

Nested Pseudo Likelihood Estimation of Continuous-Time Dynamic Discrete Games

Jason R. Blevins and Minhae Kim*

The Ohio State University

August 3, 2021

Abstract

We introduce a sequential estimator for continuous time dynamic discrete choice models (single-agent models and games) by adapting the nested pseudo likelihood (NPL) estimator of Aguirregabiria and Mira (2002, 2007), developed for discrete time models with discrete time data, to the continuous time case with data sampled either discretely (i.e., uniformly-spaced snapshot data) or continuously. We establish conditions for consistency and asymptotic normality of the estimator, a local convergence condition, and, for single agent models, a zero Jacobian property assuring local convergence. We carry out a series of Monte Carlo experiments using an entry-exit game with five heterogeneous firms to confirm the large-sample properties and demonstrate finite-sample bias reduction via iteration. In our simulations we show that the convergence issues documented for the NPL estimator in discrete time models are less likely to affect comparable continuous-time models. We also show that there can be large bias in economically-relevant parameters, such as the competitive effect and entry cost, from estimating a misspecified discrete time model when in fact the data generating process is a continuous time model.

Keywords: continuous time, dynamic discrete games, dynamic discrete choice, nested pseudo likelihood, entry games

JEL Classification: C13, C35, C73

*Department of Economics, Email: blevins.141@osu.edu and kim.6488@osu.edu

1 Introduction

This paper introduces a new sequential estimator for continuous-time dynamic discrete choice models. In *dynamic* models, as compared to static models, forward-looking agents make decisions based on expected future payoffs each period. *Discrete choice* models focus on agents’ discrete choices such as a firm’s entry-exit decision or a worker’s retirement. Dynamic discrete choice models have been rigorously studied in labor economics and empirical industrial organization since they were pioneered by Miller (1984), Wolpin (1984), Pakes (1986), and Rust (1987). They have been used subsequently in the areas of consumer inventory (Song and Chintagunta, 2003; Hendel and Nevo, 2006), dynamic demand of durable goods (Schiraldi, 2011; Gowrisankaran and Rysman, 2012), and firm entry and exit (Benkard, 2004; Holmes, 2011).¹ These examples, however, are all based on discrete time models.

More recently, authors have worked with *continuous time* dynamic discrete choice models due to their flexibility in separating the timing of actions in the model from the timing of data sampling and to their computational advantages in large-scale games. In such models, the state variables change in a stochastic, sequential manner which helps avoid the “curse of dimensionality” arising from the need to compute players’ expectations over future states (Doraszelski and Judd, 2012).² Arcidiacono, Bayer, Blevins, and Ellickson (2016)—henceforth ABBE (2016)—showed that continuous time models remain empirically tractable, can be estimated with continuous or discrete time data, can better approximate the economic reality of certain industries, and that the computational advantages can allow researchers to solve and simulate counterfactuals in large-scale models. Blevins (2016) presents theoretical, computational and econometric properties of continuous time dynamic discrete choice games and extends results from ABBE (2016). Since then, CCP estimation in dynamic discrete choice models has been applied in various areas, including the transportation industry (Mazur, 2017; Qin, 2017), supermarket industry (Schiraldi et al., 2012), TV advertising (Deng and Mela, 2018), nightlife (Cosman, 2019), online games (Nevskaya and Albuquerque, 2019), and job search (Arcidiacono et al., 2019).

The goal of this paper is to introduce a nested pseudo likelihood (NPL) estimator for continuous time models, which we refer to as the CTNPL estimator. ABBE (2016) developed a two-step estimator for continuous time discrete choice models, but here we extend this estimator using insights from the discrete time NPL estimator of Aguirregabiria and Mira (2002,

¹Doraszelski and Pakes (2007) and Aguirregabiria and Mira (2010) provide the survey of applied works in dynamic discrete choice models in discrete time.

²In discrete time models all state variables change at the same time, so if there are n state variables (i.e., one for each firm) and each state variable can change to one of κ alternative values, the number of possible future states is κ^n . In contrast, in continuous time models only one state variable can change at any given instant, meaning that there are only $\kappa(n - 1)$ possible future states for a player to consider.

2007) to iteratively impose the equilibrium conditions from the game to improve estimates of the structural parameters. Aguirregabiria and Mira (2002) originally developed the NPL algorithm and estimator in a single agent, discrete time setting and Aguirregabiria and Mira (2007) extended the algorithm and estimator to discrete time dynamic discrete choice games. Bugni and Bunting (2021) introduced the K -MD estimator, a sequential version of the minimum-distance estimator proposed by Pesendorfer and Schmidt-Dengler (2008), based on the NPL mapping.

In this paper we introduce the CTNPL estimator, derive the large sample properties, and evaluate its performance in the setting of a familiar dynamic oligopoly model. Compared to the two-step pseudo maximum likelihood (PML) estimator, the NPL estimator has a number of advantages. First, it does not require consistent estimates for the initial conditional choice probabilities. Second, the estimator performs better in finite sample as nonparametric estimates of choice probabilities (CCPs) can be imprecise, leading two-step estimators to be biased. This paper shows that these advantages also carry over to the CTNPL estimator for continuous time models. We also compare continuous time and discrete time models and (i) show that the CTNPL estimator has advantages in terms of convergence³ and (ii) evaluate the potential for biased estimates and misleading economic conclusions when estimating a discrete time model using discrete time snapshot data that was actually generated by a continuous time model.

The rest of the paper is organized as follows. Section 2 develops a dynamic discrete game in continuous time with an example. Section 3 introduces the CTNPL algorithm to estimate the parameters in the model. In Section 4, we simulate a five-player game in retail industry and present the performance of the CTNPL estimator. Section 5 concludes.

2 Model

In this section, we briefly review the continuous time dynamic discrete game framework of ABBE (2016). To provide a concrete example, we also introduce a continuous-time version of the five-player entry/exit model of Aguirregabiria and Mira (2007). This model serves

³Despite its advantages, Pesendorfer and Schmidt-Dengler (2010) showed that the NPL algorithm may fail to converge to the NPL estimator when the NPL mapping is unstable. They provided a simple two-firm example model illustrating the issue and subsequent authors such as Egesdal et al. (2015) and Aguirregabiria and Marcoux (2021) have studied the problem in more complex models. Kasahara and Shimotsu (2012) provided a local convergence condition for the NPL estimator in discrete time models, however, Aguirregabiria and Marcoux (2021) show that even when the population mapping is stable the algorithm can fail to converge when the sample counterpart mapping is unstable. We derive a local convergence condition for continuous time models, which is motivated by that of Kasahara and Shimotsu (2012), and in our simulation results the condition is more likely to be satisfied in continuous time.

not only as a running example but also as the basis for our Monte Carlo experiments in Section 4.

2.1 Setting

There are N agents in the model indexed by $i = 1, \dots, N$ who choose actions j from a choice set $\mathcal{A} = \{0, 1, \dots, J - 1\}$ at decision times that occur in continuous time indexed by $t \in [0, \infty)$. Such decisions by agents result in endogenous state changes. The model also allows exogenous state variables. Decision times and exogenous state changes occur according to Poisson processes as detailed below.

State space. The states are discrete and finite, so every state at time t can be represented by a state vector x_k in a finite state space $\mathcal{X} = \{x_1, \dots, x_K\}$ where each state is indexed by $k = 1, 2, \dots, K$.

Example 1. *Suppose that there are $N = 5$ firms competing in a small local retail market where firms have at most one store. Based on incumbency status $a_i \in \{0, 1\}$ each firm i can either choose to enter/exit ($j = 1$) or do nothing ($j = 0$). Market size is denoted by an exogenous state variable $s_k \in \{1, 2, \dots, 5\}$. Therefore, there are $K = 5 \times 2^5 = 160$ distinct states in the model. Each state can be represented by a 6×1 vector x_k :*

$$x_k = (s_k, a_{1k}, a_{2k}, a_{3k}, a_{4k}, a_{5k}).$$

State changes. There are two types of state changes in the model—endogenous state changes which are relying on each agent’s decision and exogenous state changes driven by “nature.” As endogenous state changes only when an agent makes a choice, we can define a state continuation function $l(i, j, k)$ which maps an agent i ’s choice j and current state k to next state l . We assume that choice $j = 0$, where $l(i, 0, k) = k$, is costless and that all choices j are meaningfully distinct $l(i, j, k) \neq l(i, j', k)$ for all states $k = 1, \dots, K$. On the other hand, although exogenous state changes are not the result of an action by a player per se, we can think of them as actions by nature. When treated as a player in this way we will index nature by $i = 0$.

Example 1 (continued). *If firm 2 ($i = 2$) decides to enter a market ($j = 1$) while the market is in state k , then the continuation state index will be denoted $l(2, 1, k)$. The new state in vector form is $x_{l(2,1,k)} = (s_k, a_{1k}, 1, \dots, a_{5k})$. To give an example of an exogenous state change, when market size increases by 1 in state k , the state moves from k to l which can be written in vector form as $x' = (s_k + 1, a_{1k}, a_{2k}, \dots, a_{5k})$. Note that only one state*

variable moves at any single instant and other state variables remain the same as in the previous state. ■

Poisson and Markov jump processes.⁴ Corresponding to the two types of state changes in the model there are also two types of Poisson processes governing the times of those state changes. First, in each state k each player i makes decisions about changing its endogenous state according to a Poisson process with a rate parameter λ_{ik} . Recall that for such a Poisson process, the expected number of events during a one-unit time interval is $1/\lambda_{ik}$. For example, if the unit of time is one calendar year and we assume $\lambda_{ik} = \lambda = 1$ for all i and k , then each agent moves on average once per year. This would correspond closely to the usual discrete-time assumption of exactly one move per year. However, the Poisson assumption is more general since it allows, stochastically, for each player to make more or fewer moves per period. Because the probabilities of each choice vary, in general, across states, the state-to-state transition rates are endogenous and are state-and-firm specific. Second, in each state k the exogenous state variables change according to another state-specific Poisson process. We assume that these Poisson processes are independent so that the probability that more than one component of the state vector x_k changes at a single instant is zero. However, the processes are not identically distributed since they have different rate parameters.

Intensity matrices. Since Poisson processes are continuous-time Markov processes and the state space of the model is finite, the overall state transition process is a finite Markov jump process. Such a process can be characterized by a $K \times K$ intensity matrix (also known as a transition rate matrix or infinitesimal generator matrix) which contains the rates of all possible state-to-state transitions. This matrix is the instantaneous counterpart of the one-step transition probability matrix in a discrete time model.

Each component q_{kl} of the intensity matrix $Q = (q_{kl})$ represents the rate of leaving state k and transitioning to state l . The difference is that the intensity rate gives the transition rate for an instant instead of over one period of time. So, when we define transition probability matrix over some small time interval h as $P(h)$, we can write $P(h) = I + Qh$. Note that as $h \rightarrow 0$, $P(h)$ approaches the identity matrix (under which no transitions occur). Then, for states k, l with $k \neq l$, we have

$$\Pr(x_{t+h} = l \mid x_t = k) = \begin{cases} q_{kl}h & l \neq k, \\ 1 - \sum_{l \neq k} q_{kl}h & l = k. \end{cases}$$

⁴Definitions and notations in this section are based on Schuette and Metzner (2009).

We can formally define the intensity matrix Q :

$$Q = \begin{bmatrix} q_{11} & q_{12} & q_{13} & \cdots & q_{1K} \\ q_{21} & q_{22} & q_{23} & \cdots & q_{2K} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ q_{K1} & q_{K2} & q_{K3} & \cdots & q_{KK} \end{bmatrix},$$

where

$$q_{kl} = \begin{cases} \lim_{h \rightarrow 0} \frac{\Pr(x_{t+h}=l | x_t=k)}{h} & \text{for } l \neq k, \\ -\sum_{l \neq k} q_{kl} & \text{otherwise.} \end{cases}$$

For endogenous state changes, we can write q_{kl} as the product of the decision rate in state k and a choice probability. First, each agent decides whether or not to change the state according to Poisson process at rate λ_{ik} . To determine the choice probability, we must look closer at the player's problem. Players are forward-looking and discount future payoffs at rate $\rho \in (0, \infty)$. They observe other players' moves and optimally choose their own action $j \in \mathcal{A}$ based on their beliefs of other agents' actions, including state changes by nature. We denote the conditional choice probability of agent i choosing action j in state k as σ_{ijk} . So, if at some state k a player makes decisions at rate λ_{ik} and chooses action j by probability σ_{ijk} , then the rate of state changes due to player i taking action j in state k is $\lambda_{ik}\sigma_{ijk}$.

Under the assumptions of the model so far the overall intensity matrix Q can be written as a sum of individual player-specific intensity matrices: $Q = Q_0 + Q_1 + \cdots + Q_N$, where $Q_0 = (q_{kl})$ denotes the intensity matrix for exogenous state changes by nature and the elements of Q_0 are denoted q_{kl} (i.e., we drop the $i = 0$ subscript for nature). For players $i = 1, \dots, N$, each element of Q_i , denoted q_{ikl} , can be written as follows:

$$q_{ikl} = \begin{cases} \lambda_{ik}\sigma_{ijk} & \text{if } l \neq k \text{ and } l(i, j, k) = l, \\ -\sum_{l \neq k} q_{ikl} & l(i, j, k) = k, \\ 0 & \text{otherwise.} \end{cases}$$

The first case represents state transitions out of the current state. The rate of a transition from k to l is the product of the decision rate in k and the probability of choosing action j which results a transition to state l . The second case is the net rate of leaving state k , and therefore is the negative of the sum of the off-diagonal elements in row k . Because firms may find it optimal to remain in the same state (choosing $j = 0$) we must allow for self-transitions, so the total rate of leaving state k due to action by player i need not equal the rate of decisions λ_{ik} . The third case covers states l for which transitions from k to l

cannot occur in a single instant.

