

Sequential Estimation of Dynamic Programming Models

Hiroyuki Kasahara
Department of Economics
University of British Columbia
hkasahar@gmail.com

Katsumi Shimotsu
Department of Economics
Hitotsubashi University
shimotsu@econ.hit-u.ac.jp

October 11, 2011

Preliminary and Incomplete

Abstract

This paper develops a new computationally attractive procedure for estimating dynamic discrete choice models that is applicable to a wide range of dynamic programming models. The proposed procedure can accommodate unobserved state variables that (i) are neither additively separable nor follow generalized extreme value distribution, (ii) are serially correlated, and (iii) affect the choice set. Our estimation algorithm sequentially updates the parameter estimate and the value function estimate. It builds upon the idea of the iterative estimation algorithm proposed by Aguirregabiria and Mira (2002, 2007) but conducts iteration using the value function mapping rather than the policy iteration mapping. Its implementation is straightforward in terms of computer programming; unlike the Hotz-Miller type estimators, there is no need to reformulate a fixed point mapping in the value function space as that in the space of probability distributions. It is also applicable to estimate models with unobserved heterogeneity. We analyze the convergence property of our sequential algorithm and derive the conditions for its convergence. We develop an approximated procedure which reduces computational cost substantially without deteriorating the convergence rate. We further extend our sequential procedure for estimating dynamic programming models with an equilibrium constraint, which include dynamic game models and dynamic macroeconomic models.

Keywords: dynamic discrete choice, value function mapping, nested pseudo likelihood, unobserved heterogeneity, equilibrium constraint.

JEL Classification Numbers: C13, C14, C63.

1 Introduction

Numerous empirical studies have demonstrated that the estimation of dynamic discrete models enhances our understanding of individual and firm behavior and provide important policy implications.¹ The literature on estimating dynamic models of discrete choice was pioneered by Gotz and McCall (1980), Wolpin (1984), Miller (1984), Pakes (1986), and Rust (1987, 1988). Standard methods for estimating infinite horizon dynamic discrete choice models require repeatedly solving the fixed point problem (i.e., Bellman equation) during optimization and can be very costly when the dimensionality of state space is large.

To reduce the computational burden, Hotz and Miller (1993) developed a simpler two-step estimator, called *Conditional Choice Probability (CCP) estimator*, by exploiting the inverse mapping from the value functions to the conditional choice probabilities.² Aguirregabiria and Mira (2002, 2007) developed a recursive extension of the CCP estimator called the *nested pseudo likelihood (NPL) algorithm*. These Hotz and Miller-type estimators have limited applicability, however, when unobserved state variables are not additively separable and (generalized-) extreme value distributed because evaluating the inverse mapping from the value functions to the conditional choice probabilities is computationally difficult. Recently, Arcidiacono and Miller (2008) develop estimators that relax some of the limitations of the CCP estimator by combining the Expectation-Maximization (EM) algorithm with the NPL algorithm in estimating models with unobserved heterogeneity. While Arcidiacono and Miller provide important contributions to the literature, little is known about the convergence property of their algorithm, and it is not clear how computationally easy it is to apply their estimation method to a model that does not exhibit finite time dependence.

This paper develops a new estimation procedure for infinite horizon dynamic discrete choice models with unobserved state variables that (i) are neither additively separable nor follow generalized extreme value distribution, (ii) are serially correlated, and (iii) affect the choice set. Our estimation method is based on the value function mapping (i.e., Bellman equation) and, hence, unlike the Hotz-Miller type estimators, there is no need to reformulate a Bellman equation as a fixed point mapping in the space of probability distributions (i.e., policy iteration operator). This is the major advantage of our method over the Hotz-Miller type estimators because evaluating the policy iteration operator is often difficult without the assumption of additively-separable unobservables with generalized extreme value distribution. Implementing our procedure is straightforward in terms of computer programming once the value iteration mapping is coded in a computer language.

¹Contributions include Berkovec and Stern (1991), Keane and Wolpin (1997), Rust and Phelan (1997), Rothwell and Rust (1997), Altug and Miller (1998), Gilleskie (1998), Eckstein and Wolpin (1999), Aguirregabiria (1999), Kasahara and Lapham (2008), and Kasahara (2009).

²A number of recent papers in empirical industrial organization build on the idea of Hotz and Miller (1993) to develop two-step estimators for models with multiple agents (e.g., Bajari, Benkard, and Levin, 2007; Pakes, Ostrovsky, and Berry, 2007; Pesendorfer and Schmidt-Dengler, 2008; Bajari and Hong, 2006).

Our estimation algorithm is analogous to the NPL algorithm [cf., Aguirregabiria and Mira (2002, 2007) and Kasahara and Shimotsu (2008a, 2008b)] but its iteration is based on the value function mapping rather than the policy iteration mapping. Our procedure iterates on the following two steps. First, given an initial estimator of the value function, we estimate the model’s parameter by solving a *finite horizon q -period model* in which the $(q+1)$ -th period’s value function is given by the initial value function estimate. Second, we update the value function estimate by solving a q -period model with the updated parameter estimate starting from the previous value function estimate as the continuation value in the q -th period. This sequential algorithm is computationally easy if we choose a small value of q ; if we choose $q = 1$, for instance, then the computational cost of solving this finite horizon model is equivalent to solving a static model. Iterating this procedure generates a sequence of estimators of the parameter and value function. Upon convergence, the limit of this sequence does not depend on an initial value function estimate. Hence, our method is applicable even when an initial consistent estimator of value function is not available.

We analyze the convergence property of our proposed sequential algorithm. The possibility of non-convergence of the original NPL algorithm (Aguirregabiria and Mira, 2002, 2007) is a concern as illustrated by Pesendorfer and Schmidt-Dengler (2008) and Collard-Wexler (2006).³ Since our algorithm is very similar to the original NPL algorithm, understanding the convergence property of our sequential algorithm is important. We show that a key determinant of the convergence is the *contraction* property of the value function mapping. By the Blackwell’s sufficient condition, the value function mapping is a contraction where a discount factor determines the contraction rate, and iterating the value function mapping improves the contraction property. As a result, our sequential algorithm achieves convergence when we choose sufficiently large q . To reduce computational cost further, we also develop an approximation procedure called the approximate q -NPL algorithm. This approximate algorithm has substantially less computational cost than the original sequential algorithm but has the same first-order convergence rate as the original sequential algorithm.

We extend our estimation procedure to a class of dynamic programming models in which the probability distribution of state variables satisfies some equilibrium constraints. This class of models includes models of dynamic games where the players’ choice probability is a fixed point of a best reply mapping and dynamic macroeconomic models with heterogeneous agents where each agent solves a dynamic optimization problem given the rationally expected price process which is consistent with the actual price process generated from the agent’s decision rule.

The rest of the paper is organized as follows. Section 2 illustrates the basic idea of our

³Pesendorfer and Schmidt-Dengler (2008) provided simulation evidence that the NPL algorithm may not necessarily converge while Collard-Wexler (2006) used the NPL algorithm to estimate a model of entry and exit for the ready-mix concrete industry and found that \hat{P}_j ’s “cycle around several values without converging.” Kasahara and Shimotsu (2008b) analyze the conditions under which the NPL algorithm achieves convergence and derive its convergence rate.

algorithm by a simple example. Section 3 introduces a class of single-agent dynamic programming models, presents our sequential estimation procedure, and derives its convergence property. Section 4 extends our estimation procedure to dynamic programming models with equilibrium constraint. Section 5 reports some simulation results.

2 Example: Machine Replacement Model

2.1 A Single-agent Dynamic Programming Model

To illustrate the basic idea of our estimator, consider the following version of Rust's machine replacement model. Let x_t denote machine age and let $a_t \in \{0, 1\}$ represent the machine replacement decision. Both x_t and a_t are observable to a researcher. There are two state variables in the model that are not observable to a researcher: an idiosyncratic productivity shock ϵ_t and a choice-dependent cost shock $\xi_t(a_t)$. The profit function is given by $u_\theta(a_t, x_t, \epsilon_t) + \xi_t(a_t)$, where $u_\theta(a_t, x_t, \epsilon_t) = \exp(\theta_1 x_t(1 - a_t) + \epsilon_t) - \theta_2 a_t$. Here, $\exp(\theta_1 x_t(1 - a_t) + \epsilon_t)$ represents the revenue function with $\theta_1 < 0$, and θ_2 is machine price (machine replacement cost). We assume that $\xi_t = (\xi_t(0), \xi_t(1))'$ follows Type 1 extreme value distribution independently across alternatives while ϵ_t is independently drawn from $N(0, \sigma_\epsilon^2)$. The transition function of machine age x_t is given by $x_{t+1} = a_t + (1 - a_t)(x_t + 1)$.

A firm maximizes the expected discounted sum of revenues, $E[\sum_{j=0}^{\infty} \beta^j (u_\theta(a_{t+j}, x_{t+j}, \epsilon_{t+j}) + \xi_{t+j}(a_{t+j})) | a_t, x_t]$. The Bellman equation for this dynamic optimization problem is written as $W(x, \epsilon, \xi) = \max_{a \in \{0, 1\}} u_\theta(a, x, \epsilon) + \xi(a) + \beta \int \int W(a + (1 - a)(x + 1), \epsilon', \xi') g_\epsilon(d\epsilon' | x) g_\xi(d\xi' | x)$. Define the integrated value function $V(x) = \int \int W(x, \epsilon, \xi) g_\epsilon(d\epsilon' | x) g_\xi(d\xi' | x)$. Then, using the properties of Type 1 extreme value distribution, the integrated Bellman equation is written as:

$$V(x) = \int \left(\gamma + \ln \left(\sum_{a \in \{0, 1\}} \exp(u_\theta(a, x, \epsilon') + \beta V(a + (1 - a)(x + 1))) \right) \right) \frac{\phi(\epsilon'/\sigma_\epsilon)}{\sigma_\epsilon} d\epsilon' \equiv [\Gamma(\theta, V)](x), \quad (1)$$

where $\phi(\cdot)$ is the standard normal density function and γ is Euler's constant. The Bellman operator $\Gamma(\theta, \cdot)$ is defined by the right hand side of this integrated Bellman equation. Denote the fixed point of the integrated Bellman equation (1) by $V_\theta [= \Gamma(\theta, V_\theta)]$. The value of $V_\theta(x)$ represents the value of a firm with machine age x . Given V_θ , the conditional choice probability of replacement (i.e., $a = 1$) is given by

$$P_\theta(a = 1 | x) = \int \left(\frac{\exp(u_\theta(1, x, \epsilon') + \beta V_\theta(1))}{\sum_{a' \in \{0, 1\}} \exp(u_\theta(a', x, \epsilon') + \beta V_\theta(a' + (1 - a')(x + 1)))} \right) \frac{\phi(\epsilon'/\sigma_\epsilon)}{\sigma_\epsilon} d\epsilon' \equiv \Lambda(\theta, V_\theta), \quad (2)$$

while $P_\theta(a = 0 | x) = 1 - P_\theta(a = 1 | x)$. Here, the operator defined by the right hand side of (2), denoted by $\Lambda(\theta, \cdot)$, is a mapping from the value function space into the choice probability space.

To estimate the unknown parameter vector θ given a cross sectional data $\{x_i, a_i\}_{i=1}^n$, where n is the sample size, we may use the nested fixed point (NFXP) algorithm (Rust, 1987) by repeatedly solving the fixed point of (1) and evaluating the conditional choice probabilities (2) for every candidate value of θ to maximize the likelihood $\sum_{i=1}^n \ln P_\theta(a_i|x_i)$, where the integral with respect to ϵ can be evaluated by quadrature methods or simulations. The NFXP algorithm is costly because it is computationally intensive to solve the fixed point of (1). Estimating this replacement model using the Hotz-Miller type estimators [cf., Hotz and Miller (1993) and Aguirregabiria and Mira (2002, 2007)] is not straightforward because evaluating the inverse mapping from the value functions to the conditional choice probabilities is computationally difficult due to the presence of normally distributed shocks, ϵ .

We propose a simple alternative estimation method applicable to models with unobserved state variables that are neither additively separable and nor extreme-value distributed. Our estimation algorithm is based on solving a finite horizon “q-period” model in which the continuation value for the q-th period is replaced with its estimate \tilde{V}_0 . Namely, we evaluate the likelihood by applying the fixed point iterations only q -times starting from the initial estimate \tilde{V}_0 as:

$$\max_{\theta \in \Theta} n^{-1} \sum_{i=1}^n \ln \Lambda(\theta, \Gamma^q(\theta, \tilde{V}_0))(a_i|x_i), \quad \text{where} \quad \Gamma^q(\theta, \tilde{V}_0) \equiv \underbrace{\Gamma(\theta, \Gamma(\theta, \dots \Gamma(\theta, \tilde{V}_0)))}_{q \text{ times}}, \quad (3)$$

where $\Gamma^q(\theta, \cdot)$ is a q -fold operator of $\Gamma(\theta, \cdot)$. Note that $\Gamma^q(\theta, \tilde{V}_0)(x)$ in (3) represents the value of a firm with machine age x when a firm makes an optimal dynamic decision over q -periods where the $(q + 1)$ -th period’s value (i.e., the continuation value in the q -th period) is \tilde{V}_0 . Solving the optimization problem (3) is much less computationally intensive than implementing the NFXP algorithm.

An estimator of θ defined by (3) is generally inconsistent unless \tilde{V}_0 is consistent. When an initial consistent estimator for V is not available, we may apply the idea of the NPL algorithm [cf., Aguirregabiria and Mira (2002, 2007) and Kasahara and Shimotsu (2008a, 2008b)] to estimate θ consistently. Once an estimate of θ is obtained from solving the optimization problem (3), one updates the value function estimate \tilde{V}_0 as $\tilde{V}_1 = \Gamma^q(\hat{\theta}, \tilde{V}_0)$. Next, one obtains the updated estimator of θ , $\tilde{\theta}_1$, by solving (3) using \tilde{V}_1 in place of \tilde{V}_0 . Iterating this procedure generates a sequence of estimators $\{\tilde{\theta}_j, \tilde{V}_j\}_{j=1}^\infty$. Upon convergence, the limit of this sequence is independent of the initial estimator \tilde{V}_0 , and the limiting $\tilde{\theta}_\infty$ is consistent and asymptotically normally distributed under certain regularity conditions.

2.2 Equilibrium Constraint

We may extend our estimation procedure to models with an equilibrium constraint. Consider the economy with a measure one continuum of ex ante identical firms, each of which max-

imizes the expected discounted sum of revenues by making a machine replacement decision as in the previous section. Now suppose that the price of machine, denoted by r , is endogenously determined in the stationary equilibrium. Let the supply function of machines be exogenously given by $S(r, \theta)$ while the demand for machine in the stationary equilibrium is equal to $D(r, \theta, P) = \sum_{x=1}^{\infty} P(a|x) f^*(x; \theta, P)$, where $f^*(x; \theta, P)$ is the stationary distribution of x when each firm's conditional choice probabilities are given by P .⁴ Define the equilibrium machine price function, $r(\theta, P)$, by the equilibrium condition $S(r(\theta, P), \theta) = D(r(\theta, P), \theta, P)$.

Each agent treats the equilibrium machine price as exogenously given, and the profit function now depends on the equilibrium machine price as: $u_{\theta}^P(a_t, x_t, \epsilon_t) = \exp(\theta_1 x_t (1 - a_t) + \epsilon_t) - r(\theta, P) a_t$. Then, the Bellman equation is given by

$$V(x) = \int \left(\gamma + \ln \left(\sum_{a \in \{0,1\}} \exp(u_{\theta}^P(a, x, \epsilon') + \beta V(a + (1-a)(x+1))) \right) \right) \frac{\phi(\epsilon'/\sigma_{\epsilon})}{\sigma_{\epsilon}} d\epsilon' \equiv [\Gamma(\theta, V, P)](x),$$

while the conditional choice probability is given by

$$P(a = 1|x) = \int \left(\frac{\exp(u_{\theta}^P(1, x, \epsilon') + \beta V(1))}{\sum_{a' \in \{0,1\}} \exp(u_{\theta}^P(a', x, \epsilon') + \beta V(a' + (1-a')(x+1)))} \right) \frac{\phi(\epsilon'/\sigma_{\epsilon})}{\sigma_{\epsilon}} d\epsilon' \equiv \Lambda(\theta, V, P).$$

Note that the mappings Γ and Λ depend on both V and P . For each value of θ , let (V_{θ}, P_{θ}) be the fixed point of the system of equations $V_{\theta} = \Gamma(\theta, V_{\theta}, P_{\theta})$ and $P_{\theta} = \Lambda(\theta, V_{\theta}, P_{\theta})$. The equilibrium machine price is then given by $r(\theta, P_{\theta})$.

Suppose that we have an initial consistent estimator of (V, P) denoted by $(\tilde{V}_0, \tilde{P}_0)$. We may consistently estimate the parameter θ as $\tilde{\theta}_1 = \arg \max_{\theta \in \Theta} n^{-1} \sum_{i=1}^n \ln \Lambda(\theta, \Gamma^q(\theta, \tilde{V}_0, \tilde{P}_0), \tilde{P}_0)(a_i|x_i)$. Once an estimate of θ is obtained, one can update the value function estimate $(\tilde{V}_0, \tilde{P}_0)$ as $\tilde{V}_1 = \Gamma^q(\tilde{\theta}_1, \tilde{V}_0, \tilde{P}_0)$ and $\tilde{P}_1 = \Lambda(\tilde{\theta}_1, \tilde{V}_1, \tilde{P}_0)$, which can provide a more accurate estimator of (V, P) than $(\tilde{V}_0, \tilde{P}_0)$. Iterating this procedure generates a sequence of estimators $\{\tilde{\theta}_j, \tilde{V}_j, \tilde{P}_j\}_{j=1}^{\infty}$. If it converges, the limit is independent of the initial estimator $(\tilde{V}_0, \tilde{P}_0)$. We analyze the conditions under which this algorithm converges.

3 Dynamic Programming Model

3.1 The model without unobserved heterogeneity

An agent maximizes the expected discounted sum of utilities, $E[\sum_{j=0}^{\infty} \beta^j U_{\theta}(a_{t+j}, s_{t+j})|a_t, s_t]$, where s_t is the vector of states and a_t is a discrete action to be chosen from the constraint set $G_{\theta}(s_t) \subset A \equiv \{1, 2, \dots, |A|\}$. The transition probabilities are given by $p_{\theta}(s_{t+1}|s_t, a_t)$. The Bell-

⁴The stationary distribution satisfies $f^*(x, \theta, P) = P(a = 0|x-1) f^*(x-1, \theta, P)$ for $x > 1$ and $f^*(1, \theta, P) = \sum_{x'=1}^{\infty} P(a = 1|x') f^*(x', \theta, P)$.

man equation for this dynamic optimization problem is written as $W(s_t) = \max_{a \in G_\theta(s_t)} U_\theta(a, s_t) + \beta \int W(s_{t+1}) p_\theta(s_{t+1}|s_t, a) ds_{t+1}$. From the viewpoint of an econometrician, the state vector can be partitioned as $s_t = (x_t, \xi_t)$, where $x_t \in X$ is observable state variable and ξ_t is idiosyncratic unobservable state variable. We make the following assumptions.

