for Markov processes.” in *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, eds. D. W. K. Andrews and J. H. Stock, Cambridge University Press.
- [6] Berkovec, J., and S. Stern (1991). “Job exit behavior of older man.” *Econometrica* 59: 189-210.
- [7] Cooper, R., J. Haltiwanger and L. Power (1999). “Machine replacement and the business cycle: lumps and bumps.” *American Economic Review* 89: 921-946.
- [8] Davidson, R. and J. G. MacKinnon (1999a). “Bootstrap testing in nonlinear models.” *International Economic Review* 40(2): 487-508.
- [9] Davidson, R. and J. G. MacKinnon (1999b). “The size distortion of bootstrap tests.” *Econometric Theory*, 15: 361-376.
- [10] Eckstein, Z. and K. Wolpin (1999). “Why youth drop out of high school: the impact of preferences, opportunities and abilities.” *Econometrica* 67(6): 1295-1339.
- [11] Gilleskie, D. (1998). “A dynamic stochastic model of medical care use and work absence.” *Econometrica* 66: 1-45.

- [12] Hotz, J. and R. A. Miller (1993). "Conditional choice probabilities and the estimation of dynamic models." *Review of Economic Studies* 60: 497-529.
- [13] Ichimura, H. and S. Lee (2004). Calculation of the asymptotic distribution of semiparametric M-estimators. Mimeographed, University College of London.
- [14] Imai, S. and M. Keane (2004). "Intertemporal labor supply and human capital accumulation," *International Economic Review* 45(2): 601-641.
- [15] Janssen, P., J. Jureckova, and N. Veraverbeke (1985). "Rate of convergence of one- and two-step M-estimators with applications to maximum likelihood and Pitman estimators." *Annals of Statistics* 13: 1222-1229.
- [16] Keane, M. P., and K. I. Wolpin (1997). "The career decisions of young men." *Journal of Political Economy* 105: 473-522.
- [17] Kim, J. H. (2005). Higher-order improvements of the restricted parametric bootstrap for tests. Mimeographed, National University of Singapore.
- [18] Miller, R. (1984). "Job matching and occupational choice," *Journal of Political Economy* 92: 1086-1120.
- [19] Pakes, A. (1986). "Patents as options: some estimates of the value of holding european patent stocks." *Econometrica* 54: 755-784.
- [20] Pfanzagl, J. (1974). "Asymptotic optimum estimation and test procedures." in Proceedings of the Prague Symposium on Asymptotics, Vol. 1., ed. J. Hájek, Prague: Charles University, 201-272.
- [21] Robinson, P. M. (1988). "The stochastic difference between econometric statistics." *Econometrica* 56(3): 531-548.
- [22] Rust, J. (1987). "Optimal replacement of GMC bus engines: an empirical model of Harold Zurcher." *Econometrica* 55(5): 999-1033.
- [23] Rust, J. (1988). "Maximum likelihood estimation of discrete control processes." *SIAM Journal of Control and Optimization* 26(5): 1006-1024.

- [24] Rust, J. (1996). “Numerical dynamic programming in economics.” in *Handbook of Computational Economics* eds. H. Amman, D. Kendrick and J. Rust. Elsevier, North Holland.
- [25] Rust, J., and C. Phelan (1997). “How social security and medicare affect retirement behavior in a world of incomplete markets.” *Econometrica* 65: 781-831.
- [26] Tauchen, G. (1986). “Finite state markov-chain approximation to univariate and vector autoregressions.” *Economics Letters* 20: 177-181.
- [27] Zeidler, E. (1986). *Nonlinear Functional Analysis and its Applications I: Fixed-Point Theorems*. New York, Springer-Verlag.

Table 1: Rejection Frequencies for Asymptotic Wald test at .10, .05, and .01 Levels

	Significance Levels		
	.10	.05	.01
$N = 500$	0.177	0.137	0.077
$N = 1000$	0.135	0.097	0.059
$N = 2000$	0.111	0.074	0.030

Notes: Based on 1000 simulated samples. The null hypothesis is $(\theta_1, \theta_2) = (-0.2, 0.1)$.