Example 1 (continued). *In the previous example, nature changes the market size according to Markov jump process characterized by intensity matrix Q_0 . We impose a natural restriction that the market size can only change by 1 level at an instant. Firms make decisions at rate λ and will then choose actions according to the choice probabilities σ_{ijk} . So, q_{ikl} for firm i 's intensity matrix can be written as:*

$$q_{ikl} = \begin{cases} \lambda\sigma_{i1k} & \text{if firm } i \text{ enters or exits and } l(i, 1, k) = l, \\ -\lambda\sigma_{i1k} & \text{if firm } i \text{ does nothing,} \\ 0 & \text{otherwise.} \end{cases}$$

■

Payoffs. In the continuous time setting we distinguish agents' flow payoffs from instantaneous choice-specific payoffs. First, agents receive flow payoffs at each state k regardless of their choices. We denote the flow payoffs by u_{ik} and assume $|u_{ik}| < \infty$ for $i = 1, \dots, N$ and $k = 1, \dots, K$. On the other hand, at a decision time when player i chooses action j in state k , the player receives a lump-sum instantaneous choice-specific payoff which is the sum of a deterministic component ψ_{ijk} , where $|\psi_{ijk}| < \infty$, and a stochastic component $\varepsilon_{ijk} \in \mathbb{R}$. The deterministic component ψ_{ijk} is common knowledge to other players but the second part ε_{ijk} is only observed by player i . However, players have common knowledge of the distribution of the stochastic components.

Example 1 (continued). *We consider a log-linear profit function for the flow payoffs including competition effects and market size.*

$$u_{ik} = \theta_{RS} \ln(s_k) - \theta_{RN} \ln\left(1 + \sum_{m \neq i} a_{mk}\right) - \theta_{FC,i} \quad (1)$$

We estimate the fixed costs $\theta_{FC,i}$ for $i = 1, \dots, 5$ and other profit function parameters θ_{RS} and θ_{RN} . When a firm decides an action j , they receive an instantaneous payoff $\psi_{ijk} + \varepsilon_{ijk}$ where

$$\psi_{ijk} = \begin{cases} -\theta_{EC} & \text{if } j = 1 \text{ and } a_{ik} = 0, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Note that there is a distinction between payoffs received while in a particular state and payoffs associated with actions. For example, the entry cost is paid as a lump sum in (2) instead of appearing in the flow profit equation (1). By comparison, the entry cost parameter appears

together with product market profits in the aggregated discrete time period profit equation given in (48) in Aguirregabiria and Mira (2007). ■

2.2 Bellman Optimality and CCP Representation

In a Markovian game, every agent plays a stationary Markov strategy based on their beliefs regarding other agents' choice probabilities. Let ς_i denote the beliefs held by player i about each of the $(N-1)$ other players playing each of their J choices in each of K states. Therefore, ς_i is a collection of $(N-1) \times J \times K$ probabilities, denoted ς_{imjk} for each player $m \neq i$, choice j , and state k .

We now introduce the value functions and CCPs. Agents establish the value function based on their expectations about endogenous state changes due to other agents' moves, exogenous state changes due to nature, and their own future decisions. We follow the derivation of instantaneous Bellman equation in Blevins (2016). The probability that agent i makes a decision during a small increment of time of length h while in state k is $\lambda_{ik}h$ under the Poisson assumption. Denoting each agent i 's discount rate as ρ_i , the discount factor for the time increment h is $1/(1 + \rho_i h)$. Given agent i 's beliefs, ς_i , we can express the Bellman equation as below:

$$V_{ik} = \frac{1}{1 + \rho_i h} \left[u_{ik}h + \sum_{l \neq k} q_{kl}h V_{il} + \sum_{m \neq i} \lambda_{mk}h \sum_{j=0}^J \varsigma_{imjk} V_{i,l(m,j,k)} \right. \\ \left. + \lambda_{ik}h E \left(\max_j \{ \psi_{ijk} + \varepsilon_{ijk} + V_{l(i,j,k)} \} \right) + \left(1 - \sum_{i=1}^N \lambda_{ik}h - \sum_{l \neq k} q_{0kl}h \right) V_{ik} + o(h) \right].$$

We now turn to each of the terms inside square brackets in the previous equation. First, over the small amount of time h that the model remains in state k player i receives an integrated flow payoff of $u_{ik}h$. Second, if the state changes from k to l due to nature, which occurs with probability proportional q_{kl} over the interval h , then player i receives the continuation value V_{il} . Third, each of player i 's rival players, indexed by $m \neq i$, may make a decision with probability proportional to λ_{mk} over the time interval h . When this occurs, player i believes that player m plays action j in state k with probability ς_{imjk} . When this happens, player i receives a continuation value $V_{i,l(m,j,k)}$ where $l(m,j,k)$ denotes the index of the continuation state following the action of player m . Fourth, at rate λ_{ik} player i makes a decision and optimally chooses an action j . However, the values of the stochastic choice-specific shocks have not yet been realized, so player i receives the expected maximum value of $\psi_{ijk} + \varepsilon_{ijk} + V_{l(i,j,k)}$ where j is chosen optimally at the future decision time. Finally, the

last (non-negligible) term accounts for the case where no state changes occur over the time interval h and so the continuation value is simply V_{ik} .

By rearranging and letting $h \rightarrow 0$, we arrive at a simpler expression for the Bellman equation:

$$V_{ik} = \frac{u_{ik} + \sum_{l \neq k} q_{kl} V_{il} + \sum_{m \neq i} \lambda_{mk} \sum_j \varsigma_{imjk} V_{i,l(m,j,k)} + \lambda_{ik} E \max_j \{ \psi_{ijk} + \epsilon_{ijk} + V_{i,l(i,j,k)} \}}{\rho + \sum_{l \neq k} q_{kl} + \sum_{i=1}^N \lambda_{ik}}.$$

Now we define the best response function for each agent. A Markov strategy for agent i is a mapping σ_i from each state $(k, \epsilon_{ijk}) \in \mathcal{X} \times \mathbb{R}^J$ to an action $j \in \mathcal{A}$. Then, the best response for an agent i is defined as:

$$\sigma_i(k, \epsilon; \varsigma_i) = j \iff \psi_{ijk} + \epsilon_{ijk} + V_{i,l(i,j,k)} \geq \psi_{ij'k} + \epsilon_{ij'k} + V_{i,l(i,j',k)} \quad \forall j' \in \mathcal{A}.$$

Then, the choice probability for each agent becomes:

$$\sigma_{ijk} = \Pr[\delta_i(k, \epsilon_i; \varsigma_i) = j | k].$$

We now collect assumptions needed for the existence of an equilibrium.

Assumption 1 (Discrete states). *The state space is finite: $K = |\mathcal{X}| < \infty$.*

Assumption 2 (Bounded rates and payoffs). *The discount rate ρ_i , move arrival rate, rates of state changes due to nature, and payoffs are all bounded for all $i = 1, \dots, N$, $j = 0, \dots, J-1$, $k = 1, \dots, K$, $l = 1, \dots, K$ with $l \neq k$: (a) $0 < \rho_i < \infty$, (b) $0 < \lambda_{ik} < \infty$, (c) $0 \leq q_{kl} < \infty$, (d) $|u_{ik}| < \infty$, and (e) $|\psi_{ijk}| < \infty$.*

Assumptions 1 and 2 are needed to define a continuous-time finite Markov jump process. In Assumption 2, we rule out infinite rates and payoffs to make it possible to estimate by restricting the value to a finite number.

Assumption 3 (Additive separability). *In each state k , the instantaneous payoff associated with choice j is additively separable as $\psi_{ijk} + \epsilon_{ijk}$.*

Assumption 4 (Distinct actions). *For all $i = 1, \dots, N$ and $k = 1, \dots, K$, the continuation state function $l(i, j, k)$ and the choice-specific payoffs ψ_{ijk} satisfy the following two properties: (a) choice $j = 0$ is a costless continuation choice with $l(i, j, k) = k$ and $\psi_{ijk} = 0$, and (b) all choices j are meaningfully distinct in the sense that the continuation states differ: $l(i, j, k) \neq l(i, j', k)$ for all $j = 0, \dots, J-1$ and $j' \neq j$.*

Assumption 5 (Private information). *The errors ε_{ik} are i.i.d. over time with joint distribution F which is absolutely continuous with respect to Lebesgue measure (with joint density f), has finite first moments, and has support equal to \mathbb{R}^J .*

Assumption 3 is common in the literature, as in the case of Assumption 1 of Aguirregabiria and Mira (2002) and Aguirregabiria and Mira (2007). Assumption 4 requires that different choices should result in observationally distinct state changes. For example, choices $j = 1$ and $j = 2$ should not be defined separately if they in fact always result in the same state changes. It also involves a normalization on the baseline payoff/cost associated with inaction. Finally, the distributional assumption on ε is widely in the previous literature. For example, see Assumption 2 of Aguirregabiria and Mira (2002) and Aguirregabiria and Mira (2007).

We define Markov perfect equilibrium (MPE) as in Aguirregabiria and Mira (2007) and ABBE (2016).

Definition 1 (Markov perfect equilibrium). *A collection of Markov strategies $\{\delta_1, \dots, \delta_N\}$ is a Markov perfect equilibrium if for all i and for any $(k, \varepsilon_{ik}) \in \mathcal{X} \times \mathbb{R}^J$, $\delta_i(k, \varepsilon_{ik})$ is a best response.*

By Proposition 5 in ABBE (2016), under Assumptions 1–5, a Markov perfect equilibrium exists. Moreover, under Assumptions 1–5 and assuming that λ_{ik} is constant across all agents and states denoted by λ , ABBE (2016) established the following linear representation of the value function, where $V_i(\theta, \sigma)$ is a K -vector:

$$V_i(\theta, \sigma) = \left[(\rho + N\lambda)I - \lambda \sum_{m=1}^N \Sigma_m(\sigma_m) - Q_0 \right]^{-1} [u_i(\theta) + \lambda_i E_i(\theta, \sigma)], \quad (3)$$

where $\Sigma_m(\sigma_m)$ is the $K \times K$ state transition matrix induced by the actions of player m given the choice probabilities σ_m and where $E_i(\theta, \sigma)$ is a $K \times 1$ vector where each element k is the ex-ante expected value of the choice-specific payoff in state k , $\sum_j \sigma_{ijk} [\psi_{ijk} + e_{ijk}(\theta, \sigma)]$ where $e_{ijk}(\theta, \sigma)$ is the expected value of ε_{ijk} given that choice j is optimal,

$$\frac{1}{\sigma_{ijk}} \int \varepsilon_{ijk} \cdot \mathbf{1}\{\varepsilon_{ij'k} - \varepsilon_{ijk} \leq \psi_{ijk} - \psi_{ij'k} + V_{l(i,j,k)}(\theta, \sigma) - V_{l(i,j',k)}(\theta, \sigma) \forall j'\} f(\varepsilon_{ik}) d\varepsilon_{ik}.$$

We now have two important equations that forms a players' decision problem.

1. Bellman optimality: We stack (3) across players $i = 1, \dots, N$ and define the policy valuation operator $\Upsilon(\theta, \sigma)$:

$$\Upsilon(\theta, \sigma) \equiv V(\theta, \sigma) = \left[(\rho + N\lambda)I - \lambda \sum_{m=1}^N \Sigma_m(\theta, \sigma_m) - Q_0 \right]^{-1} [u + \lambda E(\theta, \sigma)]. \quad (4)$$

2. Conditional choice probability

$$\Gamma(v) \equiv \sigma$$

where σ is a $NJK \times 1$ vector with $\sigma_{ijk} = \Pr[\delta_i(k, \varepsilon) = j|k]$. Assuming ε follows i.i.d. $TIEV(0, 1)$,

$$\sigma_{ijk} = \frac{\exp(\psi_{ijk} + V_{l(i,j,k)})}{\sum_{j'} \exp(\psi_{ij'k} + V_{l(i,j',k)})} \quad (5)$$

We now introduce a policy iteration operator Ψ in the space of conditional choice probabilities. Substituting (1) into (2),

$$\sigma = \Psi(\theta, \sigma) \equiv \Gamma(\Upsilon(\theta, \sigma)) \quad (6)$$

where Υ is a policy valuation operator that maps an $NJK \times 1$ vector of conditional choice probabilities into an $NK \times 1$ vector in value function space. Γ is a policy improvement operator that maps an $NK \times 1$ vector in value function space into $NJK \times 1$ vector of conditional choice probabilities.

3 The CTNPL Estimator

3.1 CTNPL Algorithm

Before describing the CTNPL algorithm, we first introduce two more assumptions for identification and estimation with continuous time data observed only at fixed time intervals (Assumptions 6, 7 in ABBE, 2016; Assumptions 8, 9 in Blevins, 2016).

Assumption 6. *The mapping $Q \rightarrow \{Q_0, Q_1, \dots, Q_N\}$ is known.*

Assumption 7. *(i) In every market $m = 1, \dots, M$, every agent expects a single equilibrium to be played which results an intensity matrix Q . (ii) The distribution of state transition at any time $t \in [0, \infty)$ is consistent with the intensity matrix $P(\Delta) = \exp(\Delta Q)$.*

For continuous time models, we consider two types of data. The first type is continuous time data which is observed every instant. Stock prices can be an example close to this type. Second, some data generated by a continuous time model can be only observed at discrete times. For example, population can change any time during the year but annual population data may only be observed once per year. We show how to construct log likelihood functions for estimation using both types of data. We assume throughout that we have independent and identically distributed data across markets and that the Markov process is irreducible.

Fully observed continuous time data. Consider a dataset of observations in market m and n -th move, $\{k_{mn}, t_{mn} : m = 1, \dots, M, n = 1, \dots, T_m\}$ sampled in interval $[0, \bar{T}]$. State k_{mn} denotes the state immediately before state change at time t_{mn} and time t_{mn} is the time of n -th state change in market m .