Assumption 1 (Conditional Independence of ξ_t) *The transition probability function of the state variables can be written as $p_\theta(s_{t+1}|s_t, a_t) = g_\theta(\xi_{t+1}|x_{t+1})f_\theta(x_{t+1}|x_t, a_t)$.*

Assumption 2 (Finite support for x) *The support of x is finite and given by $X = \{1, \dots, |X|\}$.*

Accordingly, P and V are represented with $L \times 1$ vectors, where $L = |A||X|$. It is assumed that the form of U_θ , G_θ , and f_θ are known up to an unknown K -dimensional vector $\theta \in \Theta \subset \mathbb{R}^K$. We are interested in estimating the parameter vector θ from the sample data $\{x_i, a_i\}_{i=1}^n$, where n is the sample size.

Define the integrated value function $V(x) = \int W(x, \xi) g_\theta(\xi|x) d\xi$, 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

$$V(x) = \int \max_{a \in G_\theta(x, \xi)} \left\{ U_\theta(a, x, \xi) + \beta \sum_{x' \in X} V(x') f_\theta(x'|x, a) \right\} g_\theta(\xi|x) d\xi \quad (4)$$

Define the Bellman operator defined by the right-hand side of the above Bellman equation as $[\Gamma(\theta, V)](x) \equiv \int \max_{a \in G_\theta(x, \xi)} \{ U_\theta(a, x, \xi) + \beta \sum_{x' \in X} V(x') f_\theta(x'|x, a) \} g_\theta(\xi|x) d\xi$. The Bellman equation (4) is compactly written as $V = \Gamma(\theta, V)$.

Let $P(a|x)$ denote the conditional choice probabilities of the action a given the 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

$$P(a|x) = \int I \left\{ a = \arg \max_{j \in G_\theta(x, \xi)} v_\theta(j, x, \xi, V) \right\} g_\theta(\xi|x) d\xi \quad (5)$$

where $v_\theta(a, x, \xi, V) = u_\theta(a, x, \xi) + \beta \sum_{x' \in X} V(x') f_\theta(x'|x, a)$ is the choice-specific value function and $I(\cdot)$ is an indicator function. The right-hand side of the equation (5) can be viewed as a mapping from one Banach (B-) space B_V to another B-space B_P . Define the mapping $\Lambda(\theta, V) : \Theta \times B_V \rightarrow B_P$ as

$$[\Lambda(\theta, V)](a|x) \equiv \int I \left\{ a = \arg \max_{j \in G_\theta(x, \xi)} v_\theta(j, x, \xi, V) \right\} g_\theta(\xi|x) d\xi. \quad (6)$$

Let θ^0 and P^0 denote the true parameter value and the true conditional choice probabilities. Let V^0 denote the true integrated value function. Then, P^0 and V^0 are related as $P^0 = \Lambda(\theta^0, V^0)$. Note that V_0 is the fixed point of $\Gamma(\theta^0, \cdot)$ and hence $V^0 = \Gamma(\theta^0, V^0)$.

Consider a cross-sectional data set $\{a_i, x_i\}_{i=1}^n$ where (a_i, x_i) is randomly drawn across i 's from the population. The maximum likelihood estimator (MLE) solves the following constrained maximization problem:

$$\max_{\theta \in \Theta} n^{-1} \sum_{i=1}^n \ln \{[\Lambda(\theta, V)](a_i|x_i)\} \quad \text{subject to} \quad V = \Gamma(\theta, V). \quad (7)$$

Computation of the MLE by the NFXP algorithm requires repeatedly solving all the fixed points of $V = \Gamma(\theta, V)$ at each parameter value to maximize the objective function with respect to θ . If evaluating the fixed point of $\Gamma(\theta, \cdot)$ is costly, this is computationally very demanding.

We propose a sequential algorithm, the q -NPL algorithm, to estimate θ . The q -NPL algorithm is similar to the algorithms by Aguirregabiria and Mira (2002, 2007) and Kasahara and Shimotsu (2008a, 2008b), but, unlike theirs, our algorithm is based on a fixed point mapping defined in the value function space rather than in the probability space. Since it is often difficult to construct a fixed point mapping in the probability space when unobserved state variables are not distributed according to generalized extreme value distribution and the number of choices are larger than three, our proposed method is applicable to a wider class of dynamic programming models than a class of models they consider.

Define a q -fold operator of Γ as

$$\Gamma^q(\theta, V) \equiv \underbrace{\Gamma(\theta, \Gamma(\theta, \dots \Gamma(\theta, \Gamma(\theta, V)) \dots))}_{q \text{ times}}.$$

Starting from an initial estimate \tilde{V}_0 , the q -NPL algorithm iterates the following steps until $j = k$:

Step 1: Given \tilde{V}_{j-1} , update θ by $\tilde{\theta}_j = \arg \max_{\theta \in \Theta} n^{-1} \sum_{i=1}^n \ln \left\{ \left[\Lambda(\theta, \Gamma^q(\theta, \tilde{V}_{j-1})) \right] (a_i|x_i) \right\}$.

Step 2: Update \tilde{V}_{j-1} using the obtained estimate $\tilde{\theta}_j$: $\tilde{V}_j = \Gamma^q(\tilde{\theta}_j, \tilde{V}_{j-1})$.

Evaluating the objective function for a value of θ involves only q evaluations of the Bellman operator $\Gamma(\theta, \cdot)$ and one evaluation of probability operator $\Lambda(\theta, \cdot)$. The computational cost of Step 1 is roughly equivalent to that of estimating a model with q periods.

This algorithm generates a sequence of estimators $\{\tilde{\theta}_j, \tilde{V}_j\}_{j=1}^k$. If this sequence converges, its limit satisfies the following conditions:

$$\check{\theta} = \arg \max_{\theta \in \Theta} n^{-1} \sum_{i=1}^n \ln \Lambda(\theta, \Gamma^q(\theta, \check{V}))(a_i|x_i) \quad \text{and} \quad \check{V} = \Gamma^q(\check{\theta}, \check{V}). \quad (8)$$

Any pair $(\check{\theta}, \check{V})$ that satisfies these two conditions in (8) is called a q -NPL fixed point. The q -NPL estimator, denoted by $(\hat{\theta}_{qNPL}, \hat{V}_{qNPL})$, is defined as the q -NPL fixed point with the highest value of the pseudo likelihood among all the q -NPL fixed points.

We state the regularity conditions for the q -NPL estimator. Let $\nabla^{(s)}f$ denote the s th order derivative of a function f with respect to its all parameters. Let \mathcal{N} denote a closed neighborhood of (θ^0, V^0) , and let \mathcal{N}_{θ^0} denote a closed neighborhood of θ^0 . Let the population counterpart of the objective function be $Q_0(\theta, V) \equiv E \ln \Psi^q(\theta, V)(a_i|x_i)$, and let $\tilde{\theta}_0(V) \equiv \arg \max_{\theta \in \Theta} Q_0(\theta, V)$ and $\phi_0(V) \equiv \Gamma^q(\tilde{\theta}_0(V), V)$. The following assumption is a straightforward counterpart of the assumptions in Aguirregabiria and Mira (2007), henceforth simply AM07.

Assumption 3 (a) *The observations $\{a_i, x_i : i = 1, \dots, n\}$ are independent and identically distributed, and $dF(x) > 0$ for any $x \in X$, where $F(x)$ is the distribution function of x_i .* (b) *$\Psi^q(\theta, V)(a|x) > 0$ for any $(a, x) \in A \times X$ and any $(\theta, V) \in \Theta \times B_V$.* (c) *$\Psi^q(\theta, V)$ is twice continuously differentiable.* (d) *Θ and B_V are compact.* (e) *There is a unique $\theta^0 \in \text{int}(\Theta)$ such that $P^0 = \Psi(\theta^0, V^0)$.* (f) *For any $\theta \neq \theta^0$ and V that solves $V = \Gamma(\theta, V)$, it is the case that $\Psi(\theta, V) \neq P^0$.* (g) *(θ^0, V^0) is an isolated population q -NPL fixed point.* (h) *$\tilde{\theta}_0(V)$ is a single-valued and continuous function of V in a neighborhood of V^0 .* (i) *the operator $\phi_0(V) - V$ has a nonsingular Jacobian matrix at V^0 .*

Assumptions 3(b)(c) imply that $\max_{(a,x) \in A \times X} \sup_{(\theta,V) \in \Theta \times B_V} \|\nabla^{(2)} \ln \Psi(\theta, V)(a|x)\| < \infty$ and hence $E \sup_{(\theta,V) \in \Theta \times B_V} \|\nabla^{(2)} \ln \Psi(\theta, V)(a_i|x_i)\|^r < \infty$ for any positive integer r . Assumption 3(h) corresponds to assumption (iv) in Proposition 2 of AM07. A sufficient condition for Assumption 3(h) is that Q_0 is globally concave in θ in a neighborhood of V^0 and $\nabla_{\theta\theta'} Q_0(\theta, V^0)$ is a nonsingular matrix. Define the Jacobian of Λ and Γ^q evaluated at (θ^0, V^0) as $\Lambda_V \equiv \nabla_{V'} \Lambda(\theta^0, V^0)$ and $\Gamma_V^q \equiv \nabla_{V'} \Gamma^q(\theta^0, V^0)$, where $\nabla_{V'} \equiv (\partial/\partial V')$. Define analogously $\Lambda_\theta \equiv \nabla_{\theta'} \Lambda(\theta^0, V^0)$, and $\Gamma_\theta^q \equiv \nabla_{\theta'} \Gamma^q(\theta^0, V^0)$.

Under Assumption 3, the proof of Proposition 2 of AM07 carries through, and the q -NPL estimator is consistent and asymptotically normally distributed. The asymptotic variance of $\hat{\theta}_{qNPL}$ is given by $\Sigma_{qNPL} = [\Omega_{\theta\theta}^q + \Omega_{\theta V}^q (I - \Gamma_V^q)^{-1} \Gamma_\theta^q]^{-1} \Omega_{\theta\theta}^q \{[\Omega_{\theta\theta}^q + \Omega_{\theta V}^q (I - \Gamma_V^q)^{-1} \Gamma_\theta^q]^{-1}\}'$, where $\Omega_{\theta\theta}^q \equiv E[\nabla_\theta \ln \Psi^q(\theta^0, V^0)(a_i|x_i) \nabla_{\theta'} \ln \Psi^q(\theta^0, V^0)(a_i|x_i)]$ and $\Omega_{\theta V}^q \equiv E[\nabla_\theta \ln \Psi^q(\theta^0, V^0)(a_i|x_i) \times \nabla_{V'} \ln \Psi^q(\theta^0, V^0)(a_i|x_i)]$ and $\Psi^q(\theta, V) \equiv \Lambda(\theta, \Gamma^q(\theta, V))$. As q increases, Σ_{qNPL} approaches to that of the limiting variance of the MLE.⁵

We now analyze the conditions under which the q -NPL algorithm achieves convergence when started from an initial consistent estimate of V^0 , and derive its convergence rate. For matrix and nonnegative scalar sequences of random variables $\{X_n, n \geq 1\}$ and $\{Y_n, n \geq 1\}$, respectively, we write $X_n = O_p(Y_n)(o_p(Y_n))$ if $\|X_n\| \leq CY_n$ for some (all) $C > 0$ with probability arbitrarily close to one for sufficiently large n .

Assumption 4 *Assumption 3 holds. Further, $\tilde{V}_0 - V^0 = o_p(1)$, $\Lambda(\theta, V)$ and $\Gamma(\theta, V)$ are three times continuously differentiable, and $\Omega_{\theta\theta}^q$ is nonsingular.*

⁵The variance of the MLE is given by $(\Omega_{\theta\theta}^\infty)^{-1}$. We may show that $\Sigma_{qNPL} \rightarrow (\Omega_{\theta\theta}^\infty)^{-1}$ as $q \rightarrow \infty$ because $\Gamma_V^q \rightarrow 0$ as $q \rightarrow \infty$.

Define $f_x(x^s) \equiv \Pr(x = x^s)$ for $s = 1, \dots, |X|$, and let f_x be an $L \times 1$ vector of $\Pr(x = x^s)$ whose elements are arranged conformably with $P_{\theta^0}(a^j|x^s)$. Let $\Delta_P \equiv \text{diag}(P^0)^{-1} \text{diag}(f_x)$. The following lemma states the local convergence rate of the q -NPL algorithm and is one of the main results of this paper.

Lemma 1 *Suppose Assumption 4 holds. Then, for $j = 1, \dots, k$,*

$$\begin{aligned}\tilde{\theta}_j - \hat{\theta}_{qNPL} &= O_p(\|\tilde{V}_{j-1} - \hat{V}_{qNPL}\|), \\ \tilde{V}_j - \hat{V}_{qNPL} &= M^q \Gamma_V^q(\tilde{V}_{j-1} - \hat{V}_{qNPL}) + O_p(n^{-1/2} \|\tilde{V}_{j-1} - \hat{V}_{qNPL}\| + \|\tilde{V}_{j-1} - \hat{V}_{qNPL}\|^2),\end{aligned}$$

where $M^q \equiv I - \Gamma_\theta^q((\Lambda_V \Gamma_\theta^q + \Lambda_\theta)' \Delta_P (\Lambda_V \Gamma_\theta^q + \Lambda_\theta))^{-1} (\Lambda_V \Gamma_\theta^q + \Lambda_\theta)' \Delta_P \Lambda_V$.

The convergence property of the q -NPL algorithm depends on the dominant eigenvalue of $M^q \Gamma_V^q$. By Blackwell's sufficient conditions, the Bellman operator Γ is a contraction with modulus β , implying that the dominant eigenvalue of Γ_V^q is at most β^q . For sufficiently large q , therefore, the dominant eigenvalue of $M^q \Gamma_V^q$ is less than one in modulus and the sequence of estimators $\{\tilde{\theta}_j, \tilde{V}_j\}$ converge.

It is possible to reduce the computational burden of implementing the q -NPL algorithm by replacing $\Lambda(\theta, \Gamma^q(\theta, V))$ with its linear approximation around (η, V) , where η is a preliminary estimate of θ . Define $\Psi^q(\theta, V) \equiv \Lambda(\theta, \Gamma^q(\theta, V))$ and let $\Psi^q(\theta, V, \eta)$ be a linear approximation of $\Psi^q(\theta, V)$ around (η, V) :

$$[\Psi^q(\theta, V, \eta)](a|x) \equiv [\Psi^q(\eta, V)](a|x) + \{[\nabla_{\theta'} \Psi^q(\eta, V)](a|x)\}(\theta - \eta). \quad (9)$$

We propose the approximate q -NPL algorithm by replacing $\Psi^q(\theta, V)$ with $\Psi^q(\theta, V, \eta)$ in the first step. Starting from an initial estimate $(\tilde{\theta}_0, \tilde{V}_0)$, the approximate q -NPL algorithm iterates the following steps until $j = k$:

Step 1: Given $(\tilde{\theta}_{j-1}, \tilde{V}_{j-1})$, update θ by $\tilde{\theta}_j = \arg \max_{\theta \in \Theta_j^q} n^{-1} \sum_{i=1}^n \ln \left\{ \left[\Psi^q(\theta, \tilde{V}_{j-1}, \tilde{\theta}_{j-1}) \right] (a_i|x_i) \right\}$, where $\Theta_j^q \equiv \{\theta \in \Theta : \Psi^q(\theta, \tilde{V}_{j-1}, \tilde{\theta}_{j-1})(a|x) \in [c, 1-c] \text{ for all } (a, x) \in A \times X\}$ for an arbitrary small $c > 0$.

Step 2: Update \tilde{V}_{j-1} using the obtained estimate $\tilde{\theta}_j$: $\tilde{V}_j = \Gamma^q(\tilde{\theta}_j, \tilde{V}_{j-1})$.

Implementing Step 1 in the approximate q -NPL algorithm is much less computationally intensive than the original q -NPL algorithm because we may evaluate $\Psi^q(\tilde{\theta}_{j-1}, \tilde{V}_{j-1})$ and $\nabla_{\theta'} \Psi^q(\tilde{\theta}_{j-1}, \tilde{V}_{j-1})$ outside of the optimization routine for θ in Step 1. Using one-sided numerical derivatives, evaluating $\nabla_{\theta'} \Psi^q(\tilde{\theta}_{j-1}, \tilde{V}_{j-1})$ requires the $(K+1)q$ function evaluations of $\Gamma(\theta, V)$ and the $(K+1)$ function evaluations of $\Lambda(\theta, V)$. Once $\Psi^q(\tilde{\theta}_{j-1}, \tilde{V}_{j-1})$ and $\nabla_{\theta'} \Psi^q(\tilde{\theta}_{j-1}, \tilde{V}_{j-1})$ are computed, evaluating $\Psi^q(\theta, \tilde{V}_{j-1}, \tilde{\theta}_{j-1})$ across different values of θ is computationally easy.

The following proposition establishes that the first-order convergence property of the approximate q -NPL algorithm is the same as that of the original q -NPL algorithm.

Assumption 5 (a) *Assumption 4 holds. hold.* (b) *For any $\nu \in \mathbb{R}^K$ such that $\nu \neq 0$, $\nabla_{\theta'} \Psi^q(\theta^0, V^0)(a_i|x_i)\nu \neq 0$ with positive probability.* (c) $\tilde{\theta}_0 - \theta^0 = o_p(1)$.

Proposition 1 *Suppose Assumption 5 holds. Suppose we obtain $\{\tilde{\theta}_j, \tilde{V}_j\}_{j=1}^k$ by the approximate q -NPL algorithm. Then, for $j = 1, \dots, k$,*

$$\begin{aligned}\tilde{\theta}_j - \hat{\theta}_{qNPL} &= O_p(\|\tilde{V}_{j-1} - \hat{V}_{qNPL}\|) + O_p(n^{-1/2}\|\tilde{\theta}_{j-1} - \hat{\theta}_{qNPL}\| + \|\tilde{\theta}_{j-1} - \hat{\theta}_{qNPL}\|^2), \\ \tilde{V}_j - \hat{V}_{qNPL} &= M^q \Gamma_V^q(\tilde{V}_{j-1} - \hat{V}_{qNPL}) + O_p(n^{-1/2}\|\tilde{\theta}_{j-1} - \hat{\theta}_{qNPL}\| + \|\tilde{\theta}_{j-1} - \hat{\theta}_{qNPL}\|^2) \\ &\quad + O_p(n^{-1/2}\|\tilde{V}_{j-1} - \hat{V}_{qNPL}\| + \|\tilde{V}_{j-1} - \hat{V}_{qNPL}\|^2),\end{aligned}$$

where $M^q \equiv I - \Gamma_\theta^q((\Lambda_V \Gamma_\theta^q + \Lambda_\theta)' \Delta_P (\Lambda_V \Gamma_\theta^q + \Lambda_\theta))^{-1} (\Lambda_V \Gamma_\theta^q + \Lambda_\theta)' \Delta_P \Lambda_V$.