Table 2: Rejection Frequencies for Bootstrap Wald test at .10, .05, and .01 Levels

		Significance Levels		
		.10	.05	.01
Asymptotic		0.135	0.097	0.059
MLE		0.088	0.040	0.016
MLE-NULL		0.084	0.039	0.006
NPL	$k = 1$	0.086	0.035	0.012
	$k = 3$	0.090	0.042	0.016
NMPL	$k = 1$	0.082	0.041	0.004
	$k = 3$	0.083	0.041	0.004
1-NPL	$k = 1$	0.026	0.007	0.000
	$k = 3$	0.087	0.046	0.003
	$k = 5$	0.090	0.048	0.007
1-NMPL	$k = 1$	0.029	0.005	0.000
	$k = 3$	0.078	0.042	0.001
	$k = 5$	0.080	0.044	0.011

Notes: Based on 1000 simulated samples. The sample size is $N = 1000$ while the number of bootstrap samples is 599. The null hypothesis is $(\theta_1, \theta_2) = (-0.2, 0.1)$.

Table 3: 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
<i>95% CI</i>	$N = 500$	0.054	0.033	0.084	0.000	0.026	0.018
	$N = 1000$	0.044	0.021	0.078	0.000	0.029	0.009
	$N = 2000$	0.026	0.030	0.043	0.000	0.022	0.011
<i>90% CI</i>	$N = 500$	0.072	0.057	0.118	0.000	0.050	0.040
	$N = 1000$	0.055	0.047	0.109	0.000	0.067	0.024
	$N = 2000$	0.042	0.058	0.072	0.007	0.053	0.031

Notes: Based on 1000 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 $\theta_0 = 2.0$, $\theta_1 = -0.2$, and $\theta_2 = 0.1$ 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 for 90% CIs and, thus, the ideal values of “Miss Left” and “Miss Right” are 0.05 for 90% CIs while they are 0.025 for 95% CIs.

Table 4: Coverage Performance of Bootstrap 90% and 95% CIs for parameters θ_0 , θ_1 , and θ_2

		95% CIs						90% CIs					
		θ_0		θ_1		θ_2		θ_0		θ_1		θ_2	
		Miss Left	Miss Right	Miss Left	Miss Right	Miss Left	Miss Right	Miss Left	Miss Right	Miss Left	Miss Right	Miss Left	Miss Right
Asymptotic CIs		0.044	0.021	0.078	0.000	0.029	0.009	0.055	0.047	0.109	0.000	0.067	0.024
MLE-SY		0.031	0.040	0.051	0.000	0.021	0.018	0.043	0.064	0.097	0.000	0.060	0.031
MLE-ET		0.037	0.040	0.033	0.057	0.040	0.023	0.041	0.067	0.051	0.059	0.063	0.058
NPL-SY	$k = 1$	0.032	0.040	0.049	0.000	0.022	0.020	0.043	0.066	0.098	0.000	0.058	0.033
	$k = 3$	0.031	0.040	0.053	0.000	0.020	0.018	0.043	0.063	0.099	0.000	0.060	0.031
NPL-ET	$k = 1$	0.037	0.037	0.031	0.057	0.043	0.022	0.041	0.060	0.049	0.059	0.065	0.050
	$k = 3$	0.037	0.040	0.034	0.057	0.041	0.023	0.041	0.067	0.053	0.059	0.063	0.059
NMPL-SY	$k = 1$	0.031	0.035	0.051	0.000	0.016	0.024	0.039	0.065	0.119	0.000	0.049	0.035
	$k = 3$	0.031	0.035	0.051	0.000	0.016	0.024	0.039	0.065	0.119	0.000	0.049	0.035
NMPL-ET	$k = 1$	0.032	0.033	0.023	0.047	0.031	0.031	0.040	0.069	0.051	0.050	0.052	0.051
	$k = 3$	0.032	0.033	0.023	0.047	0.031	0.031	0.040	0.069	0.051	0.050	0.052	0.051
1-NPL-SY	$k = 1$	0.022	0.030	0.013	0.000	0.019	0.026	0.035	0.052	0.063	0.000	0.043	0.045
	$k = 3$	0.018	0.034	0.065	0.000	0.012	0.026	0.032	0.055	0.116	0.000	0.049	0.037
	$k = 5$	0.017	0.033	0.065	0.000	0.010	0.023	0.032	0.052	0.114	0.000	0.050	0.036
1-NPL-ET	$k = 1$	0.037	0.021	0.003	0.192	0.040	0.020	0.059	0.041	0.013	0.236	0.077	0.031
	$k = 3$	0.029	0.025	0.028	0.040	0.022	0.025	0.037	0.053	0.065	0.041	0.057	0.043
	$k = 5$	0.025	0.029	0.029	0.040	0.026	0.027	0.036	0.053	0.065	0.041	0.057	0.049
1-NMPL-SY	$k = 1$	0.025	0.022	0.015	0.000	0.019	0.020	0.034	0.041	0.068	0.000	0.045	0.039
	$k = 3$	0.019	0.027	0.067	0.000	0.014	0.018	0.034	0.045	0.124	0.000	0.054	0.036
	$k = 5$	0.021	0.026	0.070	0.000	0.015	0.017	0.031	0.044	0.121	0.000	0.056	0.032
1-NMPL-ET	$k = 1$	0.050	0.016	0.001	0.186	0.046	0.014	0.072	0.033	0.015	0.22	0.087	0.026
	$k = 3$	0.031	0.021	0.028	0.034	0.026	0.020	0.048	0.038	0.067	0.036	0.066	0.043
	$k = 5$	0.026	0.023	0.032	0.034	0.035	0.022	0.037	0.039	0.070	0.037	0.067	0.049