We follow the derivation in Section 6.1. in ABBE (2016) to construct the likelihood function. Let h denote a $K(K-1) + NJK$ vector of hazard rates for state change:

$$h = (q_{12}, q_{13}, \dots, q_{K-1,K}, \lambda_{11}\sigma_{111}, \dots, \lambda_{11}\sigma_{1JK}, \dots, \lambda_{N1}\sigma_{N11}, \dots, \lambda_{NK}\sigma_{NJK}).$$

The probability that there is a state change within τ units of time is:

$$\left(\sum_{l \neq k} q_{kl} + \sum_i \lambda_{ik} \sum_{j \neq 0} \sigma_{ijk} \right) \exp \left[-\tau \left(\sum_{l \neq k} q_{kl} + \sum_i \lambda_{ik} \sum_{j \neq 0} \sigma_{ijk} \right) \right].$$

This is because state changes due to nature occur in state k according to Poisson processes with rate parameters q_{kl} , for each state $l \neq k$, and each player i takes some action according to a Poisson process with rate parameter λ_{ik} .

When there is a state change from state k , conditional probability that the change is due to agent i 's action j is

$$\frac{\lambda_{ik}\sigma_{ijk}}{\sum_{l \neq k} q_{kl} + \sum_i \lambda_{ik} \sum_{j \neq 0} \sigma_{ijk}}$$

and conditional probability that it is due to nature is

$$\frac{q_{kl}}{\sum_{l \neq k} q_{kl} + \sum_i \lambda_{ik} \sum_{j \neq 0} \sigma_{ijk}}.$$

Defining $g(\tau, k; h) = \exp \left[-\tau (\sum_{l \neq k} q_{kl} + \sum_i \lambda_{ik} \sum_{j \neq 0} \sigma_{ijk}) \right]$, the likelihood of state change is

$$\begin{cases} \lambda_{ik}\sigma_{ijk}g(\tau, k; h) & \text{if agent } i \text{ chooses action } j, \\ q_{kl}g(\tau, k; h) & \text{if nature moves.} \end{cases} \quad (7)$$

Let $\tau_{m,T+1}$ be the time interval between the last state change at time T and \bar{T} , and $k_{m,T+1} \equiv k_{m,T}$. We introduce an indicator function $I_{mn}(i, j)$ which is 1 when n -th state change at market m is induced by agent i 's action j and 0 otherwise. We also define $I_{mn}(0, l)$ as 1 when n -th state change at market m is by nature's move to state l and 0 otherwise.

Combining the results above, the likelihood function becomes

$$L(\theta, h) = \frac{1}{M} \sum_{m=1}^M \left[\sum_{n=1}^T \left\{ \ln g(\tau, k; h) + \sum_{l \neq k_{mn}} I_{mn}(0, l) \ln q_{k_{mn}, l} + \sum_i \sum_{j \neq 0} I_{mn}(i, j) \ln(\lambda_{imk} \sigma_{ijk}) \right\} + \ln g(\tau_{m, T+1}, k_{m, T+1}; h) \right]. \quad (8)$$

The first term inside the summation is the common term in (7) and the second term is from the likelihood of state changes due to nature. The third term is considering agents' moves and the last term is the likelihood of observing no state change after time $t_{m, T}$.

Since state changes due to nature are exogenous, the log-likelihood function of the model can be decomposed into terms involving the conditional choice probabilities of players and nature's transition probabilities. Consistent estimates for elements of Q_0 , $q = (q_{12}, \dots, q_{K-1, K})$, can be obtained from transition data without having to solve the dynamic game. Therefore, we assume in the remainder of the paper that q is known and focus on the estimation of θ and $\sigma = (\sigma_{111}, \dots, \sigma_{NJK})$. Then, the likelihood function becomes

$$L_M(\theta, \sigma) = \frac{1}{M} \sum_{m=1}^M \left[\sum_{n=1}^T \left\{ \ln g(\tau, k; \sigma) + \sum_i \sum_{j \neq 0} I_{mn}(i, j) \ln(\lambda_{imk} \sigma_{ijk}) \right\} + \ln g(\tau_{m, T+1}, k_{m, T+1}; \sigma) \right].$$

Discretely observed data. Consider a dataset $\{k_{mn} : m = 1, \dots, M, n = 0, \dots, T\}$ consisting of observations that are sampled at times on the lattice $\{n\Delta : n = 0, \dots, T\}$. As before, we assume that q is known. Given q and σ , we can construct a transition matrix. Define the pseudo likelihood function

$$L_M(\theta, \sigma) = \frac{1}{M} \sum_{m=1}^M \sum_{n=1}^T \ln P_{k_{m, n-1}, k_{mn}}(\Delta; \Psi(\theta, \sigma)) \quad (9)$$

where $P_{k, l}(\Delta; \Psi(\theta, \sigma))$ denotes the (k, l) element of the transition matrix induced by $\sigma = \Psi(\theta, \sigma)$.

We calculate $P(\Delta)$ using matrix exponential as in ABBE (2016). First, assume that we have a continuous time data observed only at discrete times with unit interval, $\Delta = 1$. Let $Z(\theta, \sigma)^r$ denote the transition matrix from some state k to another state k' after an unit interval of time $\Delta = 1$ in r steps. Then, let $a_{m, n}$ denote an indicator vector for an observation n in market m which is one in position $k_{m, n}$ and zero elsewhere. We can write the likelihood

function as below:

$$L_M(\theta, \sigma) = \frac{1}{M} \sum_{m=1}^M \sum_{n=1}^T \ln \left[\sum_{r=0}^{\infty} \frac{\lambda^r \exp(-\lambda)}{r!} a'_{m,n} Z(\theta, \sigma)^r a_{m,n+1} \right]. \quad (10)$$

Then, using the definition of matrix exponential, we can rewrite the previous equation as:

$$L_M(\theta, \sigma) = \frac{1}{M} \sum_{m=1}^M \sum_{n=1}^T \ln [a'_{m,n} \exp(Q(\theta, \sigma)) a_{m,n+1}], \quad (11)$$

where Q is the aggregate intensity matrix for the model.

CTNPL algorithm. We adapt the NPL algorithm to obtain the CTNPL estimator in the present continuous time model. Let $\hat{\sigma}^0$ be an initial estimate or guess of the vector of players' choice probabilities. Given $\hat{\sigma}^0$, for $l \geq 1$, perform the following steps:

1. Given $\hat{\sigma}^{l-1}$, update $\hat{\theta}$ by

$$\hat{\theta}^l = \underset{\theta \in \Theta}{\operatorname{argmax}} L_M(\theta, \hat{\sigma}^{l-1}) \quad (12)$$

2. Update $\hat{\sigma}^l$ using the equilibrium condition, i.e.

$$\hat{\sigma}^l = \Psi(\hat{\theta}^l, \hat{\sigma}^{l-1}) \quad (13)$$

Iterate in l until convergence in σ and θ is reached. Since we usually have observe continuous time data at discrete times in practice, we focus on properties of the CTNPL estimator using the likelihood function in (9) in next section.

Remark on the definition of the NPL estimator. Aguirregabiria and Mira (2002) define $\hat{\theta}^l$ as policy iteration (PI) estimator and the estimation procedure as the NPL algorithm. In Aguirregabiria and Mira (2007), they define the same estimator as K -stage estimator⁵ and define the NPL estimator as the limit of the sequence $\{\hat{\theta}^l, \hat{\sigma}^l\}_{l=1}^L$ with certain conditions. Specifically, if the sequence $\{\hat{\theta}^l, \hat{\sigma}^l\}_{l=1}^L$ converges, its limit $(\tilde{\theta}, \tilde{\sigma})$, which is called NPL fixed point, satisfies two conditions:

$$\tilde{\theta} = \underset{\theta \in \Theta}{\operatorname{argmax}} L(\theta, \tilde{\sigma}) \quad \text{and} \quad \tilde{\sigma} = \Psi(\tilde{\theta}, \tilde{\sigma}).$$

It is possible that there exist multiple NPL fixed points, so they choose the one with highest maximum likelihood and define it as the NPL estimator, $(\hat{\theta}_{\text{NPL}}, \hat{\sigma}_{\text{NPL}})$. Subsequent studies

⁵They call it K -stage estimator as they use K as a notation for the number of stages. They also define $(\hat{\theta}^1, \hat{\sigma}^1)$ as the Pseudo Maximum Likelihood (PML) estimator.

use the definition of NPL estimator in Aguirregabiria and Mira (2007), so we follow the same definition (Pesendorfer and Schmidt-Dengler, 2010; Kasahara and Shimotsu, 2012). Henceforth, we define $(\hat{\theta}^l, \hat{\sigma}^l)$ as the PI estimator and $(\hat{\theta}, \hat{\sigma}) = (\hat{\theta}_{\text{CTNPL}}, \hat{\sigma}_{\text{CTNPL}})$ as the CTNPL estimator, dropping the subscript for simplicity.

3.2 Large Sample Properties

In this section, we consider the large sample properties of the CTNPL estimator. We start with the CTNPL estimator in the case of single agent continuous time models, then we consider both the PI estimator (for a finite number of iterations l) and the CTNPL estimator continuous time dynamic games.

3.2.1 Single Agent Dynamic Discrete Choice Models

The asymptotic properties of the CTNPL estimator for single agent models depend on a zero Jacobian property for $\Psi(\theta, \sigma)$. Specifically, we show that the Jacobian matrix $\Psi_\theta \equiv \nabla_\theta \Psi(\theta^*, \sigma^*)$ is equal to the zero matrix in agent models. This result is presented formally in Proposition 1 immediately below, and we then show that it facilitates adaptive two-step estimation in Proposition 2, where we show the estimator is efficient. The same zero Jacobian property will be used later, in Section 3.3, where we show that it also ensures local stability of the estimator. All proofs are provided in Appendix A.

Proposition 1. *In a single agent model, the Jacobian matrix Ψ_σ is zero at the fixed point (θ^*, σ^*) .*

Next, Proposition 2 shows that, under certain regularity conditions, the NPL estimator in a continuous time single agent discrete choice model is consistent, asymptotically normal, and efficient.

Proposition 2. *Let Θ denote the parameter space, let σ^* denote the true CCPs, and define $\Xi \equiv \underbrace{\Delta^J \times \dots \times \Delta^J}_K$ to be the set of possible values of σ , where Δ^J denotes the standard J -simplex. Consider the following regularity conditions:*

- (a) Θ is a compact set.
- (b) $\Psi(\theta, \sigma)$ is continuous and twice continuously differentiable in θ and σ .
- (c) Let $\Psi_{jk}(\theta, \sigma)$ be the corresponding element of $\Psi(\theta, \sigma)$ for some choice j and state k . Then, $\Psi_{jk}(\theta, \sigma) > 0$ for any $j = 0, 1, \dots, J - 1$, $k = 1, \dots, K$, and any $(\theta, \sigma) \in \Theta \times \Xi$.

(d) There is a unique $\theta^* \in \text{int}(\Theta)$ such that $\sigma^* = \Psi(\theta^*, \sigma^*)$.

(e) $\hat{\sigma}^0$ is a strongly consistent estimator of σ^* such that

$$[\sqrt{M}\nabla_{\theta}L_M(\theta^*, \sigma^*); \sqrt{M}(\hat{\sigma}^0 - \sigma^*)'] \xrightarrow{d} \text{N}(0, \Omega)$$

Then,

$$\sqrt{M}(\hat{\theta}^l - \theta^*) \xrightarrow{d} \text{N}(0, \Omega_{\theta\theta'}^{-1}) \quad (14)$$

where $\Omega_{\theta\theta'} = \text{E}[\nabla_{\theta}s_m \nabla_{\theta'}s_m]$ and s_m is defined as

$$s_m \equiv \sum_{n=1}^T \ln P_{k_m, n-1 k_{mn}}(\Delta; \Psi(\theta^*, \sigma^*)).$$

Conditions (a), (b), and (c) are usual regularity conditions, while (d) is needed for identification. Condition (e) requires that the initial CCP estimates are consistent. As can be seen from the limiting distribution in (14), the asymptotic variance is not affected by the first-step estimation of the CCPs. In other words, in a single agent model there is no efficiency loss from two-step estimation, as was established in discrete time models by Aguirregabiria and Mira (2002). This is due to the so-called zero Jacobian property, which we show also holds in continuous-time models below in Proposition 1, where we also relate the property to convergence of the algorithm.

3.2.2 Dynamic Discrete Games

We now extend the model to a dynamic discrete game with N players and present the large sample theory first for the PI estimator for $l = 1, \dots, L$ in Proposition 3 and then the CTNPL estimator in Proposition 4.

Proposition 3. Let σ^* be the true set of conditional choice probabilities and let Θ and $\Xi \equiv \underbrace{\Delta^J \times \dots \times \Delta^J}_{NK}$ be the set of possible values of θ and σ , where Δ^J denotes the standard J -simplex. Consider the following regularity conditions:

(a) Θ is a compact set.

(b) $\Psi(\theta, \sigma)$ is continuous and twice continuously differentiable in θ and σ .

(c) $\Psi_{ijk}(\theta, \sigma)$ be the corresponding element of $\Psi(\theta, \sigma)$ for player i 's choice j at state k .

Then, $\Psi_{ijk}(\theta, \sigma) > 0$ for any $i = 1, \dots, N$, $j = 0, 1, \dots, J - 1$, $k = 1, \dots, K$, and any $\{\theta, \sigma\} \in \Theta \times \Xi$.

(d) There is a unique $\theta^* \in \text{int}(\Theta)$ such that $\sigma^* = \Psi(\theta^*, \sigma^*)$.