Assumption 5(b) is an identification condition for the probability limit of our objective function. It is required because we use an approximation of $\Psi^q(\theta, V)(a|x)$ in the objective function.

3.2 The model with permanent unobserved heterogeneity

Suppose that there are M types of agents, where type m is characterized by a type-specific parameter θ^m , and the population probability of being type m is π^m with $\sum_{m=1}^M \pi^m = 1$. These types capture time-invariant state variables that are unobserved by the researcher. With a slight abuse of notation, denote $\theta = (\theta^1, \dots, \theta^M)' \in \Theta^M$ and $\pi = (\pi^1, \dots, \pi^M)' \in \Theta_\pi$. Then, $\zeta = (\theta', \pi')$ is the parameter to be estimated, and let $\Theta_\zeta = \Theta^M \times \Theta_\pi$ denote the set of possible values of ζ . The true parameter is denoted by ζ^0 .

Consider a panel data set $\{\{a_{it}, x_{it}, x_{i,t+1}\}_{t=1}^T\}_{i=1}^n$ such that $w_i = \{a_{it}, x_{it}, x_{i,t+1}\}_{t=1}^T \in W \equiv (A \times X \times X)^T$ is randomly drawn across i 's from the population. The conditional probability distribution of a_{it} given x_{it} for a type m agent is given by $P_{\theta^m} = \Lambda(\theta^m, V_{\theta^m})$, where V_{θ^m} is a fixed point $V_{\theta^m} = \Gamma(\theta^m, V_{\theta^m})$. To simplify our analysis, we assume that the transition probability function of x_{it} is independent of types and given by $f_x(x_{i,t+1}|a_{it}, x_{it})$ and is known to the researcher. An extension to the case where the transition probability function is also type-dependent is straightforward.

In this framework, the initial state x_{i1} is correlated with unobserved type (i.e., the initial conditions problem of Heckman (1981)). We assume that x_{i1} for type m is randomly drawn from the type m stationary distribution characterized by a fixed point of the following equation: $p^*(x) = \sum_{x' \in X} p^*(x') (\sum_{a' \in A} P_{\theta^m}(a'|x') f_x(x|x', a')) \equiv [T(p^*, P_{\theta^m})](x)$. Since solving the fixed point of $T(\cdot, P)$ for given P is often less computationally intensive than computing the fixed

point of $\Psi(\cdot, \theta)$, we assume the full solution of the fixed point of $T(\cdot, P)$ is available given P .⁶

Let P^m and V^m denote type m 's conditional choice probabilities and type m 's value function, stack the P^m 's and the V^m 's as $\mathbf{P} = (P^1, \dots, P^M)'$ and $\mathbf{V} = (V^1, \dots, V^M)'$, respectively. Let \mathbf{P}^0 and \mathbf{V}^0 denote their true values. Let $\mathbf{\Gamma}(\theta, \mathbf{V}) = (\Gamma(\theta^1, V^1)', \dots, \Gamma(\theta^M, V^M)')$ and let $\mathbf{\Lambda}(\theta, \mathbf{V}) = (\Lambda(\theta^1, V^1)', \dots, \Lambda(\theta^M, V^M)')$. Then, the maximum likelihood estimator for a model with unobserved heterogeneity is:

$$\begin{aligned} \hat{\zeta}_{MLE} &= \arg \max_{\zeta \in \Theta_\zeta} \ln ([L(\pi, \mathbf{P})](w_i)), \\ \text{s.t.} \quad &\mathbf{P} = \mathbf{\Lambda}(\theta, \mathbf{V}), \quad \mathbf{V} = \mathbf{\Gamma}(\theta, \mathbf{V}) \end{aligned} \quad (10)$$

where

$$[L(\pi, \mathbf{P})](w_i) = \sum_{m=1}^M \pi^m p_{P^m}^*(x_{i1}) \prod_{t=1}^T P^m(a_{it}|x_{it}) f_x(x_{i,t+1}|a_{it}, x_{it}),$$

and $p_{P^m}^* = T(p_{P^m}^*, P^m)$ is the type m stationary distribution of x when the conditional choice probability is P^m . If $\mathbf{P}^0 = \mathbf{\Lambda}(\theta^0, \mathbf{V}^0)$ is the true conditional choice probability distribution and π^0 is the true mixing distribution, then $L^0 = L(\pi^0, \mathbf{P}^0)$ represents the true probability distribution of w .

We consider the following sequential algorithm for models with unobserved heterogeneity. Let $\mathbf{\Gamma}^q(\theta, \mathbf{V}) = (\Gamma^q(\theta^1, V^1)', \dots, \Gamma^q(\theta^M, V^M)')$. Define $\Psi^q(\theta^m, V^m) = \Lambda(\theta^m, \Gamma^q(\theta^m, V^m))$ for $m = 1, \dots, M$ and let $\mathbf{\Psi}^q(\theta, \mathbf{V}) = (\Psi^q(\theta^1, V^1)', \dots, \Psi^q(\theta^M, V^M)')$. Assume that an initial consistent estimator $\tilde{\mathbf{V}}_0 = (\tilde{V}_0^1, \dots, \tilde{V}_0^M)'$ is available. For $j = 1, 2, \dots$, iterate

Step 1: Given $\tilde{\mathbf{V}}_{j-1} = (\tilde{V}_{j-1}^1, \dots, \tilde{V}_{j-1}^M)'$, update $\zeta = (\theta', \pi)'$ by

$$\tilde{\zeta}_j = \arg \max_{\zeta \in \Theta_\zeta} n^{-1} \sum_{i=1}^n \ln \left([L(\pi, \mathbf{\Psi}^q(\theta, \tilde{\mathbf{V}}_{j-1}))](w_i) \right).$$

Step 2: Update \mathbf{V} using the obtained estimate $\tilde{\theta}_j$ by $\tilde{\mathbf{V}}_j = \mathbf{\Gamma}^q(\tilde{\theta}_j, \tilde{\mathbf{V}}_{j-1})$ for $m = 1, \dots, M$,

until $j = k$. If iterations converge, its limit satisfies $\hat{\zeta} = \arg \max_{\zeta \in \Theta_\zeta} n^{-1} \sum_{i=1}^n \ln ([L(\pi, \mathbf{\Psi}^q(\theta, \hat{\mathbf{V}}))](w_i))$ and $\hat{\mathbf{V}} = \mathbf{\Gamma}^q(\hat{\theta}, \hat{\mathbf{V}})$. Among the pairs that satisfy these two conditions, the one that maximizes the pseudo likelihood is called the q -NPL estimator, which we denote by $(\hat{\zeta}_{qNPL}, \hat{\mathbf{V}}_{qNPL})$.

Let us introduce the assumptions for the consistency and asymptotic normality of the q -NPL estimator. They are analogous to the assumptions used in Aguirregabiria and Mira (2007). define $Q_0(\zeta, \mathbf{V}) \equiv E \ln ([L(\pi, \mathbf{\Psi}^q(\theta, \mathbf{V}))](w_i))$, $\tilde{\zeta}_0(\mathbf{V}) \equiv \arg \max_{\zeta \in \Theta_\zeta} Q_0(\theta, \mathbf{V})$, and $\phi_0(\mathbf{V}) \equiv$

⁶It is possible to relax the stationarity assumption on the initial states by estimating the type-specific initial distributions of x , denoted by $\{p^{*m}\}_{m=1}^M$, without imposing stationarity restriction in Step 1 of the q -NPL algorithm. In this case, the q -NPL algorithm has the convergence rate similar to that of Proposition 2.

$\Gamma^q(\tilde{\theta}_0(\mathbf{V}), \mathbf{V})$. Define the set of population q -NPL fixed points as $\mathcal{Y}_0 \equiv \{(\theta, \mathbf{V}) \in \Theta \times B_V^M : \zeta = \tilde{\zeta}_0(\mathbf{V}) \text{ and } \mathbf{V} = \phi_0(\mathbf{V})\}$.

Assumption 6 (a) $w_i = \{(a_{it}, x_{it}, x_{i,t+1}) : t = 1, \dots, T\}$ 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 . (b) $[L(\pi, \mathbf{P})](w) > 0$ for any w and for any $(\pi, \mathbf{P}) \in \Theta_\pi \times B_P^M$. (c) $\Lambda(\theta, V)$ and $\Gamma(\theta, V)$ are twice continuously differentiable. (d) Θ_ζ and B_P^M are compact. (e) There is a unique $\zeta^0 \in \text{int}(\Theta_\zeta)$ such that $[L(\pi^0, \mathbf{P}^0)](w) = [L(\pi^0, \Psi(\theta^0, \mathbf{V}^0))](w)$. (f) For any $\zeta \neq \zeta^0$ and \mathbf{V} that solves $\mathbf{V} = \Gamma(\theta, \mathbf{V})$, it is the case that $\Pr(\{w : [L(\pi, \Psi(\theta, \mathbf{V}))](w) \neq L^0(w)\}) > 0$. (g) (ζ^0, \mathbf{V}^0) is an isolated population q -NPL fixed point. (h) $\tilde{\zeta}_0(\mathbf{V})$ is a single-valued and continuous function of \mathbf{V} in a neighborhood of \mathbf{V}^0 . (i) the operator $\phi_0(\mathbf{V}) - \mathbf{V}$ has a nonsingular Jacobian matrix at \mathbf{V}^0 . (j) For any $P \in B_P$, there exists a unique fixed point for $T(\cdot, P)$.

Under Assumption 6, the consistency and asymptotic normality of the q -NPL estimator can be shown by following the proof of Proposition 2 of Aguirregabiria and Mira (2007).

We now establish the convergence property of the q -NPL algorithm for models with unobserved heterogeneity.

Assumption 7 Assumption 6 holds. Further, $\tilde{\mathbf{V}}_0 - \mathbf{V}^0 = o_p(1)$, $\Lambda(\theta, V)$ and $\Gamma(\theta, V)$ are three times continuously differentiable, and $\Omega_{\zeta\zeta}^q$ is nonsingular.

Assumption 7 requires an initial consistent estimator of the value functions. As Aguirregabiria and Mira (2007) argue, if the q -NPL algorithm converges, then the limit may provide a consistent estimate of the parameter ζ even when $\tilde{\mathbf{V}}_0$ is not consistent.

The following proposition states the convergence properties of the q -NPL algorithm for models with unobserved heterogeneity.

Proposition 2 Suppose Assumption 7 holds. Then, for $j = 1, \dots, k$,

$$\begin{aligned} \tilde{\zeta}_j - \hat{\zeta}_{qNPL} &= O_p(\|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}_{qNPL}\|), \\ \tilde{\mathbf{V}}_j - \hat{\mathbf{V}}_{qNPL} &= \mathbf{M}^q \Gamma_V^q(\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}_{qNPL}) + O_p(n^{-1/2} \|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}_{qNPL}\|) + O_p(\|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}_{qNPL}\|^2). \end{aligned}$$

where $\mathbf{M}^q \equiv I - \Gamma_\theta^q D (\Psi_\theta^q)' L_P \Delta_L^{1/2} M_{L_\pi} \Delta_L^{1/2} L_P \Lambda_V$ with $D = ((\Psi_\theta^q)' L_P \Delta_L^{1/2} M_{L_\pi} \Delta_L^{1/2} L_P \Psi_\theta^q)^{-1}$, $M_{L_\pi} \equiv I - \Delta_L^{1/2} L_\pi (L_\pi' \Delta_L L_\pi)^{-1} L_\pi \Delta_L^{1/2}$, and $\Psi_\theta^q \equiv \nabla_{\theta'} \Psi^q(\theta^0, \mathbf{V}^0)$, $\Gamma_\theta^q \equiv \nabla_{\theta'} \Gamma^q(\theta^0, \mathbf{V}^0)$, $\Gamma_V^q \equiv \nabla_{\mathbf{V}'} \Gamma^q(\theta^0, \mathbf{V}^0)$, $\Lambda_V \equiv \nabla_{\mathbf{V}'} \Lambda(\theta^0, \mathbf{V}^0)$, $\Delta_L = \text{diag}((L^0)^{-1})$, $L_P = \nabla_{\mathbf{P}'} L(\pi^0, \mathbf{P}^0)$, and $L_\pi = \nabla_{\pi'} L(\pi^0, \mathbf{P}^0)$.

Since $\Gamma_V^q \rightarrow 0$ as $q \rightarrow \infty$, the algorithm is converging for sufficiently large q .

To reduce the computational cost of implementing the q -NPL algorithm, we may apply the approximate q -NPL algorithm to models with unobserved heterogeneity by replacing $\Psi^q(\theta, V)$ with $\Psi^q(\theta, V, \eta)$ in the first step. Let $\eta = (\eta^1, \dots, \eta^M)'$ be a preliminary estimate of $\theta =$

$(\theta^1, \dots, \theta^M)'$. Let $\Psi^q(\theta, V, \eta) = (\Psi^q(\theta^1, V^1, \eta^1)', \dots, \Psi^q(\theta^M, V^M, \eta^1)')'$, where $\Psi^q(\theta, V, \eta)$ is defined in (9). Assume that initial consistent estimators $\tilde{\theta}_0 = (\tilde{\theta}_0^1, \dots, \tilde{\theta}_0^M)'$ and $\tilde{\mathbf{V}}_0 = (\tilde{V}_0^1, \dots, \tilde{V}_0^M)'$ are available. The approximate q -NPL algorithm iterates the following steps until $j = k$:

Step 1: Given $\tilde{\theta}_{j-1} = (\tilde{\theta}_{j-1}^1, \dots, \tilde{\theta}_{j-1}^M)'$ and $\tilde{\mathbf{V}}_{j-1} = (\tilde{V}_{j-1}^1, \dots, \tilde{V}_{j-1}^M)'$, update $\zeta = (\theta', \pi)'$ by $\tilde{\zeta}_j = \arg \max_{\zeta \in \Theta_{\zeta, j}^q} n^{-1} \sum_{i=1}^n \ln \left([L(\pi, \Psi^q(\theta, \tilde{\mathbf{V}}_{j-1}, \tilde{\theta}_{j-1}))](w_i) \right)$, where $\Theta_{\zeta, j}^q \equiv \{(\pi, \theta) : 0 < \pi^m < 1, [\Psi^q(\theta^m, \tilde{V}_{j-1}^m, \tilde{\theta}_{j-1}^m)](w) \in [c, 1 - c] \text{ for all } w \in W \text{ for } m = 1, \dots, M\}$ for an arbitrary small $c > 0$.

Step 2: Given $(\tilde{\theta}_j, \tilde{\mathbf{V}}_{j-1})$, update \mathbf{V} by $\tilde{\mathbf{V}}_j = \Gamma^q(\tilde{\theta}_j, \tilde{\mathbf{V}}_{j-1})$ for $m = 1, \dots, M$.

The following proposition establishes that the dominant term for the convergence rate of the approximate q -NPL algorithm is the same as that of the q -NPL algorithm for models with unobserved heterogeneity.

Assumption 8 (a) Assumptions 6-7 hold. (b) For any $\nu \in \mathbb{R}^K$ such that $\nu \neq 0$, $\nabla_{\theta'} \Psi^q(\theta^0, V^0)(a_i | x_i) \nu \neq 0$ with positive probability. (c) $\tilde{\theta}_0 - \theta^0 = o_p(1)$.

Proposition 3 Suppose Assumption 8 hold. Suppose we obtain $\{\tilde{\zeta}_j, \tilde{\mathbf{V}}_j\}_{j=1}^k$ by the approximate q -NPL algorithm. Then, for $j = 1, \dots, k$,

$$\begin{aligned} \tilde{\zeta}_j - \hat{\zeta}_{NPL} &= O_p(\|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}_{NPL}\|) + O_p(n^{-1/2} \|\tilde{\zeta}_{j-1} - \hat{\zeta}\| + \|\tilde{\zeta}_{j-1} - \hat{\zeta}\|^2), \\ \tilde{\mathbf{V}}_j - \hat{\mathbf{V}}_{NPL} &= \mathbf{M}^q \Gamma_V^q(\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}_{NPL}) + O_p(n^{-1/2} \|\tilde{\zeta}_{j-1} - \hat{\zeta}\| + \|\tilde{\zeta}_{j-1} - \hat{\zeta}\|^2) \\ &\quad + O_p(n^{-1/2} \|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}_{NPL}\| + \|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}_{NPL}\|^2). \end{aligned}$$

where \mathbf{M}^q is defined in Proposition 2.

4 Dynamic Programming Model with an Equilibrium Constraint

In many dynamic game models and dynamic macroeconomic models, their equilibrium condition is characterized by the solution to the following dual fixed point problems: (i) given the equilibrium probability distribution $P \in B_P$, an agent solves the dynamic programming problem $V = \Gamma(\theta, V, P)$, and (ii) given the solution to the agent's dynamic programming problem V , the probability distribution P satisfies the equilibrium constraint $P = \Lambda(\theta, V, P)$. For instance, in a dynamic game model, P corresponds to the equilibrium strategy and Λ is the best reply mapping. Each player solves the dynamic programming problem given the other players' strategy, $V = \Gamma(\theta, V, P)$, while the equilibrium strategy is a fixed point of the best reply mapping, $P = \Lambda(\theta, V, P)$.

In the following, we extend our sequential estimation algorithm to dynamic programming models with such an equilibrium constraint. We also provide examples of dynamic games and dynamic macro models.

4.1 The Basic Model with an Equilibrium Constraint

As before, an agent maximizes the expected discounted sum of utilities but her utility, the constraint set, and the transition probabilities depend on the equilibrium probability distribution. Importantly, when the agent makes her decision, she treats the equilibrium probability distribution as exogenous: in the dynamic macro model, there are a large number of ex ante identical agents so that each agent's effect on the equilibrium probability distribution is infinitesimal while, in dynamic games, each player treats the other players' strategy as given. Denote the dependence of the equilibrium choice probabilities P on utility function, constraint set, and transition probabilities by the superscript P as $U_\theta^P(a, x, \xi)$, $G_\theta^P(x, \xi)$, and $f_\theta^P(x'|x, a)$. Then, the Bellman equation and the conditional choice probabilities for the agent dynamic optimization problem are written, respectively, as:

$$V(x) = \int \max_{a \in G_\theta^P(x, \xi)} \left\{ U_\theta^P(a, x, \xi) + \beta \sum_X V(x') f_\theta^P(x'|x, a) \right\} g(\xi|x) d\xi \equiv \Gamma(\theta, V, P)(x),$$

and

$$P(a|x) = \int I \left\{ a = \arg \max_{j \in G_\theta^P(x, \xi)} v_\theta(j, x, \xi, V, P) \right\} g_\theta(\xi|x, \epsilon) d\xi \equiv [\Lambda(\theta, V, P)](a|x),$$

where $v_\theta(a, x, \xi, V, P) = U_\theta^P(a, x, \xi) + \beta \sum_X V(x') f_\theta^P(x'|x, a)$ is the choice-specific value function and $I(\cdot)$ is an indicator function.