Notes: Based on 1000 simulated samples, each with the sample size of 1000. The number of bootstrap samples is 599.

Table 5: Average Replacement Rates, Average Revenues, and Revenue Dispersions under Counterfactual Experiments

	The Model	Counterfactual Experiments			
		Permanent Change (Stead State)		Temporary Change (Initial Year)	
		30 % More	30 % Less	30 % More	30 % Less
Ave. Replacement Rate	0.336	0.406	0.282	0.468	0.224
Ave. Revenue	0.830	0.873	0.789	0.881	0.783
Revenue Dispersion	0.142	0.132	0.154	0.132	0.155

Table 6: Coverage Performance of Asymptotic and Bootstrap 95% CIs for Average Revenues under Counterfactual Policies

Counterfactual θ_0		Permanent Change (Steady State)				Temporary Change (Initial Year)			
		30% Less		30% More		30% Less		30% More	
		Miss Left	Miss Right	Miss Left	Miss Right	Miss Left	Miss Right	Miss Left	Miss Right
Asymptotic CIs		0.035	0.041	0.035	0.044	0.032	0.044	0.039	0.043
MLE-SY		0.031	0.035	0.037	0.037	0.028	0.039	0.038	0.035
MLE-ET		0.025	0.036	0.029	0.038	0.024	0.039	0.028	0.035
NPL-SY	$k = 1$	0.029	0.035	0.036	0.037	0.027	0.039	0.039	0.034
	$k = 3$	0.032	0.035	0.038	0.037	0.028	0.039	0.039	0.035
NPL-ET	$k = 1$	0.020	0.037	0.025	0.038	0.022	0.039	0.026	0.035
	$k = 3$	0.025	0.037	0.029	0.038	0.024	0.039	0.029	0.035
NMPL-SY	$k = 1$	0.024	0.031	0.027	0.034	0.021	0.035	0.031	0.031
	$k = 3$	0.024	0.031	0.027	0.034	0.021	0.035	0.031	0.031
NMPL-ET	$k = 1$	0.012	0.033	0.022	0.033	0.015	0.035	0.020	0.031
	$k = 3$	0.012	0.033	0.022	0.033	0.015	0.035	0.020	0.031
1-NPL-SY	$k = 1$	0.015	0.012	0.021	0.018	0.011	0.011	0.016	0.007
	$k = 3$	0.035	0.019	0.038	0.022	0.030	0.016	0.037	0.010
	$k = 5$	0.034	0.021	0.037	0.024	0.031	0.024	0.041	0.022
1-NPL-ET	$k = 1$	0.005	0.031	0.015	0.032	0.004	0.011	0.005	0.007
	$k = 3$	0.020	0.027	0.027	0.027	0.022	0.015	0.026	0.015
	$k = 5$	0.023	0.025	0.028	0.025	0.024	0.027	0.027	0.025
1-NMPL-SY	$k = 1$	0.007	0.018	0.021	0.022	0.008	0.020	0.011	0.013
	$k = 3$	0.038	0.022	0.040	0.024	0.032	0.022	0.040	0.018
	$k = 5$	0.039	0.024	0.041	0.025	0.035	0.026	0.044	0.024
1-NMPL-ET	$k = 1$	0.000	0.028	0.003	0.028	0.000	0.020	0.001	0.013
	$k = 3$	0.016	0.028	0.026	0.028	0.021	0.025	0.027	0.022
	$k = 5$	0.023	0.028	0.030	0.028	0.024	0.028	0.030	0.028