(e) $\hat{\sigma}^0 = M^{-1} \sum_{m=1}^M r_m$ is a strongly consistent nonparametric estimator of σ^* such that $\sqrt{M}(\hat{\sigma}^0 - \sigma^*) \xrightarrow{d} N(0, \Sigma^0)$.

Then, for all $l \leq L$,

$$\sqrt{M}(\hat{\theta}^l - \theta^*) \xrightarrow{d} N\left(0, \Omega_{\theta\theta'}^{-1} + \Omega_{\theta\theta'}^{-1} \Omega_{\theta\sigma'} \Sigma^{l-1} \Omega'_{\theta\sigma'} \Omega_{\theta\theta'}^{-1}\right)$$

where $s_m = \sum_{n=1}^T P_{k_{m,n-1}, k_{m,n}}(\Delta; \Psi(\theta, \sigma))$, $\Omega_{\theta\theta'} = E[\nabla_{\theta} s_m \nabla_{\theta'} s_m]$, $\Omega_{\theta\sigma'} = E[\nabla_{\theta} s_m \nabla_{\sigma'} s_m]$, and $\Sigma^l = \nabla_{\theta} \Psi(\Omega_{\theta\theta'}^{-1} + \Omega_{\theta\theta'}^{-1} \Omega_{\theta\sigma'} \Sigma^{l-1} \Omega'_{\theta\sigma'} \Omega_{\theta\theta'}^{-1}) \nabla_{\theta} \Psi' + \nabla_{\sigma} \Psi \Sigma^{l-1} \nabla_{\sigma} \Psi'$ is the asymptotic variance of $\sqrt{M}(\hat{\sigma}^l - \sigma^*)$ with $\Psi_{\theta} \equiv \nabla_{\theta} \Psi(\theta^*, \sigma^*)$ and $\Psi_{\sigma} \equiv \nabla_{\sigma} \Psi(\theta^*, \sigma^*)$ denoting Jacobian matrices.

The results above are for the PI estimator $\hat{\theta}^l$ for $l = 1, \dots, L$. Aguirregabiria and Mira (2007) provide large sample properties for the NPL estimator (i.e., at convergence, if the algorithm converges). We now present similar properties for the CTNPL estimator $\hat{\theta}$ in our continuous time model. A local convergence condition for $\hat{\theta}^l$ will be provided in the next section.

Proposition 4 below establishes the large sample properties for $(\hat{\theta}, \hat{\sigma})$, but we first introduce a few definitions and notation from Aguirregabiria and Mira (2007). Let $\phi_M(\cdot)$ denote the NPL operator used to represent an iteration of the algorithm:

$$\phi_M(\sigma) = \Psi(\tilde{\theta}_M(\sigma), \sigma) \quad \text{where} \quad \tilde{\theta}_M(\sigma) \equiv \underset{\theta \in \Theta}{\operatorname{argmax}} L_M(\theta, \sigma).$$

Let Υ_M be the set of fixed points of ϕ , $\Upsilon_M \equiv \{(\theta, \sigma) \in \Theta \times [0, 1]^{NJK} : \theta = \tilde{\theta}(\sigma) \text{ and } \sigma = \phi_M(\sigma)\}$. Then, choosing the fixed point that has the maximum value of the likelihood,

$$(\hat{\theta}, \hat{\sigma}) \equiv \underset{(\theta, \sigma) \in \Upsilon_M}{\operatorname{argmax}} L_M(\theta, \sigma).$$

We denote the population counterparts of $\tilde{\theta}_M(\sigma)$ and $\tilde{\phi}_M(\sigma)$ as $\tilde{\theta}^*(\sigma) \equiv \underset{\theta \in \Theta}{\operatorname{argmax}} L(\theta, \sigma)$ and $\phi^*(\sigma) = \Psi(\tilde{\theta}^*(\sigma), \sigma)$.

Proposition 4. *Let σ^* be the true set of conditional choice probabilities and let Θ and $\Xi \equiv [0, 1]^{NJK}$ be the set of possible values of θ and σ , respectively. Consider the following regularity conditions:*

- (a) Θ is a compact set.
- (b) $\Psi(\theta, \sigma)$ is continuous and twice continuously differentiable in θ and σ .
- (c) $\Psi_{ijk}(\theta, \sigma)$ be the corresponding element of $\Psi(\theta, \sigma)$ for player i 's choice j at state k . Then, $\Psi_{ijk}(\theta, \sigma) > 0$ for any $i = 1, \dots, N$, $j = 1, \dots, J - 1$, $k = 1, \dots, K$, and any $\{\theta, \sigma\} \in \Theta \times \Xi$.

(d) There is a unique $\theta^* \in \text{int}(\Theta)$ such that $\sigma^* = \Psi(\theta^*, \sigma^*)$.

(e) (θ^*, σ^*) is an isolated population NPL fixed point.

(f) There exists a closed neighborhood of σ^* , $\mathcal{N}(\sigma)$, such that, for all $\sigma \in \mathcal{N}(\sigma^*)$, $L(\theta, \sigma)$ is globally concave in θ and $\partial^2 L(\theta, \sigma^*) / \partial \theta \partial \theta'$ is a nonsingular matrix.

(g) The operator $\phi^*(\sigma) - \sigma$ has a nonsingular Jacobian matrix at σ^* .

Then,

$$\sqrt{M}(\hat{\theta} - \theta^*) \xrightarrow{d} \text{N} \left(0, [\Omega_{\theta\theta'} + \Omega_{\theta\sigma'}(I - \Psi_\sigma)^{-1}\Psi_\theta]^{-1}\Omega_{\theta\theta'}[\Omega_{\theta\theta'} + \Psi'_\theta(I - \Psi'_\sigma)^{-1}\Omega'_{\theta\sigma'}]^{-1} \right)$$

where $\Psi_\theta \equiv \nabla_\theta \Psi(\theta^*, \sigma^*)$ and $\Psi_\sigma \equiv \nabla_\sigma \Psi(\theta^*, \sigma^*)$ denote Jacobian matrices.

Assumptions from (a)–(d) are the same as Assumptions (a)–(d) in Proposition 3. Assumptions (e)–(g) correspond to Assumptions (v)–(vii) of Proposition 2 in Aguirregabiria and Mira (2007). Note that consistency of $\hat{\sigma}^0$ is not required.

3.3 Convergence of the CTNPL Algorithm

Originally introduced for discrete time models, the NPL algorithm is a sequential method to find the NPL estimator. This is why it is possible for the NPL algorithm to fail to converge to a consistent estimator even if the NPL estimator is consistent. Ideally, the NPL algorithm could converge to the consistent estimator if it can be evaluated at every fixed point, but this may be infeasible in practice.

Starting from Pesendorfer and Schmidt-Dengler (2010), several papers have investigated the convergence of the NPL algorithm. Kasahara and Shimotsu (2012) show that a key determinant of the convergence of the NPL algorithm is the contraction property of the mapping Ψ and propose a local convergence condition for NPL estimator in discrete time. Aguirregabiria and Marcoux (2021) investigate this further and show that even when the population mapping is stable the algorithm can fail to converge when the sample NPL mapping is unstable. We extend these arguments to the CTNPL estimator and derive a local convergence condition for continuous time models that is analogous to the condition of Kasahara and Shimotsu (2012).

For simplicity, we define $\Psi_\theta \equiv \nabla_\theta \Psi(\theta^*, \sigma^*)$ and $\Psi_\sigma \equiv \nabla_\sigma \Psi(\theta^*, \sigma^*)$. Let P be an aggregate transition probability matrix between states where the elements P_{kl} of P are

$$P_{kl} = \begin{cases} \lambda_{ik}\sigma_{ijk} & \text{where } l(i, j, k) = l \quad \text{if state change from } k \text{ to } l \text{ is endogenous.} \\ [\exp(Q_0)]_{kl} & \text{if state change from } k \text{ to } l \text{ is exogenous.} \end{cases}$$

Let $P^* = \vec{P}(\Delta; \Psi(\theta^*, \sigma^*))$ denote the $K^2 \times 1$ vectorized version of $P(\Delta; \Psi(\theta^*, \sigma^*))$. Then,

define a scalar $a_{ij,kl}$ as:

$$a_{ij,kl} = \begin{cases} 1 & \text{if } l(i, j, k) = l \\ 0 & \text{otherwise.} \end{cases}$$

We then construct a $NJK \times K^2$ matrix, denoted A , by stacking $a_{ij,kl}$ accordingly, and define $\Delta_\sigma = A \text{diag}(P^*)^{-1} A'$. Then, we can rewrite $\Omega_{\theta\theta'} = \Psi'_\theta \Delta_\sigma \Psi_\theta$ and $\Omega_{\theta\sigma'} = \Psi'_\theta \Delta_\sigma \Psi_{\sigma'}$. We also define $M_{\Psi_\theta} \equiv I - \Psi_\theta (\Psi_{\theta'} \Delta_\sigma \Psi_\theta)^{-1} \Psi_{\theta'} \Delta_\sigma$.

Assumption 8. (a) The conditions of Proposition 4 hold. (b) $\Psi(\theta, \sigma)$ is three times continuously differentiable in a closed neighborhood \mathcal{N} of (θ^*, σ^*) . (c) $\Omega_{\theta\theta'}$ is nonsingular.

We now present a local convergence condition for the CTNPL estimator $\hat{\sigma}^l$. If $\hat{\sigma}^l$ converges to $\hat{\sigma}$, $\hat{\theta}^l$ converges to $\hat{\theta}$ by the continuity of $L(\theta, \sigma)$ and applying the Theorem of the Maximum in (12). So, it is sufficient to show the convergence for $\hat{\sigma}^l$. The convergence condition below is analogous to Proposition 2 in Kasahara and Shimotsu (2012).

Proposition 5. Suppose that Assumption 8 holds and that $\rho(M_{\Psi_\theta} \Psi_\sigma) < 1$. Then, there exists a neighborhood \mathcal{N}_1 of σ^* such that, for any initial value $\hat{\sigma}^0 \in \mathcal{N}_1$, we have $\lim_{L \rightarrow \infty} \hat{\sigma}^L = \hat{\sigma}$ almost surely.

In single agent models, the Jacobian matrix Ψ_σ is zero at the fixed point $\sigma^* = \Psi(\theta^*, \sigma^*)$ by Proposition 1, it follows that $\rho(M_{\Psi} \Psi_\sigma)$ is zero. Then, by Proposition 5, the CTNPL estimator is stable and always converges locally to $(\hat{\theta}, \hat{\sigma})$ in single agent models. Next, we will investigate many of the properties we have derived in the context of a simulation study.

4 Monte Carlo Experiments

In this section we present the results of a series of Monte Carlo experiments carried out using the model of Example 1 from Section 2, which is a continuous time version of the dynamic discrete game described in Section 4 of Aguirregabiria and Mira (2007). We also provide further implementation details of the CTNPL algorithm for this model in Appendix B.

4.1 Data Generating Process (DGP)

We assume the profit functions of firms are:

$$u_{ik} = \theta_{RS} \ln(s_k) - \theta_{RN} \ln \left(1 + \sum_{m \neq i} a_{mk} \right) - \theta_{FC,i}.$$

The parameters we estimate are the fixed costs $\theta_{FC,i}$, entry costs θ_{EC} , competition effect θ_{RN} , and market effect θ_{RS} . Fixed costs $\theta_{FC,i}$ are set as $(\theta_{FC,1}, \theta_{FC,2}, \theta_{FC,3}, \theta_{FC,4}, \theta_{FC,5}) = (-1.9, -1.8, -1.7, -1.6, -1.5)$. The parameter θ_{RS} equals to 1. We run six experiments varying by entry costs $\theta_{EC} = (1.0, 2.0, 4.0)$ and $\theta_{RN} = (0.0, 1.0, 2.0)$. The stochastic component in instantaneous payoff ϵ_{ijk} are independently and identically distributed and follow Type 1 extreme value distribution.

The support of the logarithm of market size is $\{1, 2, 3, 4, 5\}$ and each player chooses an action from $a_{ik} = \{0, 1\}$. So, the size of state space is $K = 5 \times 2^5 = 160$. We assume that market size increases or decreases by at most 1 step at a single instant and that the rate is constant across states k . Let q_1 and q_2 , respectively, denote the rate of market size increases and decreases, respectively. These two parameters fully characterize nature's intensity matrix Q_0 . As mentioned above, we can obtain consistent estimates for q_1 and q_2 , so for our experiments we assume that Q_0 is known and focus on estimating firms' payoff parameters. In particular, we choose the intensity matrix Q_0 so that the discrete-time transition probability matrix for market size is as in Aguirregabiria and Mira (2007)⁶. Throughout the experiments, we fix the discount rate at $\rho_{ik} = 0.05$ and the move arrival rate at $\lambda_{ik} = 1$ for all $i = 1, \dots, N$ and $k = 1, \dots, K$.

We assume in this section (unless specified otherwise), that the DGP is a continuous time model, however, we consider two distinct sampling regimes for estimation. First, we consider continuous time data observed every instant. Since the model is a finite-state jump process, it is equivalent to observe each jump time and the resulting state. We refer to this as continuous time data. To generate continuous time data, we calculate the MPE of the game, obtain the steady-state distribution, and draw the initial states $k_{m,0}$ for each market m . From these initial states, we draw the subsequent actions and state $(a_{k_{m,n}}, x_{k_{m,n-1}})$ using the equilibrium conditional choice probabilities. Second, we also consider discrete time data observed at fixed intervals of length $\Delta = 1$. We refer to this as discrete time data, noting that the underlying DGP is a continuous time model. In this case, initial states are also drawn using the steady-state distribution and then we draw the future states according to the transition probability matrix over the time interval, $P(\Delta)$.