Consider a cross-sectional data set $\{a_i, x_i\}_{i=1}^n$ where (a_i, x_i) is randomly drawn across i 's from the population. The maximum likelihood estimator (MLE) solves the following constrained maximization problem:

$$\max_{\theta \in \Theta} n^{-1} \sum_{i=1}^n \ln \{P(a_i|x_i)\} \quad \text{subject to} \quad P = \Lambda(\theta, V, P), \quad V = \Gamma(\theta, V, P). \quad (11)$$

Computation of the MLE by the NFXP algorithm requires repeatedly solving all the fixed points of $P = \Lambda(\theta, V, P)$ and $V = \Gamma(\theta, V, P)$ at each parameter value to maximize the objective function with respect to θ . If evaluating the fixed point of $V = \Gamma(\theta, V, P)$ and $P = \Lambda(\theta, V, P)$ is costly, then the MLE is computationally very demanding.

Define $\Psi^q(\theta, V, P) \equiv \Lambda^q(\theta, \Gamma^q(\theta, V, P), P)$, where $\Lambda^q(\theta, V, \cdot)$ is a q -fold operator of Λ defined as

$$\Lambda^q(\theta, V, P) \equiv \underbrace{\Lambda(\theta, V, \Lambda(\theta, V, \dots \Lambda(\theta, V, \Lambda(\theta, V, P)) \dots))}_{q \text{ times}}.$$

Let $Q_0(\theta, V, P) \equiv E \ln \Psi^q(\theta, V, P)(a_i|x_i)$, $\tilde{\theta}_0(V, P) \equiv \arg \max_{\theta \in \Theta} Q_0(\theta, V, P)$, and $\phi_0(V, P) \equiv [\Gamma^q(\tilde{\theta}_0(V, P), V, P), \Psi^q(\tilde{\theta}_0(V, P), V, P)]$.

Assumption 9 (a) The observations $\{a_i, x_i : i = 1, \dots, n\}$ are independent and identically distributed, and $dF(x) > 0$ for any $x \in X$, where $F(x)$ is the distribution function of x_i . (b) $\Psi^q(\theta, V, P)(a|x) > 0$ for any $(a, x) \in A \times X$ and any $(\theta, V, P) \in \Theta \times B_V \times B_P$. (c) $\Psi^q(\theta, V, P)$ is twice continuously differentiable. (d) Θ , B_V , and B_P are compact. (e) There is a unique $\theta^0 \in \text{int}(\Theta)$ such that $P^0 = \Psi(\theta^0, V^0, P^0)$. (f) For any $\theta \neq \theta^0$, (V, P) that solves $V = \Gamma(\theta, V, P)$ and $P = \Lambda(\theta, V, P)$, it is the case that $\Psi(\theta, V, P) \neq P^0$. (g) (θ^0, V^0, P^0) is an isolated population q -NPL fixed point. (h) $\tilde{\theta}_0(V, P)$ is a single-valued and continuous function of V and P in a neighborhood of (V^0, P^0) . (i) the operator $\phi_0(V, P) - (V, P)$ have a nonsingular Jacobian matrix at (V^0, P^0) .

Based on the mapping $\Psi^q(\theta, V, P)$, we propose the following computationally attractive algorithm that does not require repeatedly solving the fixed points of the Bellman operator Γ and the equilibrium mapping Λ . Starting from an initial estimator \tilde{V}^0 and \tilde{P}^0 , iterate the following steps until $j = k$:

Step 1: Given \tilde{V}^{j-1} and \tilde{P}^{j-1} , update θ by $\tilde{\theta}_j = \arg \max_{\theta \in \Theta} n^{-1} \sum_{i=1}^n \ln \left\{ \left[\Psi^q(\theta, \tilde{V}_{j-1}, \tilde{P}_{j-1}) \right] (a_i | x_i) \right\}$.

Step 2: Update \tilde{V}_{j-1} and \tilde{P}_{j-1} using the obtained estimate $\tilde{\theta}_j$: $\tilde{P}_j = \Psi^q(\tilde{\theta}_j, \tilde{V}_{j-1}, \tilde{P}_{j-1})$ and $\tilde{V}_j = \Gamma^q(\tilde{\theta}_j, \tilde{V}_{j-1}, \tilde{P}_{j-1})$.

If the sequence of estimators $\{\tilde{\theta}_j, \tilde{V}_j, \tilde{P}_j\}$ converges, its limit satisfies the conditions:

$$\check{\theta} = \arg \max_{\theta \in \Theta} n^{-1} \sum_{i=1}^n \ln \left[\Psi^q(\theta, \check{V}, \check{P}) \right] (a_i | x_i), \quad \check{P} = \Lambda(\check{\theta}, \check{V}, \check{P}), \quad \text{and} \quad \check{V} = \Gamma(\check{\theta}, \check{V}, \check{P}).$$

Any triplet $(\check{\theta}, \check{V}, \check{P})$ that satisfies the above three conditions is called an q -NPL fixed point. The q -NPL estimator, denoted by $(\hat{\theta}_{qNPL}, \hat{V}_{qNPL}, \hat{P}_{qNPL})$, is defined as the q -NPL fixed point with the highest value of the pseudo likelihood among all the q -NPL fixed points.

Define $\Omega_{\theta\theta}^q \equiv E[\nabla_{\theta} \ln \Psi^q(\theta^0, V^0, P^0)(a_i | x_i) \nabla_{\theta'} \ln \Psi^q(\theta^0, V^0, P^0)(a_i | x_i)]$, $\Omega_{\theta V}^q \equiv E[\nabla_{\theta} \ln \Psi^q(\theta^0, V^0, P^0)(a_i | x_i) \times \nabla_{V'} \ln \Psi^q(\theta^0, V^0, P^0)(a_i | x_i)]$, and $\Omega_{\theta P}^q \equiv E[\nabla_{\theta} \ln \Psi^q(\theta^0, V^0, P^0)(a_i | x_i) \nabla_{P'} \ln \Psi^q(\theta^0, V^0, P^0)(a_i | x_i)]$. Then, the q -NPL estimator $(\hat{\theta}_{qNPL}, \hat{V}_{qNPL}, \hat{P}_{qNPL})$ is consistent (See AM07 for details) and its asymptotic distribution is given by: $\sqrt{n}(\hat{\theta}_{qNPL} - \theta^0) \rightarrow_d N(0, \bar{\Sigma}_{qNPL})$, where $\bar{\Sigma}_{qNPL} = [\Omega_{\theta\theta}^q + S^q]^{-1} \Omega_{\theta\theta}^q \{[\Omega_{\theta\theta}^q + S^q]^{-1}\}'$ with

$$S^q = (\Omega_{\theta V}^q \quad \Omega_{\theta P}^q) \begin{pmatrix} I - \Gamma_V^q & -\Gamma_P^q \\ -\Psi_V^q & I - \Psi_P^q \end{pmatrix}^{-1} \begin{pmatrix} \Gamma_{\theta}^q \\ \Psi_{\theta}^q \end{pmatrix}.$$

As the value of q increases, the variance $\bar{\Sigma}_{qNPL}$ approaches the variance of the MLE when the dominant eigenvalue of Ψ_P is inside the unit circle. Since the computational cost increases with q , there is a trade off in the choice of q between the computational cost and the efficiency of the q -NPL estimator.

The following proposition states the local convergence property of the q -NPL algorithm.

Assumption 10 *Assumption 9 holds. Further, $\tilde{V}_0 - V^0 = o_p(1)$, $\tilde{P}_0 - P^0 = o_p(1)$, $\Lambda(\theta, V, P)$ and $\Gamma(\theta, V, P)$ are three times continuously differentiable, and $\Omega_{\theta\theta}^q$ is nonsingular.*

Proposition 4 *Suppose Assumption 10 holds. Then, for $j = 1, \dots, k$,*

$$\tilde{\theta}_j - \hat{\theta}_{qNPL} = O_p(\|\tilde{V}_j - \hat{V}_{qNPL}\|) + O_p(\|\tilde{P}_j - \hat{P}_{qNPL}\|),$$

$$\begin{pmatrix} \tilde{V}_j - \hat{V}_{qNPL} \\ \tilde{P}_j - \hat{P}_{qNPL} \end{pmatrix} = \begin{pmatrix} \Gamma_V^q - \Gamma_\theta^q(\Omega_{\theta\theta}^q)^{-1}\Omega_{\theta V}^q & \Gamma_P^q - \Gamma_\theta^q(\Omega_{\theta\theta}^q)^{-1}\Omega_{\theta P}^q \\ \Psi_V^q - \Psi_\theta^q(\Omega_{\theta\theta}^q)^{-1}\Omega_{\theta V}^q & \Psi_P^q - \Psi_\theta^q(\Omega_{\theta\theta}^q)^{-1}\Omega_{\theta P}^q \end{pmatrix} \begin{pmatrix} \tilde{V}_{j-1} - \hat{V}_{qNPL} \\ \tilde{P}_{j-1} - \hat{P}_{qNPL} \end{pmatrix} + R_{n,j},$$

where $R_{n,j} = O_p(n^{-1/2}\|\tilde{V}_{j-1} - \hat{V}_{qNPL}\| + \|\tilde{V}_{j-1} - \hat{V}_{qNPL}\|^2) + O_p(n^{-1/2}\|\tilde{P}_{j-1} - \hat{P}_{qNPL}\| + \|\tilde{P}_{j-1} - \hat{P}_{qNPL}\|^2)$.

As $q \rightarrow \infty$, both $\Gamma_V^q - \Gamma_\theta^q(\Omega_{\theta\theta}^q)^{-1}\Omega_{\theta V}^q$ and $\Psi_V^q - \Psi_\theta^q(\Omega_{\theta\theta}^q)^{-1}\Omega_{\theta V}^q$ approach zero. Thus, for sufficiently large q , the convergence property of the q -NPL algorithm is determined by the dominant eigenvalue of $\Psi_P^q - \Psi_\theta^q(\Omega_{\theta\theta}^q)^{-1}\Omega_{\theta P}^q = M^q\Psi_P^q$, where $M^q = I - \Psi_\theta^q(\Omega_{\theta\theta}^q)^{-1}\Psi_\theta^q\Delta_P$ is a projection matrix.

As before, we may reduce the computational burden of implementing the q -NPL algorithm by replacing $\Psi^q(\theta, V, P)$ with its linear approximation around (η, V, P) , where η is a preliminary estimate of θ . Let

$$\Psi^q(\theta, V, P, \eta)(a|x) \equiv [\Psi^q(\eta, V, P)](a|x) + \{[\nabla_{\theta'}\Psi^q(\eta, V, P)](a|x)\}(\theta - \eta).$$

Starting from an initial estimate $(\tilde{\theta}_0, \tilde{V}_0, \tilde{P}_0)$, the approximate q -NPL algorithm iterates the following steps until $j = k$:

Step 1: Given $(\tilde{\theta}_{j-1}, \tilde{V}_{j-1}, \tilde{P}_{j-1})$, update θ by $\tilde{\theta}_j = \arg \max_{\theta \in \Theta_j^q} n^{-1} \sum_{i=1}^n \ln \left\{ \left[\Psi^q(\theta, \tilde{V}_{j-1}, \tilde{P}_{j-1}, \tilde{\theta}_{j-1}) \right] (a_i|x_i) \right\}$, where $\Theta_j^q \equiv \{\theta \in \Theta : \Psi^q(\theta, \tilde{V}_{j-1}, \tilde{P}_{j-1}, \tilde{\theta}_{j-1})(a|x) \in [c, 1-c] \text{ for all } (a, x) \in A \times X\}$ for an arbitrary small $c > 0$.

Step 2: Update \tilde{V}_{j-1} and \tilde{P}_{j-1} using the obtained estimate $\tilde{\theta}_j$: $\tilde{P}_j = \Psi^q(\tilde{\theta}_j, \tilde{V}_{j-1}, \tilde{P}_{j-1})$ and $\tilde{V}_j = \Gamma^q(\tilde{\theta}_j, \tilde{V}_{j-1}, \tilde{P}_{j-1})$.

The repeated evaluations of the objective function in Step 1 across different values of θ is easy because we evaluate $\Psi^q(\tilde{\theta}_{j-1}, \tilde{V}_{j-1}, \tilde{P}_{j-1}, \tilde{\theta}_{j-1})$ and $\nabla_{\theta'}\Psi^q(\tilde{\theta}_{j-1}, \tilde{V}_{j-1}, \tilde{P}_{j-1}, \tilde{\theta}_{j-1})$ outside of the optimization routine. Using one-sided numerical derivatives, evaluating $\nabla_{\theta'}\Psi^q(\tilde{\theta}_{j-1}, \tilde{V}_{j-1}, \tilde{P}_{j-1}, \tilde{\theta}_{j-1})$ requires the $(K+1)q$ function evaluations of $\Gamma(\theta, V, P)$ and the $(K+1)q$ function evaluations of $\Lambda(\theta, V, P)$.

The following proposition shows that the first-order convergence property of the approximate q -NPL algorithm is the same as that of the original q -NPL algorithm.

Assumption 11 (a) Assumption 10 holds. (b) For any $\nu \in \mathbb{R}^K$ such that $\nu \neq 0$, $\nabla_{\theta'} \Psi^q(\theta^0, V^0, P^0)(a_i|x_i)\nu \neq 0$ with positive probability. (c) $\tilde{\theta}_0 - \theta^0 = o_p(1)$.

Proposition 5 Suppose Assumption 11 holds. Suppose we obtain $\{\tilde{\theta}_j, \tilde{V}_j, \tilde{P}_j\}$ by the approximate q -NPL algorithm. Then, for $j = 1, \dots, k$, Then, for $j = 1, \dots, k$,

$$\tilde{\theta}_j - \hat{\theta}_{qNPL} = O_p(\|\tilde{V}_j - \hat{V}_{qNPL}\|) + O_p(\|\tilde{P}_j - \hat{P}_{qNPL}\|) + O_p(n^{-1/2} \|\tilde{\theta}_{j-1} - \hat{\theta}_{qNPL}\| + \|\tilde{\theta}_{j-1} - \hat{\theta}_{qNPL}\|^2),$$

$$\begin{pmatrix} \tilde{V}_j - \hat{V}_{qNPL} \\ \tilde{P}_j - \hat{P}_{qNPL} \end{pmatrix} = \begin{pmatrix} \Gamma_V^q - \Gamma_{\theta}^q(\Omega_{\theta\theta}^q)^{-1}\Omega_{\theta V}^q & \Gamma_P^q - \Gamma_{\theta}^q(\Omega_{\theta\theta}^q)^{-1}\Omega_{\theta P}^q \\ \Psi_V^q - \Psi_{\theta}^q(\Omega_{\theta\theta}^q)^{-1}\Omega_{\theta V}^q & \Psi_P^q - \Psi_{\theta}^q(\Omega_{\theta\theta}^q)^{-1}\Omega_{\theta P}^q \end{pmatrix} \begin{pmatrix} \tilde{V}_{j-1} - \hat{V}_{qNPL} \\ \tilde{P}_{j-1} - \hat{P}_{qNPL} \end{pmatrix} + \bar{R}_{n,j},$$

where $\bar{R}_{n,j} = O_p(n^{-1/2} \|\tilde{\theta}_{j-1} - \hat{\theta}_{qNPL}\| + \|\tilde{\theta}_{j-1} - \hat{\theta}_{qNPL}\|^2) + O_p(n^{-1/2} \|\tilde{V}_{j-1} - \hat{V}_{qNPL}\| + \|\tilde{V}_{j-1} - \hat{V}_{qNPL}\|^2) + O_p(n^{-1/2} \|\tilde{P}_{j-1} - \hat{P}_{qNPL}\| + \|\tilde{P}_{j-1} - \hat{P}_{qNPL}\|^2)$.

4.2 Discrete Dynamic Game

Consider the model of dynamic discrete games studied by Aguirregabiria and Mira (2007), section 3.5. There are I ‘‘global’’ firms competing in N local markets. In market h , a firm i maximizes the expected discounted sum of profits $E[\sum_{s=t}^{\infty} \beta^{s-t} \{\Pi(x_{hs}, a_{hs}, \xi_{his}; \theta)\} | a_{ht}, x_{ht}; \theta]$, where x_{ht} is state variable that is common knowledge for all firms, while ξ_{hit} is state variable that is private information to firm i . The state variable x_{ht} may contain the past choice $a_{h,t-1}$. The researcher observes x_{ht} but not ξ_{hit} . There is no interaction across different markets.

Denote the strategy of firm i by σ^i . Given a set of strategy functions $\sigma = \{\sigma^i(x, \xi_i) : i = 1, \dots, I\}$, the expected behavior of firm i from the viewpoint of the rest of the firms is summarized by the conditional choice probabilities $P^i(a_i|x) = \int 1\{\sigma^i(x, \xi_i) = a_i\} g(\xi_i|x) d\xi_i$, where $g(\xi_i|x)$ is a density function for ξ_i .

Let $P^{-i} = \{P^j : j \neq i\}$ and let $G_i(x, \xi_i; \theta)$ represent a set of feasible choices for firm i when the state is (x, ξ_i) . By assuming that ξ_i 's are iid across firms, the expected profit and the transition probability of x for firm i is given by $U_i^{P^{-i}}(x, a_i, \xi_i; \theta) = \sum_{a_{-i} \in A_{-i}} \left(\prod_{j \neq i} P^j(a_j|x) \right) \Pi(x, a_i, a_{-i}, \xi_i; \theta)$ and $f_i^{P^{-i}}(x'|x, a_i; \theta) = \sum_{a_{-i} \in A_{-i}} \left(\prod_{j \neq i} P^j(a_j|x) \right) f(x'|x, a_i, a_{-i}; \theta)$, respectively. Then, given (θ, P^{-i}) , the Bellman equation and the choice probabilities are

$$V^i(x; \theta) = \int \max_{a_i \in G_i(x, \xi_i; \theta)} \left\{ U_i^{P^{-i}}(x, a_i, \xi_i; \theta) + \beta \sum_{x' \in X} V^i(x') f^{P^{-i}}(x'|x, a_i; \theta) \right\} g(d\xi_i|x) \equiv \Gamma_i(\theta, V^i, P^{-i}),$$

$$P^i(a_i|x; \theta) = \int 1 \left\{ a_i = \arg \max_{a'_i \in G_i(x, \xi_i; \theta)} U_i^{P^{-i}}(x, a'_i, \xi_i; \theta) + \beta \sum_{x' \in X} V^i(x') f^{P^{-i}}(x'|x, a'_i; \theta) \right\} g(d\xi_i|x) \equiv \Lambda_i(\theta, V^i, P^{-i})$$

Let θ^0 denote the true parameter value. Then the true conditional choice probabilities $P^{i,0}$'s and the true value function $V^{i,0}$'s are obtained as the fixed point of the following system of

equations:

$$V^{i,0} = \Gamma_i(\theta^0, V^{i,0}, P^{-i,0}) \quad \text{and} \quad P_i^0 = \Lambda_i(\theta^0, V^{i,0}, P^{-i,0}), \quad (12)$$

for $i = 1, \dots, I$.