Notes: Based on 1000 simulated samples.

Table 7: Coverage Performance of Asymptotic and Bootstrap 95% CIs for Revenue Dispersions under Counterfactual Policies

Counterfactual θ_0		Permanent Change (Steady State)				Temporary Change (Initial Year)			
		30% Less		30% More		30% Less		30% More	
		Miss Left	Miss Right	Miss Left	Miss Right	Miss Left	Miss Right	Miss Left	Miss Right
Asymptotic CIs		0.000	0.140	0.000	0.149	0.000	0.137	0.000	0.151
MLE-SY		0.000	0.041	0.000	0.043	0.000	0.072	0.000	0.074
MLE-ET		0.065	0.026	0.049	0.027	0.059	0.034	0.048	0.035
NPL-SY	$k = 1$	0.000	0.042	0.000	0.043	0.000	0.040	0.000	0.044
	$k = 3$	0.000	0.043	0.000	0.045	0.000	0.040	0.000	0.043
NPL-ET	$k = 1$	0.063	0.027	0.049	0.028	0.058	0.027	0.048	0.028
	$k = 3$	0.064	0.027	0.049	0.028	0.060	0.026	0.048	0.027
NMPL-SY	$k = 1$	0.000	0.033	0.000	0.034	0.000	0.033	0.000	0.034
	$k = 3$	0.000	0.033	0.000	0.034	0.000	0.033	0.000	0.034
NMPL-ET	$k = 1$	0.056	0.019	0.042	0.022	0.050	0.018	0.042	0.021
	$k = 3$	0.056	0.019	0.042	0.022	0.050	0.018	0.042	0.021
1-NPL-SY	$k = 1$	0.000	0.053	0.000	0.054	0.000	0.038	0.000	0.045
	$k = 3$	0.000	0.045	0.000	0.046	0.000	0.038	0.000	0.047
	$k = 5$	0.000	0.045	0.000	0.048	0.000	0.042	0.000	0.048
1-NPL-ET	$k = 1$	0.142	0.026	0.069	0.029	0.003	0.022	0.001	0.025
	$k = 3$	0.042	0.024	0.033	0.025	0.029	0.022	0.023	0.025
	$k = 5$	0.039	0.023	0.033	0.025	0.043	0.022	0.032	0.025
1-NMPL-SY	$k = 1$	0.000	0.054	0.000	0.059	0.000	0.040	0.000	0.047
	$k = 3$	0.000	0.051	0.000	0.053	0.000	0.044	0.000	0.051
	$k = 5$	0.000	0.050	0.000	0.052	0.000	0.049	0.000	0.052
1-NMPL-ET	$k = 1$	0.134	0.034	0.068	0.034	0.004	0.027	0.002	0.028
	$k = 3$	0.037	0.029	0.030	0.031	0.028	0.027	0.026	0.030
	$k = 5$	0.033	0.027	0.029	0.030	0.035	0.027	0.028	0.030

Notes: Based on 1000 simulated samples.