For each experiment, we carry out 100 replications using a sample of $M = 400$ markets

⁶The transition probability matrix used for market size s_k in Aguirregabiria and Mira (2007) is

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

in each. For simplicity, we focus on the case where there is one observed continuous-time event or discrete-time observation per market. Table 1 summarizes the DGP in the case of discretely-observed data. The parameters for the first three experiments differ only by the strategic interaction parameter, θ_{RN} , which is increasing from Experiment 1 to 3. Increasing θ_{RN} decreases the average number of active firms. In this setting, both the average number of entries and exits are increasing when θ_{RN} increases. As excess turnover increases, there will be a larger change in the number of active firms between periods. As a result, the AR(1) coefficient for the number of current active firms decreases. As θ_{RN} decreases the flow payoff for all firms, firms are less likely to be active regardless of fixed costs.

In Experiments 4–6, θ_{EC} varies from 0 to 4. The number of active firms increases as θ_{EC} increases, which might be surprising, but this is because the probability of being active differs across firms. Specifically, more efficient firms tend to stay active as the entry cost increases. This can be seen with Firms 4 and 5, which have lower fixed costs.

Table 1: DGP Steady State (Discrete Time Data)

	Exp. 1	Exp. 2	Exp. 3	Exp. 4	Exp. 5	Exp. 6
Descriptive Statistics	$\theta_{EC} = 1.0$ $\theta_{RN} = 0.0$	$\theta_{EC} = 1.0$ $\theta_{RN} = 1.0$	$\theta_{EC} = 1.0$ $\theta_{RN} = 2.0$	$\theta_{EC} = 0.0$ $\theta_{RN} = 1.0$	$\theta_{EC} = 2.0$ $\theta_{RN} = 1.0$	$\theta_{EC} = 4.0$ $\theta_{RN} = 1.0$
Average #active firms ¹	3.7107 (1.4427)	2.7744 (1.5338)	2.0468 (1.2510)	2.7351 (1.3921)	2.8027 (1.6612)	2.8214 (1.8139)
AR(1) for #active firms ²	0.8012	0.7879	0.6909	0.6720	0.8648	0.9381
Average #entrants	0.3783	0.5024	0.5388	0.6514	0.3653	0.1861
Average #exits	0.3779	0.5008	0.5385	0.6464	0.3667	0.1870
Excess turnover ³	0.2025	0.3096	0.3798	0.4770	0.1768	0.0413
Correlation between entries and exits	-0.0030	-0.0859	-0.0669	-0.1240	-0.0607	-0.0545
Prob. of being active						
Firm 1	0.7030	0.4980	0.3352	0.5032	0.4878	0.4567
Firm 2	0.7237	0.5286	0.3694	0.5263	0.5244	0.5045
Firm 3	0.7449	0.5530	0.4115	0.5468	0.5611	0.5549
Firm 4	0.7602	0.5806	0.4443	0.5687	0.5953	0.6179
Firm 5	0.7790	0.6141	0.4863	0.5902	0.6341	0.6875

¹ Values in parentheses are standard deviations.

² AR(1) for #active is the autoregressive coefficient regressing the number of current active firms on the number of active firms in previous period.

³ Excess turnover is defined as $(\#entrants + \#exits) - \text{abs}(\#entrants - \#exits)$.

4.2 Estimation Results

We consider several choices for the initial values of the CCPs: (i) true CCPs, (ii) frequency estimates, (iii) semi-parametric logit estimates (using market size and number of rival firms as explanatory variables), and (iv) random $U(0, 1)$ draws. Clearly (i) is infeasible and (iv) is inconsistent, but both are useful for comparison. For (ii), (iii), and (iv), we consider the two-step estimator ($l = 1$) and the converged CTNPL ($l = 20$).

Table 2 shows the means and standard deviations, in parentheses, of the Monte Carlo replications in the case of continuous time data. Based on the same replications, Table 3 displays the square-root of the mean square error (RMSE) of each estimator relative to the infeasible two-step PML estimates. Using these two tables we can compare the performance of the various estimators.

First, we note that the CTNPL estimator converged to the same estimates within 20 iterations regardless of which CCP estimates were used for initialization (frequency estimates, semi-parametric logit estimates, or random $U(0, 1)$ draws). Therefore, we report these in a single row denoted “CTNPL” in the table.

Of the two-step estimators, initializing using the semi-parametric logit estimates appears to be the best choice in terms of low biases, standard deviations, and therefore low relative RMSE values. Not surprisingly, using inconsistent, random $U(0, 1)$ starting values is a worse choice. Perhaps surprisingly, the frequency estimates are the worst choice for the initial CCPs. Even though frequency estimates are consistent, they result in a large bias with low variance across all parameters in Table 2 which in turn leads to high relative RMSE values in Table 3.

Overall, the performance of the CTNPL estimator—regardless of how it is initialized—is on par with the infeasible two-step PML estimator (which is initialized using the true CCPs). The RMSE values for the CTNPL estimator are close to one, meaning the estimators have similar performance, and in most cases the RMSEs are lower than those of the other feasible two-step estimators. For some parameters in some experiments, the two-step semi-parametric logit estimator performs slightly better than CTNPL, but the CTNPL estimator is much more robust than the two-step PML estimator in terms of how it is initialized. Even if the frequency estimates are used to initialize the CTNPL estimator, the finite sample bias can be reduced by repeatedly imposing the MPE conditions through iteration.

Table 4 and 5 report the similar results for estimation using discretely observed data. In this case, the results are even more distinct: the best two-step PML estimator (semi-parametric logit) exhibits more bias while the CTNPL estimator remains closer to the infeasible two-step estimator. The bias is particularly acute for the strategic interaction parameter

θ_{RN} .

4.3 Strategic Interaction and the Convergence Issue

In the specification of Experiment 1, the strategic interaction between agents, denoted by θ_{RN} , is set to zero which means that this setting corresponds to a collection of single agent models. In light of the zero Jacobian property of Proposition 1, the Ψ mapping is stable in this case and the algorithm will converge. However, as θ_{RN} increases it becomes more likely that the Ψ mapping is unstable (both in a finite sample and in the population). Kasahara and Shimotsu (2012) showed, using a two-firm example, that a MPE for which the local convergence holds exists if the contemporaneous and dynamic interaction between firms is small.

In Section 2.3, Kasahara and Shimotsu (2012) showed that $\rho(M_{\Psi_P}\Psi_P)$ is close to $\rho(\Psi_P)$ in a typical setting where the number of states is larger than the number of parameters to estimate. This is also true in continuous-time games, so similarly we will focus here on $\rho(\Psi_\sigma)$ given that Ψ_σ is often related to the characteristics of the economic model (Kasahara and Shimotsu, 2012).

Aguirregabiria and Marcoux (2021) study the relationship between θ_{RN} and the spectral radius $\rho(\Psi_P)$ for the discrete time NPL estimator using the five player example game of Aguirregabiria and Mira (2007). We conjecture that $\rho(\Psi_\sigma)$ is smaller in continuous time games than the spectral radius in a comparable discrete time game since continuous time games do not allow simultaneous moves between agents, making it more likely that the best response mapping is more stable. We investigate this in the context of the model used for our Monte Carlo experiments.

Figure 1 shows that the conjecture holds in the present model. Displayed in the top panel is our replication of a similar figure in Aguirregabiria and Marcoux (2021). This figure shows the relationship between the spectral radii for different values of θ_{RN} using a discrete time version of the five player game in Aguirregabiria and Mira (2007). In the figure, the spectral radius becomes larger than 1 when θ_{RN} is larger than 2.2. The lower panel of Figure 1 shows the relationship between the spectral radius and strategic interaction parameter in the continuous time model. To construct the figure, we use the same setting of Experiment 1 above and examine how $\rho(\Psi_\sigma)$ changes with respect to θ_{RN} . To produce the figure, we generate different data sets for each value of θ_{RN} and calculate the spectral radius of Ψ_σ . Even when θ_{RN} increases to 5 the spectral radius is 0.6, still far below the threshold for instability. On the right y-axis, we also report the average number of active firms in each game. As the strategic interaction parameter increases, the number of firms varies from

Table 2: Monte Carlo Results (Continuous Time Data)

Exp.	Estimator	Parameters			
		$\theta_{FC,1}$	θ_{RS}	θ_{EC}	θ_{RN}
1	True values	-1.9000	1.0000	1.0000	0.0000
	2S-True	-1.9112 (0.2933)	1.0052 (0.0954)	1.0274 (0.1342)	0.0169 (0.2802)
	2S-Freq	-0.3623 (0.2410)	0.3700 (0.1126)	1.7908 (0.1435)	0.3567 (0.3015)
	2S-Logit	-1.9136 (0.2948)	1.0071 (0.0952)	1.0272 (0.1346)	0.0194 (0.2805)
	2S-Random	-2.1992 (0.3977)	1.1214 (0.1031)	1.0314 (0.1352)	-0.0427 (0.3415)
	CTNPL	-1.9144 (0.2950)	1.0072 (0.0953)	1.0273 (0.1345)	0.0190 (0.2819)
2	True values	-1.9000	1.0000	1.0000	1.0000
	2S-True	-1.9099 (0.2009)	1.0016 (0.0953)	1.0183 (0.1306)	1.0137 (0.2643)
	2S-Freq	-0.8255 (0.1896)	0.4671 (0.1117)	1.4694 (0.1211)	0.5877 (0.2860)
	2S-Logit	-1.8932 (0.2113)	1.0005 (0.1023)	1.0185 (0.1304)	1.0251 (0.2807)
	2S-Random	-1.7843 (0.2889)	0.9587 (0.0724)	1.0224 (0.1309)	1.0140 (0.3030)
	CTNPL	-1.9163 (0.2078)	1.0032 (0.1030)	1.0184 (0.1304)	1.0145 (0.2761)
3	True values	-1.9000	1.0000	1.0000	2.0000
	2S-True	-1.9401 (0.1770)	1.0117 (0.0808)	1.0070 (0.0960)	2.0182 (0.2665)
	2S-Freq	-1.1756 (0.2523)	0.5865 (0.1211)	1.3317 (0.1055)	1.1303 (0.3315)
	2S-Logit	-1.8700 (0.2034)	1.0077 (0.0916)	1.0070 (0.0962)	2.0714 (0.2999)
	2S-Random	-1.2428 (0.2885)	0.8001 (0.0624)	1.0215 (0.0928)	1.8789 (0.3228)
	CTNPL	-1.9488 (0.2013)	1.0147 (0.0930)	1.0070 (0.0961)	2.0196 (0.2867)
4	True values	-1.9000	1.0000	0.0000	1.0000
	2S-True	-1.9208 (0.2382)	1.0219 (0.0950)	-0.0071 (0.1149)	1.0665 (0.3163)
	2S-Freq	-0.9979 (0.2103)	0.5025 (0.1096)	0.4399 (0.1097)	0.6593 (0.3496)
	2S-Logit	-1.9081 (0.2458)	1.0233 (0.1015)	-0.0073 (0.1146)	1.0835 (0.3347)
	2S-Random	-1.5546 (0.2965)	0.8156 (0.0697)	0.0249 (0.1138)	0.8344 (0.2922)
	CTNPL	-1.9260 (0.2410)	1.0247 (0.1021)	-0.0071 (0.1149)	1.0727 (0.3297)
5	True values	-1.9000	1.0000	2.0000	1.0000
	2S-True	-1.9320 (0.1902)	1.0175 (0.1033)	2.0155 (0.1334)	1.0336 (0.2189)
	2S-Freq	-0.7240 (0.2086)	0.4921 (0.1171)	2.4653 (0.1393)	0.6563 (0.2428)
	2S-Logit	-1.9079 (0.2029)	1.0165 (0.1118)	2.0152 (0.1337)	1.0493 (0.2381)
	2S-Random	-2.1011 (0.3297)	1.1627 (0.0987)	1.9982 (0.1292)	1.3001 (0.3330)
	CTNPL	-1.9416 (0.2045)	1.0219 (0.1132)	2.0155 (0.1336)	1.0390 (0.2346)
6	True values	-1.9000	1.0000	4.0000	1.0000
	2S-True	-1.9277 (0.1966)	1.0167 (0.1094)	4.0347 (0.2232)	1.0180 (0.1934)
	2S-Freq	-0.5088 (0.1907)	0.4460 (0.1182)	4.4598 (0.2255)	0.5811 (0.2107)
	2S-Logit	-1.8957 (0.2064)	1.0113 (0.1205)	4.0319 (0.2237)	1.0291 (0.2098)
	2S-Random	-2.8593 (0.4994)	1.6734 (0.1611)	3.9831 (0.2210)	2.0784 (0.5526)
	CTNPL	-1.9448 (0.2087)	1.0214 (0.1206)	4.0332 (0.2232)	1.0180 (0.2040)

Displayed values are means with standard deviations in parentheses.