Let $\{\{a_{ih}, x_{ih}\}_{i=1}^I\}_{h=1}^n$ be a cross-sectional data set where $\{a_{ih}, x_{ih}\}_{i=1}^I$ is randomly drawn across h 's from the population. Suppose we have an initial estimator of the choice probabilities, $\tilde{P}_0 = (\tilde{P}_0^1, \dots, \tilde{P}_0^I)$. Then, we may obtain a two-step estimator of θ as $\tilde{\theta} = \arg \max_{\theta \in \Theta} n^{-1} \sum_{h=1}^n \sum_{i=1}^I \ln \left\{ \left[\Lambda_i(\theta, V^i, \tilde{P}_0^{-i}) \right] (a_{ih} | x_{ih}) \right\}$ subject to $V^i = \Gamma_i(\theta, V^i, \tilde{P}_0^{-i})$ by solving the fixed point problem $V^i = \Gamma_i(\theta, V^i, \tilde{P}_0^{-i})$ at each parameter value to maximize the log pseudo-likelihood with respect to θ . When the initial estimator \tilde{P}_0 is consistent, this two-step estimator $\tilde{\theta}$ is consistent. Since each agent solves her dynamic programming problem given an estimate of choice probabilities of other agents, $\tilde{P}_0^{-i} = \{\tilde{P}_0^j : j \neq i\}$, the computational cost is similar to that of implementing the NFXP algorithm I times for a single agent problem. We may reduce the computational burden by iterating the following step until $j = k$ starting from an initial value of $\tilde{V}_0 = (\tilde{V}_0^1, \dots, \tilde{V}_0^I)$:

Step 1: Given $\tilde{V}_{j-1} = (\tilde{V}_{j-1}^1, \dots, \tilde{V}_{j-1}^I)$, update θ by

$$\tilde{\theta}_j = \arg \max_{\theta \in \Theta} n^{-1} \sum_{h=1}^n \sum_{i=1}^I \ln \left\{ \left[\Lambda_i(\theta, \Gamma_i^q(\tilde{\theta}_j, \tilde{V}_{j-1}^i, \tilde{P}_0^{-i}), \tilde{P}_0^{-i}) \right] (a_{ih} | x_{ih}) \right\}.$$

Step 2: Update \tilde{V}_{j-1} using the obtained estimate $\tilde{\theta}_j$: $\tilde{V}_j^i = \Gamma_i^q(\tilde{\theta}_j, \tilde{V}_{j-1}^i, \tilde{P}_0^{-i})$, for $i = 1, \dots, I$.

The estimators obtained by the above algorithm may suffer from the finite sample bias when the first-stage estimator \tilde{P}_0 is imprecisely estimated.

To obtain more efficient estimators, we may apply the q -NPL algorithm discussed in the previous section as follows. Define $[\Psi_i^q(\theta, V^i, P^{-i})](a_{ih} | x_{ih}) \equiv [\Lambda_i^q(\theta, \Gamma_i^q(\theta, V^i, P^{-i}), P^{-i})](a_{ih} | x_{ih})$. Starting from an initial estimate \tilde{P}_0 and an initial value of \tilde{V}_0 , the q -NPL algorithm iterates the following step until $j = k$:

Step 1: Given $\tilde{P}_{j-1} = (\tilde{P}_{j-1}^1, \dots, \tilde{P}_{j-1}^I)$ and $\tilde{V}_{j-1} = (\tilde{V}_{j-1}^1, \dots, \tilde{V}_{j-1}^I)$, update $\tilde{\theta}_j$ by

$$\tilde{\theta}_j = \arg \max_{\theta \in \Theta} n^{-1} \sum_{h=1}^n \sum_{i=1}^I \ln \left\{ \left[\Psi_i^q(\theta, \tilde{V}_{j-1}^i, \tilde{P}_{j-1}^{-i}) \right] (a_{ih} | x_{ih}) \right\}.$$

Step 2: Update \tilde{P}_{j-1} and \tilde{V}_{j-1} using the obtained estimate $\tilde{\theta}_j$ as $\tilde{P}_j^i = \Psi_i^q(\tilde{\theta}_j, \tilde{V}_{j-1}^i, \tilde{P}_{j-1}^{-i})$ and $\tilde{V}_j^i = \Gamma_i^q(\tilde{\theta}_j, \tilde{V}_{j-1}^i, \tilde{P}_{j-1}^{-i})$ for $i = 1, \dots, I$.

As Proposition 4 indicates, this q -NPL algorithm converges for sufficiently large value of q when the dominant eigenvalue of $\Psi_{i,P}$ is within a unit circle.

In the q -NPL algorithm, Step 1 is computationally intensive when the evaluations of Λ_i and Γ_i is costly and the value of q is large. We may reduce the computational burden of implementing the approximate q -NPL algorithm that replaces $\Psi_i^q(\theta, V_i, P_{-i})$ with its linear application around (η, V_i, P_{-i}) , $[\Psi_i^q(\theta, V^i, P^{-i})](a|x) \equiv [\Psi_i^q(\eta, V^i, P^{-i})](a|x) + \{[\nabla_{\theta'} \Psi_i^q(\eta, V^i, P^{-i})](a|x)\}(\theta - \eta)$.

That is, we may update θ in Step 1 as $\tilde{\theta}_j = \arg \max_{\theta \in \Theta} n^{-1} \sum_{h=1}^n \sum_{i=1}^I \ln \left\{ \left[\Psi_i^q(\theta, \tilde{V}_{j-1}^i, \tilde{P}_{j-1}^{-i}, \tilde{\theta}_{j-1}) \right] (a_{ih} | x_{ih}) \right\}$. The resulting approximate algorithm is much less computationally intensive than the q -NPL algorithm while its first-order convergence property is the same as that of the q -NPL algorithm (Proposition 5).

5 Monte Carlo Experiments

5.1 Experiment 1: Machine Replacement Model

We consider a version of machine replacement model with unobserved heterogeneity. The observed state variable is machine age denoted by $x_t \in \mathbb{N}_x$, and the unobserved state variables include production shock ϵ_t and choice-specific cost shock $\xi_t = (\xi_t(0), \xi_t(1))'$, where ϵ_t is independently drawn from $N(0, \sigma_\epsilon^2)$ while $\xi_t(a)$'s are independently drawn from Type 1 extreme value distribution. The replacement decision is denoted by $a_t \in \{0, 1\}$ and the transition function of x_t is given by $x_{t+1} = a_t + (1 - a_t)(x_t + 1)$. The profit function is given by $u_\theta(x_t, \epsilon_t, a_t) + \xi(a_t)$, where $u_\theta(x_t, \epsilon_t, a_t) = \exp(\theta_1 x_t(1 - a_t) + \epsilon_t) - \theta_2 a_t$.

We assume that $\theta = (\theta_1, \theta_2)$ is multinomially distributed with the number of support points equal to M , where the m -th type is characterized by a type-specific parameter $\theta^m = (\theta_1^m, \theta_2^m)'$ and the fraction of the m -th type in the population is π^m . We also assume that revenue is observable but with measurement error as: $\ln y_t = \theta_1^m x_t(1 - a_t) + \xi_t + \eta_t$, where η_t is the measurement error and is assumed to be independent of ξ_t and drawn from $N(0, \sigma_\eta^2)$.

The Bellman equation for this firm's dynamic optimization problem is written as

$$V(x, \epsilon) = \gamma + \ln \left(\sum_{a=0}^1 \exp(u_{\theta^m}(x, \epsilon, a) + \beta E_{\epsilon'}[V(x', \epsilon') | x, a]) \right) \equiv [\Gamma(\theta^m, V)](x, \epsilon)$$

while the mapping from the value function to the conditional choice probability is given by

$$[\Lambda(\theta^m, V)](a | x, \epsilon) \equiv \frac{\exp(u_{\theta^m}(x, \epsilon, a) + \beta E_{\epsilon'}[V(x', \epsilon') | x, a])}{\sum_{a'=0}^1 \exp(u_{\theta^m}(x, \epsilon, a') + \beta E_{\epsilon'}[V(x', \epsilon') | x, a'])},$$

where $E_{\epsilon'}[V(x', \epsilon') | x, a] = \int V(a + (1 - a)(x + 1), \epsilon') \phi(\epsilon' / \sigma_\epsilon) / \sigma_\epsilon d\epsilon'$.

In our experiment, we set $M = 2$ and estimate the five structural parameters $\theta \equiv (\theta^1, \theta^2, \pi^1)'$, of which true value is given by $\theta^1 = (-0.3, 4.0)'$, $\theta^2 = (-0.1, 2.0)'$, and $\pi^1 = \pi^2 = 1/2$. We assume that the other parameters in the model are known and common across unobserved types at $(\beta, \sigma_\epsilon, \sigma_\eta) = (0.96, 0.4, 0.2)$.

We generate a panel data set of sample size n with T periods from a parametric model. We first draw types of firms $\{m_i : i = 1, \dots, n\}$ from the multinomial distribution and, then, we draw the initial states $\{(x_{i1}, \xi_{i1}) : i = 1, \dots, n\}$ from the type-specific stationary distributions of (x, ϵ) given θ^{m_i} 's. For firm i , starting from the initial state (x_{i1}, ξ_{i1}) , a_{i1} 's are drawn from the

type-specific conditional choice probabilities $P_{\theta^{m_i}}(a|x_{i1}, \xi_{i1})$ while η_{i1} 's are simulated to generate y_{i1} 's. Then, starting from the initial state (x_{i1}, a_{i1}) , firm i 's time-series data is generated from the model under θ^{m_i} . The data set consists of $\{(x_{it}, y_{it}, a_{it})\}_{t=1}^T : i = 1, \dots, n\}$.⁷

To compute the likelihood, let $w_t = \xi_t + \eta_t$ and define $\sigma_w^2 = \sigma_\epsilon^2 + \sigma_\eta^2$ and $\rho^2 = \sigma_\epsilon^2/\sigma_w^2$. Then, the density of ϵ conditional on w is given by $g(\epsilon|w) = \phi[(\epsilon - \rho^2 w)/(\sigma_\epsilon \sqrt{1 - \rho^2})]/(\sigma_\epsilon \sqrt{1 - \rho^2})$, where $\phi(\cdot)$ is the standard normal density function. Denoting the joint density of ϵ and w by $g(\epsilon, w) = g(\epsilon|w)\phi(w/\sigma_w)/\sigma_w$, firm i 's likelihood contribution is computed by integrating out the unobserved heterogeneity, ϵ 's and θ^m 's, as

$$L(\theta|\{(x_{it}, y_{it}, a_{it})\}_{t=1}^T) = \sum_{m=1}^m \pi^m p_{P_{\theta^m}}^*(x_{i1}) \prod_{t=1}^T P_{\theta^m}(a_{it}|x_{it}, \epsilon') g(\epsilon', \tilde{w}_{it}(\theta^m)) d\epsilon',$$

where $\tilde{w}_{it}(\theta^m) = \ln y_{it} - \theta_1^m x_{it}(1 - a_{it})$ and $p_{P_{\theta^m}}^*(x)$ is the stationary distribution of x implied by the conditional choice probability P_{θ^m} , where $P_{\theta^m} = \Lambda(\theta, V_{\theta^m})$ given the fixed point $V_{\theta^m} = \Gamma(V_{\theta^m}, \theta^m)$.⁸ The maximum likelihood estimator is obtained by maximizing $\sum_{i=1}^n \ln L(\theta|\{(x_{it}, y_{it}, a_{it})\}_{t=1}^T)$.

The q -NPL algorithm is implemented by iterating the following Steps 1 and 2. In Step 1, given \tilde{V}_{j-1}^m for $m = 1, 2$, we update (θ^1, θ^2) by

$$(\tilde{\theta}_j^1, \tilde{\theta}_j^2) = \arg \max_{(\theta^1, \theta^2) \in \Theta^2} n^{-1} \sum_{i=1}^n \ln \left\{ \sum_{m=1}^2 \pi^m p_{\Psi^q(\theta^m, \tilde{V}_{j-1}^m)}^*(x_{i1}) \prod_{t=1}^T \int [\Psi^q(\theta^m, \tilde{V}_{j-1}^m)](a_{it}|x_{it}, \epsilon') g(\epsilon', \tilde{w}_{it}(\theta^m)) d\epsilon' \right\}.$$

Here, $p_{\Psi^q(\theta^m, \tilde{V}_{j-1}^m)}^*(x)$ is the stationary distribution of x when a firm follows the decision rule specified by the choice probabilities $\Psi^q(\theta^m, \tilde{V}_{j-1}^m)$. In Step 2, \tilde{V}_{j-1}^m 's are updated using $\tilde{\theta}_j^m$'s as $\tilde{V}_j^m = \Gamma^q(\tilde{\theta}_j^m, \tilde{V}_{j-1}^m)$ for $m = 1, 2$. The approximate q -NPL algorithm is similarly implemented by replacing $\Psi^q(\theta^m, \tilde{V}_{j-1}^m)$ with its linear approximation around $\theta^m = \tilde{\theta}_{j-1}^m$ in Step 1.

We first examine the finite sample performance of our proposed estimators based on the q -NPL and approximate q -NPL algorithm for $q = 2, 4, 6$, and 8 . We simulate 200 samples, each of which consists of $(n, T) = (400, 5)$ observations. To use the q -NPL algorithm, we set the initial value of the expected value function to zeros. Since applying the approximate q -NPL algorithm also requires the initial estimate of θ , we use the q -NPL algorithm at the initial iteration ($k = 1$) to obtain an initial estimate of (θ, V) , and then we examine the performance of the approximate q -NPL algorithm starting from the second iteration ($k = 2$).

Table 1 reports the bias and the square roots of the mean squared errors. The bias and the

⁷To 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 [cf., Rust (1996)] to evaluate conditional choice probabilities at points outside of the grids. This allows us to generate a sample with continuously distributed ϵ from the approximated model and to evaluate a likelihood function at points outside of the grids. We approximate the state space of ϵ by 10 grid points using the method of Tauchen (1986) while the state space of x is given by $\{1, \dots, 20\}$.

⁸To compute the integral with respect to ϵ given w_i , we approximate the distribution of ϵ conditional on the realized value of w_i for $i = 1, \dots, n$ using Tauchen's method.

mean squared errors of the estimators from the q -NPL algorithm improve with the number of iterations, k , given the value of $q = 2, 4, 6$, and 8 , while they improve with q given the value of k . When k is small, the bias and the mean squared errors of the estimates from the q -NPL algorithm tend to be larger than those of the MLE. The performance of the q -NPL estimators is very similar to that of the q -NPL algorithm across different values of k and q , indicating that our proposed approximation method works in this experiment.

Table 2 reports the average absolute percentage difference between our proposed estimator and the MLE. For both q -NPL estimator and approximate q -NPL estimator, the distance between our proposed estimator and the MLE becomes smaller as k and q increase.

Table 3 shows how the q -NPL estimators after $k = 10$ iterations improve with the sample size across different values of q .

5.2 Experiment 2: Dynamic Game

We apply our proposed method to a dynamic model of entry and exit studied by Aguirregabiria and Mira (2007) and compare its performance with the performance of the original NPL algorithm. The profit of firm i operating in market m in period t is equal to

$$\theta_{RS} \ln S_{mt} - \theta_{RN} \ln(1 + \sum_{j \neq i} a_{jmt}) - \theta_{FC,i} - \theta_{EC}(1 - a_{im,t-1}) + \xi_{imt}(1),$$

whereas its profit is $\xi_{imt}(0)$ if the firm is not operating. We assume that $\{\xi_{imt}(0), \xi_{imt}(1)\}$ follow i.i.d. type I extreme value distribution with zero mean and unit variance, and S_{mt} is the market demand that follows an exogenous first-order Markov process $f_S(S_{m,t+1}|S_{mt})$. We set the number of firms $I = 3$. The state space for the market size S_{mt} is $\{2, 6, 10\}$.⁹ The discount factor is set to $\beta = 0.96$ while we set $(\theta_{RS}, \theta_{RN}, \theta_{EC}) = (1, 1, 1)$. Fixed operating costs are $\theta_{FC,1} = 1.0$, $\theta_{FC,2} = 0.9$, and $\theta_{FC,3} = 0.8$. We compare the performance of the estimators generated by the NPL algorithm of AM07 with those of the estimators generated by the q -NPL and the approximate q -NPL algorithms.

We set $q = 1$ and $q = 2$ in the q -NPL and the approximate q -NPL algorithm. We use a frequency estimator as our initial estimator for P while we set an initial value of V to zero. The sample size is set to $n = 500$. Table 4 presents the bias and the square root of mean squared errors for the AM's NPL estimators together with those for the q -NPL and the approximate q -NPL estimators across different iteration values of $q = 1, 2$, and 3 .

Even for $q = 1$, the overall performance of the q -NPL estimator becomes similar to that of

⁹The transition probability matrix of S_{mt} is given by

$$\begin{bmatrix} 0.8 & 0.2 & 0.0 \\ 0.2 & 0.6 & 0.2 \\ 0.0 & 0.2 & 0.8 \end{bmatrix}.$$

the NPL estimator after $j = 5$ iterations across different values of q . By looking at the bias and the RMSE across different value of iterations, the q -NPL algorithm appears to be largely converged after $j = 10$ iterations. The RMSE at $j = 20$ of the q -NPL estimator improves as the value of q increases from one to four, suggesting that an increase in the value of q leads to an efficiency gain.

The approximate q -NPL algorithm has a convergence problem when $q = 1$. However, The convergence property of the approximate q -NPL algorithm improves as the value of q increases. This is consistent with our analysis on the convergence rate—for small value of q , the (approximate) q -NPL algorithm may not converge unless the dominant eigenvalue of Ψ_P is sufficiently close to zero. For $q = 2$, the performance of the approximate q -NPL algorithm at $j = 20$ iterations is the same as that of the q -NPL algorithm while, for $q = 4$, the approximate q -NPL algorithm converges at $j = 10$ iterations.

6 Proofs

6.1 Proof of Lemma 1

Define $\bar{\psi}^q(\theta, V) \equiv n^{-1} \sum_{i=1}^n \ln \Psi^q(\theta, V)(a_i | x_i)$. With these notations, we may write $\Omega_{\theta\theta}^q = (\Psi_\theta^q)' \Delta_P \Psi_\theta^q$ and $\Omega_{\theta V}^q = (\Psi_\theta^q)' \Delta_P \Psi_V^q$, where $\Psi_\theta^q = \Lambda_V \Gamma_\theta^q + \Lambda_\theta$ and $\Psi_V^q = \Lambda_V \Gamma_V^q$.

First, $\tilde{\theta}_j$ satisfies the first order condition $\nabla_{\theta} \bar{\psi}^q(\tilde{\theta}_j, \tilde{V}_{j-1}) = 0$. Expanding this around $(\hat{\theta}, \hat{V})$ and using $\nabla_{\theta} \bar{\psi}^q(\hat{\theta}, \hat{V}) = 0$ gives

$$0 = \nabla_{\theta\theta'} \bar{\psi}^q(\bar{\theta}, \bar{V})(\tilde{\theta}_j - \hat{\theta}) + \nabla_{\theta V'} \bar{\psi}^q(\bar{\theta}, \bar{V})(\tilde{V}_{j-1} - \hat{V}), \quad (13)$$

where $(\bar{\theta}, \bar{V})$ lie between $(\tilde{\theta}_j, \tilde{V}_{j-1})$ and $(\hat{\theta}, \hat{V})$. It follows from the information matrix equality and the consistency of $(\bar{\theta}, \bar{V})$ that $\nabla_{\theta\theta'} \bar{\psi}^q(\bar{\theta}, \bar{V}) = -\Omega_{\theta\theta}^q + o_p(1)$ and $\nabla_{\theta V'} \bar{\psi}^q(\bar{\theta}, \bar{V}) = -\Omega_{\theta V}^q + o_p(1)$. Since $\Omega_{\theta\theta}^q$ is positive definite, we obtain $\tilde{\theta}_j - \hat{\theta} = O_p(\|\tilde{V}_{j-1} - \hat{V}\|)$, giving the first result.

For the updating equation of V , note that the second derivatives of $\Gamma^q(\theta, V)$ are uniformly bounded in $(\theta, V) \in \Theta \times B_V$ from Assumption. Hence, expanding the right hand side of $\tilde{V}_j = \Gamma^q(\tilde{\theta}_j, \tilde{V}_{j-1})$ twice around $(\hat{\theta}, \hat{V})$ and using $\Gamma^q(\hat{\theta}, \hat{V}) = \hat{V}$, root- n consistency of $(\hat{\theta}, \hat{V})$, and $\tilde{\theta}_j - \hat{\theta} = O_p(\|\tilde{V}_{j-1} - \hat{V}\|)$, we obtain

$$\tilde{V}_j - \hat{V} = \Gamma_{\theta}^q(\tilde{\theta}_j - \hat{\theta}) + \Gamma_V^q(\tilde{V}_{j-1} - \hat{V}) + O_p(n^{-1/2} \|\tilde{V}_{j-1} - \hat{V}\| + \|\tilde{V}_{j-1} - \hat{V}\|^2). \quad (14)$$

Refine (13) as $\tilde{\theta}_j - \hat{\theta} = -\Omega_{\theta\theta}^{-1} \Omega_{\theta V}(\tilde{V}_{j-1} - \hat{V}) + O_p(n^{-1/2} \|\tilde{V}_{j-1} - \hat{V}\| + \|\tilde{V}_{j-1} - \hat{V}\|^2)$ by using $\nabla_{\theta V'} \bar{\psi}^q(\bar{\theta}, \bar{V}) = -\Omega_{\theta V}^q + O_p(\|\tilde{V}_{j-1} - \hat{V}\|) + O_p(n^{-1/2})$ and $\nabla_{\theta\theta'} \bar{\psi}^q(\bar{\theta}, \bar{V}) = -\Omega_{\theta\theta}^q + O_p(\|\tilde{V}_{j-1} - \hat{V}\|) + O_p(n^{-1/2})$. Substituting this into (14) in conjunction with $(\Omega_{\theta\theta}^q)^{-1} \Omega_{\theta V}^q = ((\Lambda_V \Gamma_\theta^q + \Lambda_\theta)' \Delta_P (\Lambda_V \Gamma_\theta^q + \Lambda_\theta))^{-1} (\Lambda_V \Gamma_\theta^q + \Lambda_\theta)' \Delta_P \Lambda_V \Gamma_V^q$ gives the stated result. \square

6.2 Proof of Proposition 1

We suppress the subscript qNPL from $\hat{\theta}_{qNPL}$ and \hat{V}_{qNPL} . Write the objective function as $\bar{\psi}^q(\theta, V, \eta) \equiv n^{-1} \sum_{i=1}^n \ln \Psi^q(\theta, V, \eta)(a_i|x_i)$, and define $\psi^q(\theta, V, \eta) \equiv E \ln \Psi^q(\theta, V, \eta)(a_i|x_i)$. We use induction. Assume $(\tilde{\theta}_{j-1}, \tilde{V}_{j-1}) \rightarrow_p (\theta^0, V^0)$.

First, we prove the consistency, i.e., $(\tilde{\theta}_j, \tilde{V}_j) \rightarrow_p (\theta^0, V^0)$ if $(\tilde{\theta}_{j-1}, \tilde{V}_{j-1}) \rightarrow_p (\theta^0, V^0)$. To show the consistency of $\tilde{\theta}_j$, we show that Θ_j^q is compact and

$$\sup_{(\theta, V, \eta) \in \Theta_j^q \times \mathcal{N}} |\bar{\psi}^q(\theta, V, \eta) - \psi^q(\theta, V, \eta)| = o_p(1), \quad (15)$$

$$\psi^q(\theta, V^0, \theta^0) \text{ is continuous in } \theta, \text{ and } \psi^q(\theta, V^0, \theta^0) \text{ is uniquely maximized at } \theta^0. \quad (16)$$

Then the consistency of $\tilde{\theta}_j$ follows from Theorem 2.1 of Newey and McFadden (1994) because (15) in conjunction with the consistency of $(\tilde{\theta}_{j-1}, \tilde{V}_{j-1})$ and the triangle inequality implies $\sup_{\theta \in \Theta_j^q} |\bar{\psi}^q(\theta, \tilde{V}_{j-1}, \tilde{\theta}_{j-1}) - \psi^q(\theta, V^0, \theta^0)| = o_p(1)$.

Θ_j^q is compact because Θ_j^q is an intersection of the compact set Θ and $|A||X|$ closed sets. Take \mathcal{N} sufficiently small, then it follows from the consistency of $(\tilde{\theta}_{j-1}, \tilde{V}_{j-1})$ and the continuity of $\Psi^q(\theta, V, \eta)$ that $\Psi^q(\theta, V, \eta)(a|x) \in [\epsilon/2, 1 - \epsilon/2]$ for all $(a, x) \in A \times X$ and $(\theta, V, \eta) \in \Theta_j^q \times \mathcal{N}$ with probability approaching one (henceforth wpa1). Observe that (i) $\Theta_j^q \times \mathcal{N}$ is compact, (ii) $\ln \Psi^q(\theta, V, \eta)$ is continuous in $(\theta, V, \eta) \in \Theta_j^q \times \mathcal{N}$, and (iii) $E \sup_{(\theta, V, \eta) \in \Theta_j^q \times \mathcal{N}} |\ln \Psi^q(\theta, V, \eta)(a_i|x_i)| \leq (|\ln(\epsilon/2)| + |\ln(1 - \epsilon/2)|) < \infty$ because of the way we choose \mathcal{N} . Therefore, (15) follows from Lemma 2.4 of Newey and McFadden (1994). Lemma 2.4 of Newey and McFadden (1994) also implies that $\psi^q(\theta, V, \eta)$ is continuous, giving the first part of (16). Finally, we show that θ^0 uniquely maximizes $\psi^q(\theta, V^0, \theta^0)$. Note that

$$\begin{aligned} \psi^q(\theta, V^0, \theta^0) - \psi^q(\theta^0, V^0, \theta^0) &= E \ln(\nabla_{\theta'} \Psi^q(\theta^0, V^0)(\theta - \theta^0) + P^0)(a_i|x_i) - E \ln P^0(a_i|x_i) \\ &= E \ln \left(\frac{\nabla_{\theta'} \Psi^q(\theta^0, V^0)(a_i|x_i)(\theta - \theta^0)}{P^0(a_i|x_i)} + 1 \right). \end{aligned}$$

Recall that $\ln(y + 1) \leq y$ for all $y > -1$ where the inequality is strict if $y \neq 0$, and that Assumption 5(b) implies $\nabla_{\theta'} \Psi^q(\theta^0, V^0)(a_i|x_i)(\theta - \theta^0)/P^0(a_i|x_i) \neq 0$ with positive probability for all $\theta \neq \theta^0$. Therefore, the right hand side of (17) is strictly smaller than

$$E \left[\frac{\nabla_{\theta'} \Psi^q(\theta^0, V^0)(a_i|x_i)(\theta - \theta^0)}{P^0(a_i|x_i)} \right] \quad \text{for all } \theta \neq \theta^0.$$

Because $E[\nabla_{\theta'} \Psi^q(\theta^0, V^0)(a_i|x_i)/P^0(a_i|x_i)] = 0$, we have $\psi^q(\theta, V^0, \theta^0) - \psi^q(\theta^0, V^0, \theta^0) < 0$ for all $\theta \neq \theta^0$, and θ^0 uniquely maximizes $\psi^q(\theta, V^0, \theta^0)$. Therefore, $\tilde{\theta}_j \rightarrow_p \theta^0$. Finally, $\tilde{V}_j \rightarrow_p V^0$ follows from $\Gamma^q(\tilde{\theta}_j, \tilde{V}_{j-1}) \rightarrow_p \Gamma^q(\theta^0, V^0) = V^0$, and we establish the consistency of $(\tilde{\theta}_j, \tilde{V}_j)$.

We proceed to derive the stated representation of $\tilde{\theta}_j - \hat{\theta}$ and $\tilde{V}_j - \hat{V}$. Expanding the first

order condition $0 = \nabla_{\theta} \psi^q(\tilde{\theta}_j, \tilde{V}_{j-1}, \tilde{\theta}_{j-1})$ twice around $(\hat{\theta}, \hat{V}_{j-1}, \hat{\theta}_{j-1})$ gives

$$0 = \nabla_{\theta} \bar{\psi}^q(\hat{\theta}, \tilde{V}_{j-1}, \tilde{\theta}_{j-1}) + \nabla_{\theta\theta'} \bar{\psi}^q(\hat{\theta}, \tilde{V}_{j-1}, \tilde{\theta}_{j-1})(\tilde{\theta}_j - \hat{\theta}) + O_p(\|\tilde{\theta}_j - \hat{\theta}\|^2), \quad (17)$$

Note that the q -NPL estimator satisfies $\nabla_{\theta} \bar{\psi}^q(\hat{\theta}, \hat{V}, \hat{\theta}) = 0$, and that $\Psi^q(\theta^0, V^0, \theta^0) = \Psi^q(\theta^0, V^0)$, $\nabla_{\theta'} \Psi^q(\theta^0, V^0, \theta^0) = \nabla_{\theta'} \Psi^q(\theta^0, V^0)$, $\nabla_{V'} \Psi^q(\theta^0, V^0, \theta^0) = \nabla_{V'} \Psi^q(\theta^0, V^0)$, and $\nabla_{\eta'} \Psi^q(\theta^0, V^0, \theta^0) = 0$. Therefore, expanding $\nabla_{\theta} \bar{\psi}^q(\hat{\theta}, \tilde{V}_{j-1}, \tilde{\theta}_{j-1})$ twice around $(\hat{\theta}, \hat{V}, \hat{\theta})$ and using the root- n consistency of $(\hat{\theta}, \hat{V})$ and the information matrix equality, we obtain $\nabla_{\theta} \bar{\psi}^q(\hat{\theta}, \tilde{V}_{j-1}, \tilde{\theta}_{j-1}) = -\Omega_{\theta V}^q(\tilde{V}_j - \hat{V}) + r_{nj}$, where r_{nj} denotes a reminder term of $O_p(n^{-1/2} \|\tilde{\theta}_{j-1} - \hat{\theta}\| + \|\tilde{\theta}_{j-1} - \hat{\theta}\|^2 + n^{-1/2} \|\tilde{V}_{j-1} - \hat{V}\| + \|\tilde{V}_{j-1} - \hat{V}\|^2)$. Then the stated bound of $\tilde{\theta}_j - \hat{\theta}$ follows from (17) by noting that $\nabla_{\theta\theta'} \psi^q(\hat{\theta}, \tilde{V}_{j-1}, \tilde{\theta}_{j-1}) = -\Omega_{\theta\theta}^q + o_p(1)$.

For the updating equation of V , expanding $\nabla_{\theta\theta'} \bar{\psi}^q(\hat{\theta}, \tilde{V}_{j-1}, \tilde{\theta}_{j-1})$ around $(\hat{\theta}, \hat{V}, \hat{\theta})$ in (17) and using the bound of $\tilde{\theta}_j - \hat{\theta}$ obtained above gives $\tilde{\theta}_j - \hat{\theta} = -(\Omega_{\theta\theta}^q)^{-1} \Omega_{\theta V}^q(\tilde{V}_j - \hat{V}) + r_{nj}$. Substituting this into the right hand side of $\tilde{V}_j - \hat{V} = \Gamma_{\theta}^q(\tilde{\theta}_j - \hat{\theta}) + \Gamma_V^q(\tilde{V}_{j-1} - \hat{V}) + r_{nj}$ and noting $(\Omega_{\theta\theta}^q)^{-1} \Omega_{\theta V}^q = ((\Psi_{\theta}^q)' \Delta_P \Psi_{\theta}^q)^{-1} (\Psi_{\theta}^q)' \Delta_P \Psi_V^q = ((\Lambda_V \Gamma_{\theta}^q + \Lambda_{\theta})' \Delta_P (\Lambda_V \Gamma_{\theta}^q + \Lambda_{\theta}))^{-1} (\Lambda_V \Gamma_{\theta}^q + \Lambda_{\theta})' \Delta_P \Lambda_V \Gamma_V^q$ gives the stated result. \square

6.3 Proof of Proposition 2

We suppress the subscript qNPL from $\hat{\zeta}_{qNPL}$ and $\hat{\mathbf{V}}_{qNPL}$. The proof follows the proof of Lemma 1. Let $l^q(\zeta, \mathbf{V})(w) \equiv \ln(L(\pi, \Psi^q(\theta, \mathbf{V}))(w))$. Define $\bar{l}_{\zeta}^q(\zeta, \mathbf{V}) = n^{-1} \sum_{i=1}^n \nabla_{\zeta} l^q(\zeta, \mathbf{V})(w_i)$, $\bar{l}_{\zeta\zeta}^q(\zeta, \mathbf{V}) = n^{-1} \sum_{i=1}^n \nabla_{\zeta\zeta'} l^q(\zeta, \mathbf{V})(w_i)$, and $\bar{l}_{\zeta\mathbf{V}}^q(\zeta, \mathbf{V}) = n^{-1} \sum_{i=1}^n \nabla_{\zeta\mathbf{V}'} l^q(\zeta, \mathbf{V})(w_i)$. Expanding the first order condition $\bar{l}_{\zeta}^q(\tilde{\zeta}_j, \tilde{\mathbf{V}}_{j-1}) = \bar{l}_{\zeta}^q(\hat{\zeta}, \hat{\mathbf{V}}) = 0$ gives

$$\tilde{\zeta}_j - \hat{\zeta} = -\bar{l}_{\zeta\zeta}^q(\bar{\zeta}, \bar{\mathbf{V}})^{-1} \bar{l}_{\zeta\mathbf{V}}^q(\bar{\zeta}, \bar{\mathbf{V}})(\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}) = O_p(\|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}\|), \quad (18)$$

where $(\bar{\zeta}, \bar{\mathbf{V}})$ is between $(\tilde{\zeta}_j, \tilde{\mathbf{V}}_{j-1})$ and $(\hat{\zeta}, \hat{\mathbf{V}})$. This gives the bound for $\tilde{\zeta}_j - \hat{\zeta}$. Rewriting this further using Assumption 7 gives

$$\tilde{\zeta}_j - \hat{\zeta} = -(\Omega_{\zeta\zeta}^q)^{-1} \Omega_{\zeta\mathbf{V}}^q(\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}) + O_p(n^{-1/2} \|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}\|) + O_p(\|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}\|^2), \quad (19)$$

where $\Omega_{\zeta\zeta}^q = E[\nabla_{\zeta} l^q(\zeta^0, \mathbf{V}^0)(w_i) \nabla_{\zeta'} l^q(\zeta^0, \mathbf{V}^0)(w_i)]$ and $\Omega_{\zeta\mathbf{V}}^q = E[\nabla_{\zeta} l^q(\zeta^0, \mathbf{V}^0)(w_i) \nabla_{\mathbf{V}'} l^q(\zeta^0, \mathbf{V}^0)(w_i)]$. On the other hand, expanding the second step equation $\tilde{\mathbf{V}}_j = \Gamma^q(\tilde{\zeta}_j, \tilde{\mathbf{V}}_{j-1})$ twice around $(\hat{\zeta}, \hat{\mathbf{V}})$, using the root- n consistency of $(\hat{\zeta}, \hat{\mathbf{V}})$ and (18) give

$$\tilde{\mathbf{V}}_j - \hat{\mathbf{V}} = \Gamma_V^q(\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}) + \Gamma_{\zeta}^q(\tilde{\zeta}_j - \hat{\zeta}) + O_p(n^{-1/2} \|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}\|) + O_p(\|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}\|^2), \quad (20)$$

where $\Gamma_{\zeta}^q \equiv \nabla_{\zeta'} \Gamma^q(\theta^0, \mathbf{V}^0) = [\Gamma_{\theta}^q, \mathbf{0}]$. Substituting (19) into (20) gives

$$\tilde{\mathbf{V}}_j - \hat{\mathbf{V}} = [\Gamma_V^q - \Gamma_{\zeta}^q (\Omega_{\zeta\zeta}^q)^{-1} \Omega_{\zeta\mathbf{V}}^q](\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}) + O_p(n^{-1/2} \|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}\|) + O_p(\|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}\|^2).$$

Note that $\Omega_{\zeta\zeta}^q$ and $\Omega_{\zeta V}^q$ are written as

$$\Omega_{\zeta\zeta}^q = \begin{bmatrix} \Omega_{\theta\theta}^q & \Omega_{\theta\pi}^q \\ \Omega_{\pi\theta}^q & \Omega_{\pi\pi}^q \end{bmatrix} = \begin{bmatrix} (\Psi_\theta^q)' L'_P \Delta_L L_P \Psi_\theta^q & (\Psi_\theta^q)' L'_P \Delta_L L_\pi \\ L'_\pi \Delta_L L_P \Psi_\theta & L'_\pi \Delta_L L_\pi \end{bmatrix},$$

$$\Omega_{\zeta V}^q = \begin{bmatrix} \Omega_{\theta V}^q \\ \Omega_{\pi V}^q \end{bmatrix} = \begin{bmatrix} (\Psi_\theta^q)' L'_P \Delta_L L_P \Psi_V^q \\ L'_\pi \Delta_L L_P \Psi_V^q \end{bmatrix},$$

and

$$(\Omega_{\zeta\zeta}^q)^{-1} = \begin{bmatrix} D & -D\Omega_{\theta\pi}^q(\Omega_{\pi\pi}^q)^{-1} \\ -(\Omega_{\pi\pi}^q)^{-1}\Omega_{\pi\theta}^q D & (\Omega_{\pi\pi}^q)^{-1} + (\Omega_{\pi\pi}^q)^{-1}\Omega_{\pi\theta}^q D\Omega_{\theta\pi}^q(\Omega_{\pi\pi}^q)^{-1} \end{bmatrix},$$

where $D = ((\Psi_\theta^q)' L'_P \Delta_L^{1/2} M_{L_\pi} \Delta_L^{1/2} L_P \Psi_\theta^q)^{-1}$ with $M_{L_\pi} = I - \Delta_L^{1/2} L_\pi (L'_\pi \Delta_L L_\pi)^{-1} L_\pi \Delta_L^{1/2}$.