Table 3: RMSE Relative to Two-Step PML with True CCPs (Continuous Time Data)

Exp.	Estimator	Parameters			
		$\theta_{FC,1}$	θ_{RS}	θ_{EC}	θ_{RN}
1	2S-Freq	5.3031	6.6989	5.8694	1.6638
	2S-Logit	1.0056	0.9992	1.0027	1.0015
	2S-Random	1.6955	1.6671	1.0137	1.2260
	CTNPL	1.0064	1.0004	1.0022	1.0063
2	2S-Freq	5.4232	5.7112	3.6768	1.8962
	2S-Logit	1.0506	1.0727	0.9987	1.0650
	2S-Random	1.5469	0.8746	1.0070	1.1461
	CTNPL	1.0362	1.0806	0.9991	1.0450
3	2S-Freq	4.2254	5.2762	3.6157	3.4849
	2S-Logit	1.1325	1.1259	1.0016	1.1543
	2S-Random	3.9534	2.5644	0.9896	1.2909
	CTNPL	1.1408	1.1534	1.0012	1.0759
4	2S-Freq	3.8739	5.2251	3.9373	1.5103
	2S-Logit	1.0287	1.0685	0.9973	1.0674
	2S-Random	1.9037	2.0219	1.0114	1.0392
	CTNPL	1.0140	1.0778	0.9993	1.0445
5	2S-Freq	6.1921	4.9729	3.6155	1.9001
	2S-Logit	1.0525	1.0779	1.0018	1.0978
	2S-Random	2.0023	1.8152	0.9619	2.0241
	CTNPL	1.0818	1.0998	1.0015	1.0738
6	2S-Freq	7.0717	5.1198	2.2676	2.4137
	2S-Logit	1.0397	1.0937	1.0005	1.0901
	2S-Random	5.4465	6.2583	0.9815	6.2374
	CTNPL	1.0752	1.1074	0.9993	1.0540

Table 4: Monte Carlo Results (Discrete Time Data)


Exp.	Estimator	Parameters			
		$\theta_{FC,1}$	θ_{RS}	θ_{EC}	θ_{RN}
1	True values	-1.9000	1.0000	1.0000	0.0000
	2S-True	-1.8696 (0.4613)	1.0321 (0.1864)	1.0126 (0.3056)	0.0710 (0.4025)
	2S-Freq	-0.2995 (0.2997)	0.3971 (0.1259)	1.9794 (0.3101)	0.3994 (0.2739)
	2S-Logit	-1.7440 (0.4069)	0.9755 (0.1655)	1.0055 (0.3129)	0.0860 (0.3060)
	2S-Random	-2.1636 (0.5416)	1.1508 (0.1539)	1.0134 (0.2922)	-0.0045 (0.4948)
	CTNPL	-1.8830 (0.4711)	1.0383 (0.1952)	1.0146 (0.3040)	0.0747 (0.4094)
2	True values	-1.9000	1.0000	1.0000	1.0000
	2S-True	-1.9433 (0.3396)	1.0385 (0.1613)	0.9871 (0.2539)	1.0938 (0.3765)
	2S-Freq	-0.6917 (0.2469)	0.4395 (0.1124)	1.6766 (0.2417)	0.5812 (0.2729)
	2S-Logit	-1.7812 (0.3104)	0.8945 (0.1339)	0.9858 (0.2555)	0.8286 (0.2973)
	2S-Random	-1.7902 (0.4500)	0.9946 (0.1114)	0.9748 (0.2538)	1.1124 (0.4437)
	CTNPL	-1.9558 (0.3573)	1.0461 (0.1775)	0.9880 (0.2527)	1.1051 (0.4033)
3	True values	-1.9000	1.0000	1.0000	2.0000
	2S-True	-1.9298 (0.2940)	1.0302 (0.1646)	0.9909 (0.2158)	2.0662 (0.4787)
	2S-Freq	-0.8447 (0.2504)	0.4236 (0.1133)	1.5584 (0.2076)	0.7669 (0.3106)
	2S-Logit	-1.6782 (0.2705)	0.8325 (0.1329)	0.9908 (0.2160)	1.5422 (0.3803)
	2S-Random	-1.2187 (0.3895)	0.8255 (0.1008)	0.9921 (0.2087)	1.9675 (0.5099)
	CTNPL	-1.9521 (0.3397)	1.0416 (0.1953)	0.9912 (0.2159)	2.0869 (0.5427)
4	True values	-1.9000	1.0000	0.0000	1.0000
	2S-True	-1.9361 (0.4290)	0.9941 (0.1813)	0.0073 (0.3143)	0.9462 (0.4605)
	2S-Freq	-0.6702 (0.2656)	0.3328 (0.0856)	0.8494 (0.2739)	0.3373 (0.2498)
	2S-Logit	-1.7521 (0.3658)	0.8327 (0.1387)	0.0044 (0.3153)	0.6578 (0.3278)
	2S-Random	-1.6005 (0.4331)	0.8264 (0.1074)	0.0254 (0.2973)	0.7897 (0.4397)
	CTNPL	-1.9443 (0.4427)	0.9993 (0.1939)	0.0068 (0.3146)	0.9524 (0.4843)
5	True values	-1.9000	1.0000	2.0000	1.0000
	2S-True	-1.9124 (0.2865)	1.0140 (0.1627)	1.9971 (0.2054)	1.0125 (0.3300)
	2S-Freq	-0.6527 (0.2805)	0.4642 (0.1309)	2.5805 (0.2464)	0.5930 (0.2421)
	2S-Logit	-1.7567 (0.2654)	0.8959 (0.1376)	1.9996 (0.2050)	0.8150 (0.2730)
	2S-Random	-2.0970 (0.3999)	1.1894 (0.1401)	1.9635 (0.2184)	1.3346 (0.4834)
	CTNPL	-1.9285 (0.3105)	1.0209 (0.1811)	1.9961 (0.2065)	1.0190 (0.3562)
6	True values	-1.9000	1.0000	4.0000	1.0000
	2S-True	-1.9634 (0.2460)	1.0438 (0.1470)	4.0477 (0.2482)	1.0649 (0.2560)
	2S-Freq	-0.7220 (0.3098)	0.5674 (0.1777)	4.4336 (0.2896)	0.7483 (0.2930)
	2S-Logit	-1.8252 (0.2365)	0.9631 (0.1388)	4.0460 (0.2452)	0.9616 (0.2545)
	2S-Random	-2.9300 (0.5703)	1.8220 (0.2097)	3.9921 (0.2574)	2.4322 (0.7285)
	CTNPL	-1.9921 (0.2694)	1.0570 (0.1675)	4.0473 (0.2497)	1.0799 (0.2934)

Displayed values are means with standard deviations in parentheses.


Table 5: RMSE Relative to Two-Step PML with True CCPs (Discrete Time Data)

Exp.	Estimator	Parameters			
		$\theta_{FC,1}$	θ_{RS}	θ_{EC}	θ_{RN}
1	2S-Freq	3.5224	3.2572	3.3591	1.1850
	2S-Logit	0.9426	0.8844	1.0232	0.7777
	2S-Random	1.3030	1.1393	0.9565	1.2108
	CTNPL	1.0197	1.0519	0.9951	1.0182
2	2S-Freq	3.6028	3.4469	2.8259	1.2885
	2S-Logit	0.9709	1.0280	1.0065	0.8844
	2S-Random	1.3533	0.6723	1.0030	1.1797
	CTNPL	1.0566	1.1059	0.9949	1.0742
3	2S-Freq	3.6707	3.5096	2.7585	2.6315
	2S-Logit	1.1837	1.2774	1.0009	1.2316
	2S-Random	2.6560	1.2038	0.9673	1.0574
	CTNPL	1.1632	1.1926	1.0006	1.1373
4	2S-Freq	2.9227	3.7093	2.8391	1.5275
	2S-Logit	0.9166	1.1981	1.0030	1.0222
	2S-Random	1.2232	1.1255	0.9492	1.0514
	CTNPL	1.0335	1.0692	1.0012	1.0497
5	2S-Freq	4.4578	3.3784	3.0697	1.4344
	2S-Logit	1.0517	1.0566	0.9978	0.9989
	2S-Random	1.5544	1.4431	1.0777	1.7804
	CTNPL	1.0871	1.1165	1.0055	1.0802
6	2S-Freq	4.7959	3.0490	2.0634	1.4625
	2S-Logit	0.9768	0.9366	0.9873	0.9747
	2S-Random	4.6357	5.5309	1.0191	6.0838
	CTNPL	1.1210	1.1537	1.0055	1.1515

Figure 1: Spectral Radius vs. θ_{RN}



(a) Discrete Time

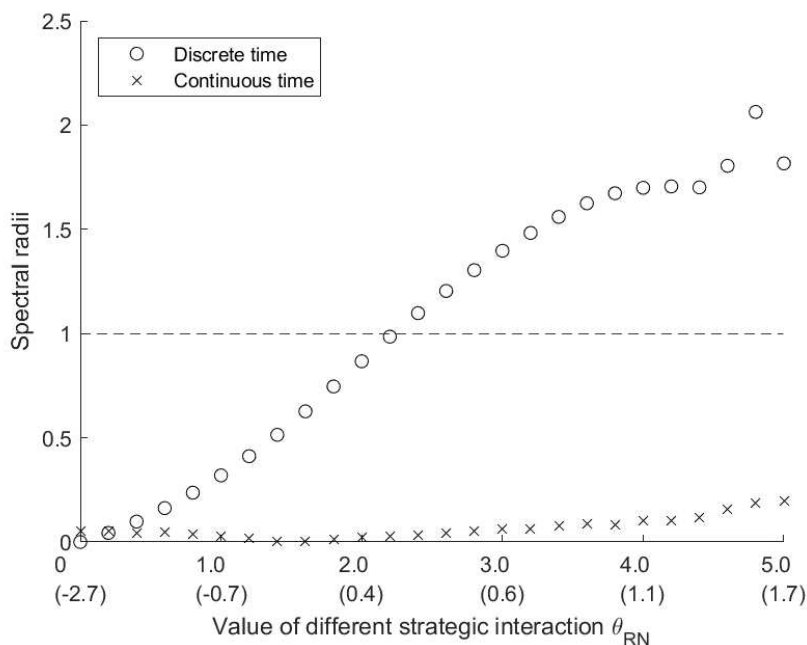


(b) Continuous Time

almost 4 to around 1.

Although we can see that the number of active firms in the discrete time and continuous time models is similar, the meaning of a particular value for the strategic interaction parameter θ_{RN} is different in the two models. In order to provide another means to compare the spectral radii, we also compare the discrete time model with the best-fit continuous time model in Figure 2. To do so, we use the same discrete time data for each replication and use it to estimate the continuous time model. Then we calculate the Jacobian Ψ_σ as before and calculate the spectral radius.

Figure 2: Spectral Radius vs θ_{RN} : Best Fit Comparison



Note: The numbers in parentheses are corresponding estimates for θ_{RN} from a continuous time game.

As such, the x-axis of Figure 2 has two labels. The upper labels denote the original values of θ_{RN} used to generate the discrete time data. The numbers in parentheses below are the corresponding estimates of θ_{RN} from the continuous time model. Note that when θ_{RN} is zero in continuous time, the spectral radius is also zero as established in Proposition 1. Note that the continuous time estimates for the strategic interaction parameter are lower, and sometimes even negative, but the spectral radii are calculated using the absolute values of the eigenvalues so they are always positive. As θ_{RN} increases to five, spectral radius is at most around 0.16, even farther from one than before, so it is highly unlikely that there will be any convergence issues. As a result, in this example we can allow for a wider range of strategic interaction values than in discrete time without worrying about convergence to an

inconsistent estimator.

4.4 Estimation of Misspecified Discrete Time Models

Finally, in this section we consider what happens if one estimates a discrete time model when in reality the DGP is a continuous time model. We report estimates for the misspecified discrete time NPL estimator in Table 6 along with the correctly specified CTNPL estimates for comparison (the latter being reproduced from Table 4).

There is bias in all parameters as expected, however it is important to note that there are severe biases in the economically important strategic interaction and entry cost parameters. It is also of interest to learn how these biases behave in different settings. In particular, the biases tend to be larger when entry costs are smaller.

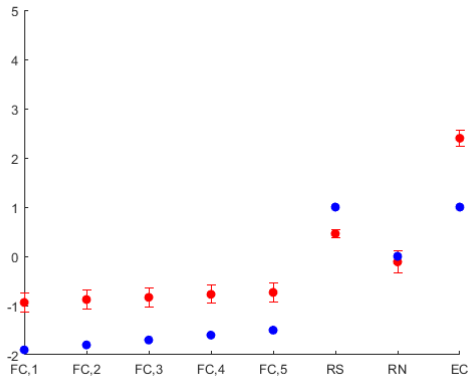
Although the estimates for θ_{RS} have small standard deviations, the biases are larger than those of estimates from correctly specified models in all experiments. This result suggests that when a researcher wrongly chooses a discrete time model to estimate parameters from continuous data, it is possible to arrive at inconsistent estimates with large finite sample bias although seeming precise with small standard errors.

The underlying issue is that the discrete time model conflates higher rates of entry due to low entry costs with higher rates of entry due to low levels of competition. To see this, consider a two-firm model and suppose the market is empty at time $t = 0$. Suppose firm 1 first enters at time $t = 0.1$ and firm 2 later decides not to enter at time $t = 0.2$. Suppose that this market structure remains until time $t = 1$. Although firm 1 has spent nearly the entire unit of time in the market, the discrete time model interprets this as no firms in the market at time $t = 0$ and a firm-1 monopoly beginning at time 1. Through the discrete time lens, firm 2 chose to not enter an empty market and therefore entry costs must be high. On the other hand, the continuous time model—even with discrete time data—allows for the possibility that firm 1 entered earlier and firm 2 may have chosen to remain out because of the competitive effect and the presence of firm 1 rather than high entry costs.

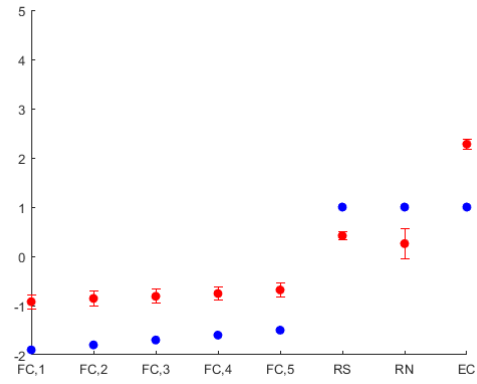
Figure 3 illustrates the possibility of misleading estimates. We plot the misspecified discrete time model estimates from Experiment 1 and the true values. Even though the confidence intervals for the discrete time estimates are small, the estimates are far from the true parameters. This suggests that a researcher should carefully specify the model by considering whether sequential moves are important and understanding how and when the state variables change.