Then, using $\Gamma_\zeta^q = [\Gamma_\theta^q, \mathbf{0}]$ and $\Psi_V^q = \Lambda_V \Gamma_V^q$ gives $\Gamma_\zeta^q (\Omega_{\zeta\zeta}^q)^{-1} \Omega_{\zeta P}^q = \Gamma_\theta^q D (\Psi_\theta^q)' L'_P \Delta_L^{1/2} M_{L_\pi} \Delta_L^{1/2} L_P \Lambda_V \Gamma_V^q$, and the stated result follows. \square

6.4 Proof of Proposition 3

We suppress the subscript qNPL from $\hat{\zeta}_{qNPL}$ and $\hat{\mathbf{V}}_{qNPL}$. Let $l^q(\zeta, \mathbf{V}, \eta)(w) \equiv \ln(L(\pi, \Psi^q(\theta, \mathbf{V}, \eta))(w))$. Define $\bar{l}_\zeta^q(\zeta, \mathbf{V}, \eta) = n^{-1} \sum_{i=1}^n \nabla_\zeta l^q(\zeta, \mathbf{V}, \eta)(w_i)$, $\bar{l}_{\zeta\zeta}^q(\zeta, \mathbf{V}, \eta) = n^{-1} \sum_{i=1}^n \nabla_{\zeta\zeta'} l^q(\zeta, \mathbf{V}, \eta)(w_i)$, and $\bar{l}_{\zeta\mathbf{V}}^q(\zeta, \mathbf{V}, \eta) = n^{-1} \sum_{i=1}^n \nabla_{\zeta\mathbf{V}'} l^q(\zeta, \mathbf{V}, \eta)(w_i)$. Note that the q -NPL estimator satisfies $\nabla_{\theta'} \bar{l}_\zeta^q(\hat{\zeta}, \hat{\mathbf{V}}, \hat{\theta}) = 0$ and that $\Psi^q(\zeta^0, \mathbf{V}^0, \theta^0) = \Psi^q(\zeta^0, \mathbf{V}^0)$, $\nabla_{\zeta'} \Psi^q(\zeta^0, \mathbf{V}^0, \theta^0) = \nabla_{\zeta'} \Psi^q(\zeta^0, \mathbf{V}^0)$, $\nabla_{\mathbf{V}'} \Psi^q(\zeta^0, \mathbf{V}^0, \theta^0) = \nabla_{\mathbf{V}'} \Psi^q(\zeta^0, \mathbf{V}^0)$, and $\nabla_{\eta'} \Psi^q(\zeta^0, \mathbf{V}^0, \theta^0) = 0$.

The consistency of $(\tilde{\zeta}_j, \tilde{\mathbf{V}}_j)$ for $j = 1, \dots, k$ can be shown by following the proof of Proposition 1 and, thus, its proof is omitted.

Expanding the first order condition $0 = \bar{l}_\zeta^q(\tilde{\zeta}_j, \tilde{\mathbf{V}}_{j-1}, \tilde{\theta}_{j-1})$ twice around $(\hat{\zeta}, \tilde{\mathbf{V}}_{j-1}, \tilde{\theta}_{j-1})$ gives

$$\begin{aligned} 0 &= \nabla_{\zeta'} \bar{l}_\zeta^q(\hat{\zeta}, \tilde{\mathbf{V}}_{j-1}, \tilde{\theta}_{j-1}) + \nabla_{\zeta\zeta'} \bar{l}_\zeta^q(\hat{\zeta}, \tilde{\mathbf{V}}_{j-1}, \tilde{\theta}_{j-1})(\tilde{\zeta}_j - \hat{\zeta}) + O_p(\|\tilde{\zeta}_j - \hat{\zeta}\|^2) \\ &= \nabla_{\zeta'} \bar{l}_\zeta^q(\hat{\zeta}, \tilde{\mathbf{V}}_{j-1}, \tilde{\theta}_{j-1}) + \left[-\Omega_{\zeta\zeta}^q + O_p(n^{-1/2} + \|\tilde{\zeta}_{j-1} - \hat{\zeta}\| + \|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}\|) \right] (\tilde{\zeta}_j - \hat{\zeta}) + O_p(\|\tilde{\zeta}_j - \hat{\zeta}\|^2), \end{aligned} \tag{21}$$

where the second equality follows from expanding $\nabla_{\zeta\zeta'} \bar{l}_\zeta^q(\hat{\zeta}, \tilde{\mathbf{V}}_{j-1}, \tilde{\theta}_{j-1})$ around $(\hat{\zeta}, \hat{\mathbf{V}}, \hat{\theta})$ and using the root- n consistency of $(\hat{\zeta}, \hat{\mathbf{V}})$ and the information matrix equality. Furthermore, expanding $\nabla_{\zeta'} \bar{l}_\zeta^q(\hat{\zeta}, \tilde{\mathbf{V}}_{j-1}, \tilde{\theta}_{j-1})$ twice around $(\hat{\zeta}, \hat{\mathbf{V}}, \hat{\theta})$ and using the root- n consistency of $(\hat{\zeta}, \hat{\mathbf{V}})$ and the information matrix equality, we obtain $\nabla_{\zeta'} \bar{l}_\zeta^q(\hat{\zeta}, \tilde{\mathbf{V}}_{j-1}, \tilde{\theta}_{j-1}) = -\Omega_{\zeta\mathbf{V}}^q(\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}) + r_{nj}$, where r_{nj} denotes a reminder term of $O_p(n^{-1/2} \|\tilde{\zeta}_{j-1} - \hat{\zeta}\| + \|\tilde{\zeta}_{j-1} - \hat{\theta}\|^2 + n^{-1/2} \|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}\| + \|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}\|^2)$. Then, the bound for $\tilde{\zeta}_j - \hat{\zeta}$ follows from writing the second and third terms on the right of (21) as $(-\Omega_{\zeta\zeta}^q + o_p(1))(\tilde{\zeta}_j - \hat{\zeta})$ and using the positive definiteness of $\Omega_{\zeta\zeta}^q$.

For the bound of $\tilde{\mathbf{V}}_j - \hat{\mathbf{V}}$, expanding $\tilde{\mathbf{V}}_j = \Gamma^q(\tilde{\zeta}_j, \tilde{\mathbf{V}}_{j-1})$ twice around $(\hat{\zeta}, \hat{\mathbf{V}})$, using the root- n

consistency of $(\hat{\zeta}, \hat{\mathbf{V}})$ and the bound for $\tilde{\zeta}_j - \hat{\zeta}$ give

$$\tilde{\mathbf{V}}_j - \hat{\mathbf{V}} = \mathbf{\Gamma}_V^q(\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}) + \mathbf{\Gamma}_\zeta^q(\tilde{\zeta}_j - \hat{\zeta}) + O_p(n^{-1/2}\|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}\| + \|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}\|^2). \quad (22)$$

On the other hand, it follows from $\tilde{\zeta}_j - \hat{\zeta} = O_p(\|\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}\|)$ and (21) that $\tilde{\zeta}_j - \hat{\zeta} = -(\Omega_{\zeta\zeta}^q)^{-1}\Omega_{\zeta\mathbf{V}}^q(\tilde{\mathbf{V}}_{j-1} - \hat{\mathbf{V}}) + r_{nj}$. Substituting this into (22) and repeating the argument of Proposition 2 give the stated bound of $\tilde{\mathbf{V}}_j - \hat{\mathbf{V}}$. \square

6.5 Proof of Proposition 4

Define $\bar{\psi}^q(\theta, V, P) \equiv n^{-1} \sum_{i=1}^n \ln \Psi^q(\theta, V, P)(a_i|x_i)$. First, $\tilde{\theta}_j$ satisfies the first order condition $\nabla_{\theta} \bar{\psi}^q(\tilde{\theta}_j, \tilde{V}_{j-1}, \tilde{P}_{j-1}) = 0$. Expanding this around $(\hat{\theta}, \hat{V}, \hat{P})$ and using $\nabla_{\theta} \bar{\psi}^q(\hat{\theta}, \hat{V}, \hat{P}) = 0$ gives

$$0 = \nabla_{\theta\theta} \bar{\psi}^q(\bar{\theta}, \bar{V}, \bar{P})(\tilde{\theta}_j - \hat{\theta}) + \nabla_{\theta V} \bar{\psi}^q(\bar{\theta}, \bar{V}, \bar{P})(\tilde{V}_{j-1} - \hat{V}) + \nabla_{\theta P} \bar{\psi}^q(\bar{\theta}, \bar{V}, \bar{P})(\tilde{P}_{j-1} - \hat{P}), \quad (23)$$

where $(\bar{\theta}, \bar{V}, \bar{P})$ lie between $(\tilde{\theta}_j, \tilde{V}_{j-1}, \tilde{P}_{j-1})$ and $(\hat{\theta}, \hat{V}, \hat{P})$. It follows from the information matrix equality and the consistency of $(\bar{\theta}, \bar{V}, \bar{P})$ that $\nabla_{\theta\theta} \bar{\psi}(\bar{\theta}, \bar{V}, \bar{P}) = -\Omega_{\theta\theta}^q + o_p(1)$, $\nabla_{\theta V} \bar{\psi}(\bar{\theta}, \bar{V}, \bar{P}) = -\Omega_{\theta V}^q + o_p(1)$, and $\nabla_{\theta P} \bar{\psi}(\bar{\theta}, \bar{V}, \bar{P}) = -\Omega_{\theta P}^q + o_p(1)$. Since $\Omega_{\theta\theta}^q$ is positive definite, we obtain $\tilde{\theta}_j - \hat{\theta} = O_p(\|\tilde{V}_{j-1} - \hat{V}\| + \|\tilde{P}_{j-1} - \hat{P}\|)$, giving the first result.

For the updating equation of V and P , note that the second derivatives of $\Gamma^q(\theta, V, P)$ and $\Psi^q(\theta, V, P)$ are uniformly bounded in $(\theta, V, P) \in \Theta \times B_V \times B_P$ from Assumption. Hence, expanding the right hand sides of $\tilde{V}_j = \Gamma^q(\tilde{\theta}_j, \tilde{V}_{j-1}, \tilde{P}_{j-1})$ and $\tilde{P}_j = \Psi^q(\tilde{\theta}_j, \tilde{V}_{j-1}, \tilde{P}_{j-1})$ twice around $(\hat{\theta}, \hat{V}, \hat{P})$ and using $\Gamma^q(\hat{\theta}, \hat{V}, \hat{P}) = \hat{V}$, $\Psi^q(\hat{\theta}, \hat{V}, \hat{P}) = \hat{P}$, root- n consistency of $(\hat{\theta}, \hat{V}, \hat{P})$, and $\tilde{\theta}_j - \hat{\theta} = O_p(\|\tilde{V}_{j-1} - \hat{V}\| + \|\tilde{P}_{j-1} - \hat{P}\|)$, we obtain

$$\tilde{V}_j - \hat{V} = \Gamma_{\theta}^q(\tilde{\theta}_j - \hat{\theta}) + \Gamma_V^q(\tilde{V}_{j-1} - \hat{V}) + \Gamma_P^q(\tilde{P}_{j-1} - \hat{P}) + R_{n,j} \quad (24)$$

$$\tilde{P}_j - \hat{P} = \Psi_{\theta}^q(\tilde{\theta}_j - \hat{\theta}) + \Psi_V^q(\tilde{V}_{j-1} - \hat{V}) + \Psi_P^q(\tilde{P}_{j-1} - \hat{P}) + R_{n,j} \quad (25)$$

where $R_{n,j}$ is a generic reminder term of $O_p(n^{-1/2}\|\tilde{V}_{j-1} - \hat{V}\| + n^{-1/2}\|\tilde{P}_{j-1} - \hat{P}\| + \|\tilde{V}_{j-1} - \hat{V}\|^2 + \|\tilde{P}_{j-1} - \hat{P}\|^2)$. Refine (23) as $\tilde{\theta}_j - \hat{\theta} = -\Omega_{\theta\theta}^{-1}\Omega_{\theta V}(\tilde{V}_{j-1} - \hat{V}) - \Omega_{\theta\theta}^{-1}\Omega_{\theta P}(\tilde{P}_{j-1} - \hat{P}) + R_{n,j}$. Substituting this into (24)-(25) gives the stated result. \square

6.6 Proof of Proposition 5

Define $\bar{\psi}^q(\theta, V, P, \eta) \equiv n^{-1} \sum_{i=1}^n \ln \tilde{\Psi}^q(\theta, V, P, \eta)(a_i|x_i)$ and $\psi^q(\theta, V, P, \eta) \equiv E \ln \tilde{\Psi}^q(\theta, V, P, \eta)(a_i|x_i)$.

We first show the consistency of $(\tilde{\theta}_j, \tilde{V}_j, \tilde{P}_j)$ for all $j = 1, 2, \dots, k$. We use induction. Assume $(\tilde{\theta}_{j-1}, \tilde{V}_{j-1}, \tilde{P}_{j-1}) \rightarrow_p (\theta^0, V^0, P^0)$. In order to show $\tilde{\theta}_j \rightarrow_p \theta^0$, it suffices to show that (15)–(16) in the proof of Proposition 1 hold if we replace $\bar{\psi}^q(\theta, V, \eta)$ and $\psi(\theta, V, \eta)$ with $\bar{\psi}^q(\theta, V, P, \eta)$ and $\psi^q(\theta, V, P, \eta)$. Let \mathcal{N}_0 be a closed neighborhood of (V^0, P^0, θ^0) and take \mathcal{N}_0 sufficiently small, then (i) $\Theta_j^q \times \mathcal{N}_0$ is compact, (ii) $\ln \tilde{\Psi}^q(\theta, V, P, \eta)$ is continuous in $(\theta, V, P, \eta) \in \Theta_j^q \times \mathcal{N}_0$, and (iii)

$E \sup_{(\theta, V, P, \eta) \in \Theta_j^q \times \mathcal{N}_0} |\ln \Psi^q(\theta, V, P, \eta)(a_i|x_i)| < \infty$. Therefore, (15) and the first result of (16) hold for $\bar{\psi}^q(\theta, V, P, \eta)$ and $\psi^q(\theta, V, P, \eta)$.

To show that θ^0 uniquely maximizes $\psi^q(\theta, V^0, P^0, \theta^0)$, note that

$$\begin{aligned} \psi^q(\theta, V^0, P^0, \theta^0) - \psi^q(\theta^0, V^0, P^0, \theta^0) &= E \ln \left(\frac{\nabla_{\theta'} \Psi^q(\theta^0, V^0, P^0)(a_i|x_i)(\theta - \theta^0)}{P^0(a_i|x_i)} + 1 \right) \\ &< E \left[\frac{\nabla_{\theta'} \Psi^q(\theta^0, V^0, P^0)(a_i|x_i)(\theta - \theta^0)}{P^0(a_i|x_i)} \right] \end{aligned}$$

for all $\theta \neq \theta^0$, where the last inequality follows from Assumption 11(b) and the inequality $\ln(y+1) > y$ for all $y > -1$ when $y \neq 0$. It follows from $E[\nabla_{\theta'} \Psi^q(\theta^0, V^0, P^0)(a_i|x_i)/P^0(a_i|x_i)] = 0$ that we have $\psi^q(\theta, V^0, P^0, \theta^0) - \psi^q(\theta^0, V^0, P^0, \theta^0) < 0$ for all $\theta \neq \theta^0$, and the second result of (16) hold for $\bar{\psi}^q(\theta, V, P, \eta)$ and $\psi^q(\theta, V, P, \eta)$. Therefore, $\tilde{\theta}_j \rightarrow_p \theta^0$. Finally, $\tilde{V}_j \rightarrow_p V^0$ and $\tilde{P}_j \rightarrow_p P^0$ follows from $\Gamma^q(\tilde{\theta}_j, \tilde{V}_{j-1}, \tilde{P}_{j-1}) \rightarrow_p \Gamma^q(\theta^0, V^0, P^0) = V^0$ and $\Psi^q(\tilde{\theta}_j, \tilde{V}_{j-1}, \tilde{P}_{j-1}) \rightarrow_p \Psi^q(\theta^0, V^0, P^0) = V^0$, and we establish the consistency of $(\tilde{\theta}_j, \tilde{V}_j, \tilde{P}_j)$.

We proceed to show the updating equation of θ, V and P . Expanding the first order condition $0 = \nabla_{\theta} \bar{\psi}^q(\tilde{\theta}_j, \tilde{V}_{j-1}, \tilde{P}_{j-1}, \tilde{\theta}_{j-1})$ twice around $(\hat{\theta}, \tilde{V}_{j-1}, \tilde{P}_{j-1}, \tilde{\theta}_{j-1})$ gives

$$0 = \nabla_{\theta} \bar{\psi}^q(\hat{\theta}, \tilde{V}_{j-1}, \tilde{P}_{j-1}, \tilde{\theta}_{j-1}) + \nabla_{\theta\theta'} \bar{\psi}^q(\hat{\theta}, \tilde{V}_{j-1}, \tilde{P}_{j-1}, \tilde{\theta}_{j-1})(\tilde{\theta}_j - \hat{\theta}) + O_p(\|\tilde{\theta}_j - \hat{\theta}\|^2). \quad (26)$$

Second, note that the approximate q -NPL estimator satisfies $\nabla_{\theta} \bar{\psi}^q(\hat{\theta}, \hat{V}, \hat{P}, \hat{\theta}) = 0$, and that $\Psi^q(\theta^0, V^0, P^0, \theta^0) = \Psi^q(\theta^0, V^0, P^0)$, $\nabla_{\theta'} \Psi^q(\theta^0, V^0, P^0, \theta^0) = \nabla_{\theta'} \Psi^q(\theta^0, V^0, P^0)$, $\nabla_{V'} \Psi^q(\theta^0, V^0, P^0, \theta^0) = \nabla_{V'} \Psi^q(\theta^0, V^0, P^0)$, $\nabla_{P'} \Psi^q(\theta^0, V^0, P^0, \theta^0) = \nabla_{P'} \Psi^q(\theta^0, V^0, P^0)$, and $\nabla_{\eta'} \Psi^q(\theta^0, V^0, P^0, \theta^0) = 0$. Therefore, expanding $\nabla_{\theta} \bar{\psi}^q(\hat{\theta}, \tilde{V}_{j-1}, \tilde{P}_{j-1}, \tilde{\theta}_{j-1})$ twice around $(\hat{\theta}, \hat{V}, \hat{P}, \hat{\theta})$ and using the root- n consistency of $(\hat{\theta}, \hat{V}, \hat{P})$ and the information matrix equality, we obtain $\nabla_{\theta} \bar{\psi}^q(\hat{\theta}, \tilde{V}_{j-1}, \tilde{P}_{j-1}, \tilde{\theta}_{j-1}) = -\Omega_{\theta V}^q(\tilde{V}_j - \hat{V}) - \Omega_{\theta P}^q(\tilde{P}_j - \hat{P}) + r_{nj}$, where r_{nj} denotes a reminder term of $O_p(n^{-1/2} \|\tilde{\theta}_{j-1} - \hat{\theta}\| + \|\tilde{\theta}_{j-1} - \hat{\theta}\|^2 + n^{-1/2} \|\tilde{V}_{j-1} - \hat{V}\| + \|\tilde{V}_{j-1} - \hat{V}\|^2 + n^{-1/2} \|\tilde{P}_{j-1} - \hat{P}\| + \|\tilde{P}_{j-1} - \hat{P}\|^2)$. Then the stated bound of $\tilde{\theta}_j - \hat{\theta}$ follows from (26) by noting that $\nabla_{\theta\theta'} \bar{\psi}^q(\hat{\theta}, \tilde{V}_{j-1}, \tilde{P}_{j-1}, \tilde{\theta}_{j-1}) = -\Omega_{\theta\theta}^q + o_p(1)$.

For the updating equation of V and P , expanding $\nabla_{\theta\theta'} \bar{\psi}^q(\hat{\theta}, \tilde{V}_{j-1}, \tilde{P}_{j-1}, \tilde{\theta}_{j-1})$ around $(\hat{\theta}, \hat{V}, \hat{P}, \hat{\theta})$ in (26) and using the bound of $\tilde{\theta}_j - \hat{\theta}$ obtained above gives $\tilde{\theta}_j - \hat{\theta} = -(\Omega_{\theta\theta}^q)^{-1} \Omega_{\theta V}^q(\tilde{V}_j - \hat{V}) - (\Omega_{\theta\theta}^q)^{-1} \Omega_{\theta P}^q(\tilde{P}_j - \hat{P}) + r_{nj}$. Substituting this into the right hand side of $\tilde{V}_j - \hat{V} = \Gamma_{\theta}^q(\tilde{\theta}_j - \hat{\theta}) + \Gamma_V^q(\tilde{V}_{j-1} - \hat{V}) + \Gamma_P^q(\tilde{P}_{j-1} - \hat{P}) + r_{nj}$ and $\tilde{P}_j - \hat{P} = \Psi_{\theta}^q(\tilde{\theta}_j - \hat{\theta}) + \Psi_V^q(\tilde{V}_{j-1} - \hat{V}) + \Psi_P^q(\tilde{P}_{j-1} - \hat{P}) + r_{nj}$ gives the stated result. \square

References

- Aguirregabiria, V. (1999). "The dynamics of markups and inventories in retailing firms." *Review of Economic Studies* 66: 275-308.
- 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.
- Aguirregabiria, V. and P. Mira (2007). "Sequential estimation of dynamic discrete games." *Econometrica* 75(1): 1-53.
- Arcidiacono, P. and R. A. Miller (2008). CCP estimation of dynamic discrete choice models with unobserved heterogeneity. Mimeographed, Duke university.
- Bajari, P., Benkard, C. L., and Levin, J. (2007). "Estimating dynamic models of imperfect competition." *Econometrica* 75(5): 1331-1370.
- Bajari, P. and H. Hong (2006). Semiparametric estimation of a dynamic game of incomplete information. NBER Technical Working Paper 320.
- Collard-Wexler, A. (2006) Demand fluctuations and plant turnover in the Ready-Mix concrete industry. Mimeographed, NYU.
- 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.
- Gilleskie, D. (1998). "A dynamic stochastic model of medical care use and work absence." *Econometrica* 66: 1-45.
- Gotz, G. A. and J. J. McCall (1980). "Estimation in sequential decisionmaking models: a methodological note." *Economics Letters* 6(2): 131-136.
- Heckman, J. (1981) "The incidental parameter problem and the problem of initial conditions in estimating a discrete time-discrete data stochastic process," in *Structural Analysis of Discrete Data with Econometric Applications*, ed. by C. Manski and D. McFadden. Cambridge: MIT Press.
- Hotz, J. and R. A. Miller (1993). "Conditional choice probabilities and the estimation of dynamic models." *Review of Economic Studies* 60: 497-529.
- Kasahara, H. (2009) "Temporary Increases in Tariffs and Investment: The Chilean Case," *Journal of Business and Economic Statistics*, 27(1): 113-127.
- Kasahara, H. and B. Lapham (2008). "Productivity and the Decision to Import and Export: Theory and Evidence." CESifo Working Paper No. 2240.

- Kasahara, H. and K. Shimotsu (2008a) “Pseudo-likelihood Estimation and Bootstrap Inference for Structural Discrete Markov Decision Models,” *Journal of Econometrics*, 146: 92-106.
- Kasahara, H. and K. Shimotsu (2008b) Sequential Estimation of Structural Models with a Fixed Point Constraint. Mimeographed, University of Western Ontario.
- Miller, R. (1984). “Job matching and occupational choice.” *Journal of Political Economy* 92: 1086-1120.
- Pakes, A. (1986). “Patents as options: some estimates of the value of holding European patent stocks.” *Econometrica* 54: 755-784.
- Pakes, A., M. Ostrovsky, and S. Berry (2007). “Simple estimators for the parameters of discrete dynamic games (with entry/exit examples).” *RAND Journal of Economics* 38(2): 373-399.
- Pesendorfer, M. and P. Schmidt-Dengler (2008). “Asymptotic least squares estimators for dynamic games,” *Review of Economic Studies*, 75, 901-928.
- Rothwell, G. and J. Rust (1997). “On the optimal lifetime of nuclear power plants.” *Journal of Business and Economic Statistics* 15(2): 195-208.
- Rust, J. (1987). “Optimal replacement of GMC bus engines: an empirical model of Harold Zurcher.” *Econometrica* 55(5): 999-1033.
- Rust, J. (1988). “Maximum likelihood estimation of discrete control processes.” *SIAM Journal of Control and Optimization* 26(5): 1006-1024.
- Rust, J. (1996). “Numerical dynamic programming in economics.” in *Handbook of Computational Economics* eds. H. Amman, D. Kendrick and J. Rust. Elsevier, North Holland.
- Wolpin, K. (1984) “An Estimable Dynamic Stochastic Model of Fertility and Child Mortality.” *Journal of Political Economy* 92: 852-874.

Table 1: Performance of q -NPL and approximate q -NPL estimator

			MLE	q -NPL				approximate q -NPL			
				q=2	q=4	q=6	q=8	q=2	q=4	q=6	q=8
θ_1^1	Bias	k=1	0.0059	0.0061	0.0060	0.0059	0.0059	0.0061	0.0060	0.0059	0.0059
		k=3	0.0059	0.0056	0.0058	0.0059	0.0059	0.0056	0.0058	0.0059	0.0059
		k=5	0.0059	0.0056	0.0058	0.0059	0.0059	0.0056	0.0058	0.0059	0.0059
		k=10	0.0059	0.0056	0.0058	0.0059	0.0059	0.0056	0.0058	0.0059	0.0059
	\sqrt{MSE}	k=1	0.0076	0.0077	0.0077	0.0076	0.0076	0.0077	0.0077	0.0076	0.0076
		k=3	0.0076	0.0074	0.0075	0.0076	0.0076	0.0074	0.0075	0.0076	0.0076
		k=5	0.0076	0.0074	0.0075	0.0076	0.0076	0.0074	0.0075	0.0076	0.0076
		k=10	0.0076	0.0074	0.0075	0.0076	0.0076	0.0074	0.0075	0.0076	0.0076
θ_2^1	Bias	k=1	-0.6887	-1.3293	-0.8081	-0.6748	-0.6774	-1.3293	-0.8081	-0.6748	-0.6774
		k=3	-0.6887	-0.7294	-0.6866	-0.6875	-0.6886	-0.7262	-0.6870	-0.6874	-0.6886
		k=5	-0.6887	-0.7002	-0.6867	-0.6875	-0.6886	-0.6998	-0.6868	-0.6874	-0.6886
		k=10	-0.6887	-0.7014	-0.6868	-0.6875	-0.6886	-0.7014	-0.6868	-0.6874	-0.6886
	\sqrt{MSE}	k=1	0.6949	1.3307	0.8118	0.6809	0.6839	1.3307	0.8118	0.6809	0.6839
		k=3	0.6949	0.7343	0.6927	0.6937	0.6947	0.7313	0.6931	0.6936	0.6947
		k=5	0.6949	0.7058	0.6929	0.6937	0.6947	0.7054	0.6930	0.6936	0.6947
		k=10	0.6949	0.7070	0.6929	0.6937	0.6947	0.7070	0.6930	0.6936	0.6947
θ_1^2	Bias	k=1	0.0034	0.0028	0.0033	0.0035	0.0035	0.0028	0.0033	0.0035	0.0035
		k=3	0.0034	0.0032	0.0034	0.0034	0.0034	0.0032	0.0034	0.0034	0.0034
		k=5	0.0034	0.0032	0.0034	0.0034	0.0034	0.0032	0.0034	0.0034	0.0034
		k=10	0.0034	0.0032	0.0034	0.0034	0.0034	0.0032	0.0034	0.0034	0.0034
	\sqrt{MSE}	k=1	0.0067	0.0064	0.0066	0.0067	0.0067	0.0064	0.0066	0.0067	0.0067
		k=3	0.0067	0.0066	0.0066	0.0067	0.0067	0.0066	0.0066	0.0067	0.0067
		k=5	0.0067	0.0066	0.0066	0.0067	0.0067	0.0066	0.0066	0.0067	0.0067
		k=10	0.0067	0.0066	0.0066	0.0067	0.0067	0.0066	0.0066	0.0067	0.0067
θ_2^2	Bias	k=1	-0.1568	-0.4253	-0.2114	-0.1552	-0.1530	-0.4253	-0.2114	-0.1552	-0.1530
		k=3	-0.1568	-0.1615	-0.1554	-0.1569	-0.1569	-0.1649	-0.1554	-0.1569	-0.1569
		k=5	-0.1568	-0.1505	-0.1554	-0.1569	-0.1569	-0.1507	-0.1554	-0.1569	-0.1569
		k=10	-0.1568	-0.1515	-0.1554	-0.1569	-0.1569	-0.1515	-0.1554	-0.1569	-0.1569
	\sqrt{MSE}	k=1	0.1878	0.4318	0.2319	0.1868	0.1852	0.4318	0.2319	0.1868	0.1852
		k=3	0.1878	0.1903	0.1867	0.1879	0.1879	0.1929	0.1867	0.1879	0.1879
		k=5	0.1878	0.1823	0.1868	0.1879	0.1879	0.1825	0.1867	0.1879	0.1879
		k=10	0.1878	0.1830	0.1868	0.1879	0.1879	0.1830	0.1867	0.1879	0.1879
π^1	Bias	k=1	0.0314	0.0374	0.0324	0.0312	0.0313	0.0374	0.0324	0.0312	0.0313
		k=3	0.0314	0.0315	0.0312	0.0314	0.0314	0.0310	0.0312	0.0314	0.0314
		k=5	0.0314	0.0313	0.0312	0.0314	0.0314	0.0313	0.0312	0.0314	0.0314
		k=10	0.0314	0.0313	0.0312	0.0314	0.0314	0.0313	0.0312	0.0314	0.0314
	\sqrt{MSE}	k=1	0.0478	0.0521	0.0484	0.0476	0.0477	0.0521	0.0484	0.0476	0.0477
		k=3	0.0478	0.0478	0.0476	0.0477	0.0478	0.0475	0.0476	0.0477	0.0478
		k=5	0.0478	0.0477	0.0476	0.0477	0.0478	0.0477	0.0476	0.0477	0.0478
		k=10	0.0478	0.0477	0.0476	0.0477	0.0478	0.0477	0.0476	0.0477	0.0478

Notes: Based on 200 simulated samples, each of which consists of $(n, T) = (400, 5)$ observations.

Table 2: Convergence of q -NPL and approximate q -NPL estimator to MLE

		q -NPL				approximate q -NPL			
		q=2	q=4	q=6	q=8	q=2	q=4	q=6	q=8
θ_1^1	k=1	0.0005	0.0002	0.0001	0.0000	0.0005	0.0002	0.0001	0.0000
	k=3	0.0010	0.0003	0.0000	0.0000	0.0010	0.0003	0.0000	0.0000
	k=5	0.0010	0.0003	0.0000	0.0000	0.0010	0.0003	0.0000	0.0000
	k=10	0.0010	0.0003	0.0000	0.0000	0.0010	0.0003	0.0000	0.0000
θ_2^1	k=1	0.1602	0.0299	0.0035	0.0028	0.1602	0.0299	0.0035	0.0028
	k=3	0.0102	0.0005	0.0003	0.0000	0.0094	0.0004	0.0003	0.0000
	k=5	0.0029	0.0005	0.0003	0.0000	0.0028	0.0005	0.0003	0.0000
	k=10	0.0032	0.0005	0.0003	0.0000	0.0032	0.0005	0.0003	0.0000
θ_1^2	k=1	0.0067	0.0009	0.0005	0.0002	0.0067	0.0009	0.0005	0.0002
	k=3	0.0025	0.0006	0.0001	0.0000	0.0025	0.0006	0.0001	0.0000
	k=5	0.0025	0.0006	0.0001	0.0000	0.0025	0.0006	0.0001	0.0000
	k=10	0.0025	0.0006	0.0001	0.0000	0.0025	0.0006	0.0001	0.0000
θ_2^2	k=1	0.1342	0.0273	0.0009	0.0019	0.1342	0.0273	0.0009	0.0019
	k=3	0.0031	0.0007	0.0001	0.0001	0.0042	0.0007	0.0001	0.0001
	k=5	0.0034	0.0007	0.0001	0.0001	0.0033	0.0007	0.0001	0.0001
	k=10	0.0030	0.0007	0.0001	0.0001	0.0030	0.0007	0.0001	0.0001
π^1	k=1	0.0120	0.0019	0.0005	0.0003	0.0120	0.0019	0.0005	0.0003
	k=3	0.0004	0.0005	0.0001	0.0000	0.0009	0.0004	0.0001	0.0000
	k=5	0.0004	0.0005	0.0001	0.0000	0.0004	0.0004	0.0001	0.0000
	k=10	0.0004	0.0005	0.0001	0.0000	0.0004	0.0004	0.0001	0.0000

Notes: The reported values, for instance, are the average of $|(\hat{\theta}_{1,q\text{-NPL}}^{1,k} - \hat{\theta}_{1,\text{MLE}}^1)/\theta_1^1|$ across 200 replications.

Table 3: Performance of q -NPL estimator at $k = 10$ for $(n, T) = (200, 5), (400, 5),$ and $(800, 5)$

θ_1^1	MLE	q -NPL			
		q=2	q=4	q=6	q=8
$(n, T) = (200, 5)$	0.0089	0.0088	0.0089	0.0089	0.0089
$(n, T) = (400, 5)$	0.0076	0.0074	0.0075	0.0076	0.0076
$(n, T) = (800, 5)$	0.0067	0.0065	0.0066	0.0067	0.0067
θ_2^1	MLE	q=2	q=4	q=6	q=8
$(n, T) = (200, 5)$	0.6887	0.7005	0.6868	0.6875	0.6886
$(n, T) = (400, 5)$	0.6949	0.7070	0.6929	0.6937	0.6947
$(n, T) = (800, 5)$	0.6853	0.6975	0.6833	0.6841	0.6851
θ_1^2	MLE	q=2	q=4	q=6	q=8
$(n, T) = (200, 5)$	0.0090	0.0090	0.0090	0.0090	0.0090
$(n, T) = (400, 5)$	0.0067	0.0066	0.0066	0.0067	0.0067
$(n, T) = (800, 5)$	0.0052	0.0050	0.0051	0.0052	0.0052
θ_2^2	MLE	q=2	q=4	q=6	q=8
$(n, T) = (200, 5)$	0.2111	0.2070	0.2105	0.2113	0.2111
$(n, T) = (400, 5)$	0.1878	0.1830	0.1868	0.1879	0.1879
$(n, T) = (800, 5)$	0.1827	0.1774	0.1815	0.1828	0.1828
π^1	MLE	q=2	q=4	q=6	q=8
$(n, T) = (200, 5)$	0.0607	0.0605	0.0606	0.0607	0.0607
$(n, T) = (400, 5)$	0.0478	0.0477	0.0476	0.0477	0.0478
$(n, T) = (800, 5)$	0.0389	0.0389	0.0388	0.0389	0.0389

Table 4: Bias and RMSE for $j = 1, 3, 5, 10,$ and $20.$

Estimation of θ_{RS}														
	AM-NPL		q-NPL						approximate q-NPL					
	Bias	RMSE	q=1		q=2		q=4		q=1		q=2		q=4	
			Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE
j=1	-0.1437	0.2307	-0.0870	0.2034	0.8542	0.8955	0.1226	0.4139	-0.0870	0.2034	0.8542	0.8955	0.1226	0.4139
j=3	-0.0127	0.1979	-0.0113	0.1997	0.0566	0.2669	-0.0093	0.2051	-0.0498	0.1837	0.1529	0.3114	0.0089	0.2325
j=5	-0.0096	0.2062	-0.0095	0.2067	-0.0082	0.2067	-0.0118	0.2018	-0.0419	0.1750	0.0043	0.2113	-0.0110	0.2021
j=10	-0.0115	0.2036	-0.0111	0.2044	-0.0112	0.2025	-0.0119	0.2018	0.0045	0.2378	-0.0116	0.2029	-0.0118	0.2018
j=20	-0.0114	0.2038	-0.0110	0.2045	-0.0112	0.2025	-0.0119	0.2018	0.0003	0.2390	-0.0112	0.2025	-0.0119	0.2018

Estimation of θ_{RN}														
	AM-NPL		q-NPL						approximate q-NPL					
	Bias	RMSE	q=1		q=2		q=4		q=1		q=2		q=4	
			Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE
j=1	-0.3655	0.6001	-0.3749	0.6120	2.4167	2.5143	0.3453	1.1264	-0.3749	0.6120	2.4167	2.5143	0.3453	1.1264
j=3	-0.0210	0.5281	-0.0235	0.5324	0.1834	0.7295	-0.0154	0.5474	-0.1546	0.4894	0.6133	1.1089	0.0447	0.6378
j=5	-0.0181	0.5489	-0.0181	0.5500	-0.0128	0.5509	-0.0251	0.5347	-0.1116	0.4558	0.1029	0.7653	-0.0201	0.5378
j=10	-0.0238	0.5407	-0.0228	0.5430	-0.0236	0.5364	-0.0252	0.5346	0.0214	0.6373	-0.0179	0.5506	-0.0251	0.5346
j=20	-0.0236	0.5411	-0.0226	0.5433	-0.0236	0.5363	-0.0252	0.5346	0.0141	0.6496	-0.0236	0.5363	-0.0252	0.5346

The result is based on 500 simulated samples.