We can also examine how this problem affects counterfactuals. Dynamic game estimation results are often used for estimating the impact of a policy affecting firms' entry-exit


Figure 3: Misspecified Discrete Time Model Estimates



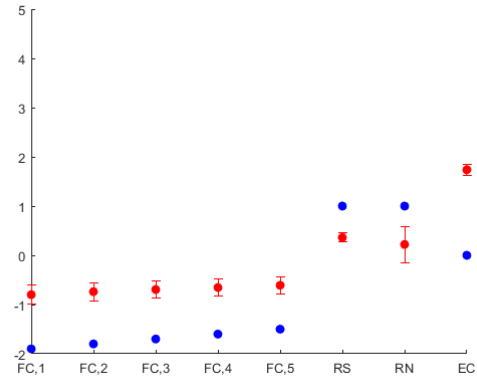
(a)



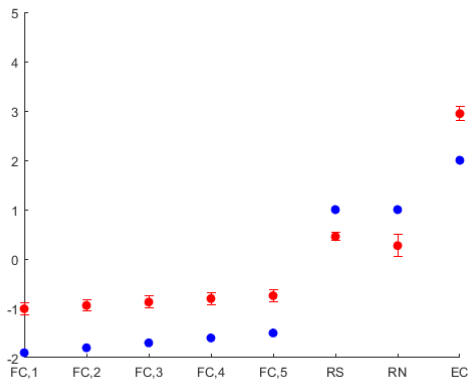
(b)



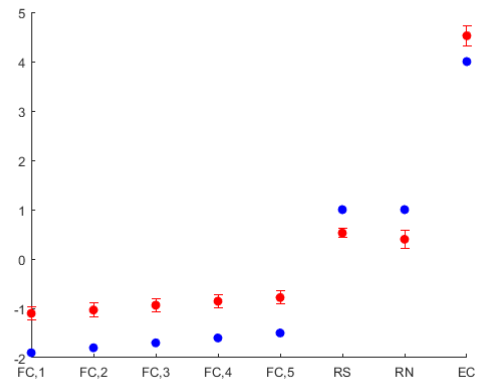
(c)



(d)



(e)



(f)

Note: Blue dots are the true parameters from the continuous time model and red dots are the estimates and 95% confidence intervals from the estimated discrete time model.

Table 6: Misspecified Monte Carlo Results (Discrete Time Estimation of Continuous Time DGP)

Exp.	Values	Parameters			
		$\theta_{FC,1}$	θ_{RS}	θ_{EC}	θ_{RN}
1	True values	-1.9000	1.0000	1.0000	0.0000
	Correct	-1.8830 (0.4711)	1.0383 (0.1952)	1.0146 (0.3040)	0.0747 (0.4094)
	Misspecified	-0.9369 (0.1928)	0.4604 (0.0762)	2.3993 (0.1555)	-0.1135 (0.2300)
	Bias	0.9631	-0.5396	1.3993	-0.1135
2	True values	-1.9000	1.0000	1.0000	1.0000
	Correct	-1.9558 (0.3573)	1.0461 (0.1775)	0.9880 (0.2527)	1.1051 (0.4033)
	Misspecified	-0.9221 (0.1444)	0.4176 (0.0889)	2.2785 (0.1083)	0.2574 (0.3024)
	Bias	0.9779	-0.5824	1.2785	-0.7426
3	True values	-1.9000	1.0000	1.0000	2.0000
	Correct	-1.9521 (0.3397)	1.0416 (0.1953)	0.9912 (0.2159)	2.0869 (0.5427)
	Misspecified	-0.8631 (0.1433)	0.4615 (0.0999)	2.1725 (0.1188)	0.9247 (0.3744)
	Bias	1.0369	-0.5385	1.1725	-1.0753
4	True values	-1.9000	1.0000	0.0000	1.0000
	Correct	-1.9443 (0.4427)	0.9993 (0.1939)	0.0068 (0.3146)	0.9524 (0.4843)
	Misspecified	-0.7994 (0.1909)	0.3607 (0.0887)	1.7383 (0.1084)	0.2212 (0.3696)
	Bias	1.1006	-0.6393	1.7383	-0.7788
5	True values	-1.9000	1.0000	2.0000	1.0000
	Correct	-1.9285 (0.3105)	1.0209 (0.1811)	1.9961 (0.2065)	1.0190 (0.3562)
	Misspecified	-1.0088 (0.1240)	0.4522 (0.0797)	2.9451 (0.1373)	0.2705 (0.2210)
	Bias	0.8912	-0.5478	0.9451	-0.7295
6	True values	-1.9000	1.0000	4.0000	1.0000
	Correct	-1.9921 (0.2694)	1.0570 (0.1675)	4.0473 (0.2497)	1.0799 (0.2934)
	Misspecified	-1.1024 (0.1418)	0.5281 (0.0937)	4.5233 (0.2036)	0.3995 (0.1889)
	Bias	0.7976	-0.4719	0.5233	-0.6005

Displayed values are means with standard deviations in parentheses. Biases are differences between the estimated parameters and the true parameter values.

decisions. Suppose that the entry costs are one million dollars and the government is considering subsidizing firms by two hundred thousand dollars. In particular, the policy would reduce the fixed costs by \$200K. Since the correct model and the misspecified model have different estimates for entry costs and fixed costs, the policy will result in different steady state outcomes.

To compare the different counterfactual predictions, we change the parameter values according to the counterfactual policy, recalculate the equilibrium, and compare the results over 50,000 simulations. Table 7 presents the average number of active firms before and after subsidization. The first column is the average number of active firms before firms receiving the subsidy. Then, in Experiment 1, if firms receive the subsidy the continuous time model predicts that the number of active firms increase by 5.6% but the misspecified discrete time model suggests 20.4% increase in the number of active firms. The prediction from the discrete time model is four times larger. This overestimation of the effect when using the misspecified model also occurs for the other experiments.

Table 7: Average Number of Active Firms Before and After Subsidization

Exp.	Before subsidization		After subsidization				Δ #Active firms (%)	
	Mean	(s.d.)	Correct		Misspecified		Correct	Misspecified
			Mean	(s.d.)	Mean	(s.d.)		
1	3.70	(1.47)	3.91	(1.33)	4.46	(0.83)	5.6	20.4
2	2.77	(1.53)	2.96	(1.51)	3.63	(1.32)	6.7	30.8
3	2.05	(1.25)	2.20	(1.27)	2.68	(1.37)	7.6	31.1
4	2.74	(1.39)	-	-	-	-	-	-
5	2.80	(1.66)	3.20	(1.58)	3.94	(1.23)	14.2	40.5
6	2.82	(1.81)	3.67	(1.48)	4.47	(0.85)	30.1	58.6

5 Conclusion

This paper introduced an NPL estimator for dynamic discrete choice models in continuous time, which we refer to as the CTNPL estimator. We derived its properties, and demonstrated the performance in a series of Monte Carlo experiments involving a model with five heterogeneous firms. Specifically, we first showed that the CTNPL estimator is consistent and asymptotically normal both with and without initial consistent and asymptotically normal CCP estimates. Second, we presented a local convergence condition in the iterative CTNPL algorithm. Researchers have documented problems regarding the convergence of the NPL algorithm in discrete time. We showed that the algorithm always converges in

continuous-time single agent models and that provided simulation evidence that convergence failures are much less likely to affect comparable continuous-time games. Third, our Monte Carlo experiments based on those of Aguirregabiria and Mira (2007) showed that the CT-NPL estimator is more robust than two-step estimators initialized from estimates or random draws for the CCPs. Finally, our results highlight the potential for economically misleading estimates and counterfactuals when estimating discrete time models in cases where events in the data generating process are unfolding in continuous time.

References

- Aguirregabiria, V. and M. Marcoux (2021). Imposing equilibrium restrictions in the estimation of dynamic discrete games. *Working Paper*.
- 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.
- Aguirregabiria, V. and P. Mira (2010). Dynamic discrete choice structural models: A survey. *Journal of Econometrics* 156(1), 38 – 67. Structural Models of Optimization Behavior in Labor, Aging, and Health.
- Amemiya, T. (1985). *Advanced Econometrics*. Harvard University Press.
- Andrews, D. W. K. (1992). Generic uniform convergence. *Econometric Theory* 8(2), 241–257.
- Arcidiacono, P., P. Bayer, J. R. Blevins, and P. B. Ellickson (2016). Estimation of dynamic discrete choice models in continuous time with an application to retail competition. *Review of Economic Studies* 83, 889–931.
- Arcidiacono, P. S., A. Gyetvai, E. Jardim, and A. Maurel (2019). Conditional choice probability estimation of continuous-time job search models. Working paper, Duke University.
- Benkard, C. L. (2004). A dynamic analysis of the market for wide-bodied commercial aircraft. *Review of Economic Studies* 71(3), 581–611.
- Blevins, J. R. (2016). Identification and estimation of continuous time dynamic discrete choice games. Working paper, Ohio State University.

- Bugni, F. A. and J. Bunting (2021). On the iterated estimation of dynamic discrete choice games. *Review of Economic Studies* 88, 1031–1073.
- Cosman, J. (2019). Industry dynamics and the value of variety in nightlife: Evidence from Chicago. *Working Paper*.
- Deng, Y. and C. F. Mela (2018). TV viewing and advertising targeting. *Journal of Marketing Research* 55(1), 99–118.
- Doraszelski, U. and K. L. Judd (2012). Avoiding the curse of dimensionality in dynamic stochastic games. *Quantitative Economics* 3(1), 53–93.
- Doraszelski, U. and A. Pakes (2007). A framework for applied dynamic analysis in IO. In M. Armstrong and R. Porter (Eds.), *Handbook of Industrial Organization* (1 ed.), Volume 3, Chapter 30, pp. 1887–1966. Elsevier.
- Egedal, M., Z. Lai, and C.-L. Su (2015). Estimating dynamic discrete-choice games of incomplete information. *Quantitative Economics* 6(3), 567–597.
- Gourieroux, C. and A. Monfort (1995). *Statistics and Econometric Models*, Volume II. Cambridge University Press.
- Gowrisankaran, G. and M. Rysman (2012). Dynamics of consumer demand for new durable goods. *Journal of Political Economy* 120(6), 1173–1219.
- Hendel, I. and A. Nevo (2006). Measuring the implications of sales and consumer inventory behavior. *Econometrica* 74(6), 1637–1673.
- Holmes, T. J. (2011). The diffusion of Wal-Mart and economies of density. *Econometrica* 79(1), 253–302.
- Horn, R. A. (1995). *Matrix Analysis*. Cambridge University Press.
- Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* 60, 497–529.
- Jennrich, R. I. (1969). Asymptotic properties of non-linear least squares estimators. *Annals of Mathematical Statistics* 40, 633–643.
- Kawahara, H. and K. Shimotsu (2012). Sequential estimation of structural models with a fixed point constraint. *Econometrica* 80(5), 2303–2319.

- Mazur, J. (2017). Can stricter bankruptcy laws discipline capital investment? Evidence from the U.S. airline industry. Working paper.
- Miller, R. A. (1984). Job matching and occupational choice. *Journal of Political Economy* 92(6), 1086–1120.
- Nevskaya, Y. and P. Albuquerque (2019). How should firms manage excessive product use? A continuous-time demand model to test reward schedules, notifications, and time limits. *Journal of Marketing Research* 56(3), 379–400.
- Pakes, A. (1986). Patents as options: Some estimates of the value of holding European patent stocks. *Econometrica* 54(4), 755–784.
- Pesendorfer, M. and P. Schmidt-Dengler (2008). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies* 75, 901–928.
- Pesendorfer, M. and P. Schmidt-Dengler (2010). Sequential estimation of dynamic discrete games: A comment. *Econometrica* 78(2), 833–842.
- Qin, M. S. (2017). Planes, trains, and co-opetition: Evidence from China. Working paper.
- Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica* 55(5), 999–1033.
- Schiraldi, P. (2011). Automobile replacement: A dynamic structural approach. *The RAND Journal of Economics* 42(2), 266–291.
- Schiraldi, P., H. Smith, and Y. Takahashi (2012). Estimating a dynamic game of spatial competition: The case of the U.K. supermarket industry. Working paper.
- Schuette, C. and P. Metzner (2009, July). Markov chains and jump processes: An introduction to Markov chains and jump processes on countable state spaces.
- Song, I. and P. K. Chintagunta (2003, Dec). A micromodel of new product adoption with heterogeneous and forward-looking consumers: Application to the digital camera category. *Quantitative Marketing and Economics* 1(4), 371–407.
- Wolpin, K. I. (1984). An estimable dynamic stochastic model of fertility and child mortality. *Journal of Political Economy* 92(5), 852–874.

A Proofs

A.1 Proof for Proposition 1

We introduce notation from ABBE (2016), prove two preliminary Lemmas, and then prove the zero Jacobian property in a single agent dynamic choice model in continuous time.

For choice j in state k , let $v_{jk} = \psi_{jk} + V_{l(j,k)}$ denote an arbitrary choice-specific valuation and $v_k = (v_{0k}, \dots, v_{J-1,k}) \in \mathbb{R}^J$ denote a J -vector of valuations in state k . Define \tilde{v}_k as a $(J-1)$ -vector of normalized valuations based on choice 0 on the $(J-1)$ -dimensional space $\mathcal{V} = \{\tilde{v}_k \in \mathbb{R}^{J-1} : \tilde{v}_{0k} = 0\}$.

Now consider the mapping $\tilde{H}_k : \mathcal{V} \rightarrow \mathcal{P}_k$ where $\mathcal{P}_k \in \Delta^{J-1}$ is the space of all CCP vectors σ_k in state k . By Proposition 1 of Hotz and Miller (1993), \tilde{H}_k is one-to-one with $\tilde{H}_k(\sigma_k) = \tilde{v}_k$ and invertible.

Lemma 1. *Let σ_k^0 denote an arbitrary vector of $J-1$ choice probabilities in state k .*

$$\frac{\partial}{\partial \sigma_k} \sum_j \sigma_{jk}^0 e_{jk}(\theta, \sigma_k^0) = -\tilde{H}_k^{-1}(\sigma_k^0) \quad (15)$$

where $e_{jk}(\theta, \sigma)$ is the expected value of ε_{jk} given that choice j is optimal at state k .

Proof. Let $W_{jk}(\tilde{v})$ represent the expectation e_{jk} as a function of $J-1$ normalized valuations. Since $\tilde{H}(\cdot)$ is invertible, we can write $e_{jk}(\sigma^0) = W(\tilde{H}_k^{-1}(\sigma^0))$.

$$\begin{aligned} \sum_{j=1}^J \sigma_j^0 e_{jk}(\theta, \sigma^0) &= \sum_{j=1}^J \sigma_{jk}^0 W_{jk}(\tilde{H}^{-1}(\sigma^0)) \\ &= \sigma^0 \cdot W(\tilde{H}^{-1}(\sigma^0)) \quad \text{in matrix form} \end{aligned}$$

Differentiating with respect to $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_J)$,

$$\frac{\partial}{\partial \sigma} \sigma^0 \cdot W(\tilde{H}^{-1}(\sigma^0)) = \frac{\partial \sigma^0}{\partial \sigma} W(\tilde{v}) + \frac{\partial W(\tilde{v})}{\partial \tilde{H}^{-1}(\sigma)} \frac{\partial \tilde{H}^{-1}(\sigma)}{\partial \sigma} \frac{\partial \sigma}{\partial \sigma^0} \sigma^0$$

Each term becomes

$$1. \quad \frac{\partial \sigma^0}{\partial \sigma} W(\tilde{v}) = \frac{\partial(\sigma_1, \dots, \sigma_{J-1})'}{\partial(\sigma_0, \dots, \sigma_{J-1})'} W(\tilde{v}) = \begin{bmatrix} -1 & -1 & -1 & \dots & -1 \\ 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \end{bmatrix} W(\tilde{v}) = [-i; I]W(\tilde{v})$$

$$\begin{aligned}
2. \quad & \frac{\partial W(\tilde{v})}{\partial \tilde{H}^{-1}(\sigma)} = \frac{\partial W(\tilde{v})}{\partial \tilde{v}} \\
3. \quad & \frac{\partial \tilde{H}^{-1}(\sigma)}{\partial \sigma} = \frac{\partial \tilde{v}}{\partial \tilde{H}(\tilde{v})} = \left[\frac{\partial \tilde{H}(\tilde{v})}{\partial \tilde{v}} \right]^{-1} \\
4. \quad & \frac{\partial \sigma}{\partial \sigma^0} \sigma^0 = [-i; I] \sigma^0 = \begin{bmatrix} -\sigma_1 - \dots - \sigma_{J-1} \\ \sigma_1 \\ \vdots \\ \sigma_{J-1} \end{bmatrix} = \begin{pmatrix} 1 - i' \sigma^0 \\ \sigma^0 \end{pmatrix}
\end{aligned}$$

where i is a $1 \times (J-1)$ vector of ones and I is $(J-1) \times (J-1)$ identity matrix. From 1–4, we get an analogous equation to (Ap.1) in Aguirregabiria and Mira (2002).

$$\frac{\partial}{\partial \sigma} \sum_{j=1}^J \sigma_j^0 e_j = [-i; I] W(\tilde{v}) + \left[\frac{\partial \tilde{H}(\tilde{v})}{\partial \tilde{v}} \right]^{-1} \frac{\partial W(\tilde{v})'}{\partial \tilde{v}} \begin{pmatrix} 1 - i' \sigma^0 \\ \sigma^0 \end{pmatrix} \quad (16)$$

From the surplus function,

$$S(\tilde{v}) = [1 - i' \tilde{H}(\tilde{v}); \tilde{H}(\tilde{v})](v + W(\tilde{v})).$$

Differentiating with respect to \tilde{v} ,

$$\begin{aligned}
\sigma &= \frac{\partial S(\tilde{v})}{\partial \tilde{v}} = \begin{bmatrix} -i' \frac{\partial \tilde{H}(\tilde{v})}{\partial \tilde{v}} \\ \frac{\partial \tilde{H}(\tilde{v})}{\partial \tilde{v}} \end{bmatrix} [v + W(\tilde{v})] + \left[\frac{\partial v}{\partial \tilde{v}} + \frac{\partial W(\tilde{v})}{\partial \tilde{v}} \right] \begin{bmatrix} 1 - i' \tilde{H}(\tilde{v}) \\ \tilde{H}(\tilde{v}) \end{bmatrix} \\
&= \frac{\partial \tilde{H}(\tilde{v})}{\partial \tilde{v}} \begin{bmatrix} -i \\ I \end{bmatrix} [v + W(\tilde{v})] + \left[\frac{\partial v}{\partial \tilde{v}} + \frac{\partial W(\tilde{v})}{\partial \tilde{v}} \right] \begin{bmatrix} 1 - i' \sigma^0 \\ \sigma^0 \end{bmatrix} \\
\left[\frac{\partial \tilde{H}(\tilde{v})}{\partial \tilde{v}} \right]^{-1} \sigma &= \begin{bmatrix} -i \\ I \end{bmatrix} [v + W(\tilde{v})] + \left[\frac{\partial \tilde{H}(\tilde{v})}{\partial \tilde{v}} \right]^{-1} \left[\frac{\partial v}{\partial \tilde{v}} + \frac{\partial W(\tilde{v})}{\partial \tilde{v}} \right] \begin{bmatrix} 1 - i' \sigma^0 \\ \sigma^0 \end{bmatrix} \\
&= \begin{bmatrix} -i \\ I \end{bmatrix} v + \left[\frac{\partial \tilde{H}(\tilde{v})}{\partial \tilde{v}} \right]^{-1} \frac{\partial v}{\partial \tilde{v}} \begin{bmatrix} 1 - i' \sigma^0 \\ \sigma^0 \end{bmatrix} + \begin{bmatrix} -i \\ I \end{bmatrix} W(\tilde{v}) + \left[\frac{\partial \tilde{H}(\tilde{v})}{\partial \tilde{v}} \right]^{-1} \frac{\partial W(\tilde{v})'}{\partial \tilde{v}} \begin{bmatrix} 1 - i' \sigma^0 \\ \sigma^0 \end{bmatrix} \\
&= \begin{bmatrix} -i \\ I \end{bmatrix} v + \left[\frac{\partial \tilde{H}(\tilde{v})}{\partial \tilde{v}} \right]^{-1} \begin{bmatrix} -i \\ I \end{bmatrix} \begin{bmatrix} 1 - i' \sigma^0 \\ \sigma^0 \end{bmatrix} + \frac{\partial}{\partial \sigma} \sum_{j=1}^J \sigma_j^0 e_j \text{ from (16)} \\
&= \begin{bmatrix} -i \\ I \end{bmatrix} v + \left[\frac{\partial \tilde{H}(\tilde{v})}{\partial \tilde{v}} \right]^{-1} \sigma + \frac{\partial}{\partial \sigma} \sum_{j=1}^J \sigma_j^0 e_j
\end{aligned}$$

Rearranging, we have the stated result:

$$\frac{\partial}{\partial \sigma} \sum_{j=1}^J \sigma_j^0 e_j(\theta, \sigma^0) = -\tilde{v} = -\tilde{H}^{-1}(\sigma^0).$$

■

Lemma 2. *Let σ^0 be an arbitrary set of conditional choice probabilities, and define the mapping*

$$G(\sigma^0, V) = [(\rho + \lambda)I - \lambda\Sigma(\sigma) - Q_0]^{-1}[\lambda\Sigma(\sigma)V + u + \lambda E(\theta, \sigma^0)].$$

Let σ_k^0 be $(J-1)$ column vector of CCPs in σ^0 associated with state k . Then (i) $\partial G_l / \partial \sigma_k^0 = 0$ for $l \neq k$ and (ii) $\partial G_k / \partial \sigma_k^0 = [(\rho + \lambda)I - \lambda\Sigma(\sigma) - Q_0]^{-1}[\lambda[\tilde{v} - \tilde{H}_k^{-1}(\sigma^0)]]$.

Proof. (i) Given that G_k does not depend on the probabilities for states different to k , (i) follows trivially.

(ii) Using the results from Lemma 1 and denoting \tilde{v}_k as $(K-1)$ column vector of differenced values corresponding to state k ,

$$\begin{aligned} \frac{\partial G_k}{\partial \sigma_k^0} &= [(\rho + \lambda)I - \lambda\Sigma(\sigma) - Q_0]^{-1} \left[\lambda \frac{\partial \Sigma_k(\theta, \sigma)}{\partial \sigma_k^0} V_k + \frac{\partial}{\partial \sigma} \sum_{j=1}^J \sigma_{jk}^0 e_{jk} \right] \\ &= [(\rho + \lambda)I - \lambda\Sigma(\sigma) - Q_0]^{-1} \left[\lambda[-i; I] V_{l(j,k)} - \lambda \tilde{H}_k^{-1}(\sigma^0) \right] \\ &= [(\rho + \lambda)I - \lambda\Sigma(\sigma) - Q_0]^{-1} [\lambda[\tilde{v}_k - \tilde{H}_k^{-1}(\sigma^0)]] \end{aligned}$$

■

We present the proof for Proposition 1 below.

Proof. Rearranging (3) and using the value function operator,

$$\Upsilon(\theta, \sigma) = [(\rho + \lambda)I - \lambda\Sigma(\sigma) - Q_0]^{-1}[\lambda\Sigma(\sigma)\Upsilon(\theta, \sigma) + u + \lambda E(\sigma)]$$

Therefore, $\Upsilon(\theta, \sigma) = G(\sigma, \Upsilon(\theta, \sigma))$. Differentiating on both sides with respect to σ ,

$$\begin{aligned} \frac{\partial \Upsilon(\sigma)}{\partial \sigma} &= \frac{\partial G(\cdot)}{\partial \sigma} + \frac{\partial G(\cdot)}{\partial \Upsilon(\cdot)} \frac{\partial \Upsilon(\sigma)}{\partial \sigma} \\ \left[I - \frac{\partial G(\cdot)}{\partial \Upsilon(\cdot)} \right] \frac{\partial \Upsilon(\sigma)}{\partial \sigma} &= \frac{\partial G(\cdot)}{\partial \sigma} \\ \frac{\partial \Upsilon(\sigma)}{\partial \sigma} &= \left[I - \frac{\partial G(\cdot)}{\partial \Upsilon(\cdot)} \right]^{-1} \frac{\partial G(\cdot)}{\partial \sigma} \end{aligned}$$

By Lemma 2, $\partial G_l / \partial \sigma_k = 0$ for $l \neq k$ and $\partial G_k / \partial \sigma_k = [(\rho + \lambda)I - \lambda \Sigma(\sigma) - Q_0]^{-1} [\lambda [\tilde{v}_k - \tilde{H}_k^{-1}(\sigma)]]$. In the latter case, let $\tilde{\sigma}$ be the fixed point of $\Upsilon(\theta, \sigma)$. Then $\tilde{v}_k = \tilde{H}^{-1}(\Upsilon(\theta, \tilde{\sigma})) = \tilde{H}^{-1}(\tilde{\sigma})$. So at the fixed point $\tilde{v}_k = \tilde{H}_k^{-1}(\sigma)$ which means $\partial G_k / \partial \sigma_k = 0$, and hence, $\partial G(\cdot) / \partial \sigma^0 = 0$. Then, at the fixed point σ^0 , $\partial \Upsilon(\sigma^0) / \partial \sigma^0 = 0$.

Since $\Gamma(v)$ is continuous, and $\Psi(\theta, \sigma) = \Gamma(\Upsilon(\theta, \sigma))$, $\Psi_\sigma = 0$ at the fixed point σ . ■

A.2 Proof of Proposition 2

We provide the proof for Proposition 2 for single agent models. The proof proceeds along similar lines as the proof of Proposition 4 in Aguirregabiria and Mira (2002). Let $g(k_m, \theta, \sigma) = \sum_{n=1}^T \ln P_{k_{m,n-1}, k_{m,n}}(\Delta; \Psi(\theta, \sigma))$ for $k_m = (k_{m,0}, k_{m,1}, \dots, k_{m,T})$ denote the likelihood function for an individual market m so that the sample log likelihood function can be written as $L_M(\theta, \sigma) = M^{-1} \sum_{m=1}^M g(x_m, \theta, \sigma)$.

Consistency. First, we prove the consistency of $\hat{\theta}^l$.

Step 1. If $\hat{\sigma}^{l-1}$ is consistent, then $L_M(\theta, \hat{\sigma}^{l-1})$ converges a.s. and uniformly in θ to a deterministic function $L(\theta, \sigma^*)$.

By Lemma 24.1 of Gourieroux and Monfort (1995), we have that if (i) $L_M(\theta, \sigma)$ converges a.s. and uniformly in (θ, σ) to $L(\theta, \sigma)$; (ii) $L(\theta, \sigma)$ is uniformly continuous in (θ, σ) ; and (iii) $\{\hat{\sigma}^{l-1}\}$ converges a.s. to σ^* ; then $L_M(\theta, \hat{\sigma}^{l-1})$ converges a.s. and uniformly in θ to $L(\theta, \sigma^*)$.

To establish condition (i), we invoke the uniform strong law of large numbers of (Jennrich, 1969, Theorem 2). This requires that (1) The data are i.i.d.; (2) Θ is compact.; (3) $g(k, \theta, \sigma)$ is continuous at each (θ, σ) for each k and a measurable function of k for each (θ, σ) ; (4) $|g(k, \theta, \sigma)| \leq h(k)$ for all k and (θ, σ) where $E[h(k)] < \infty$. Condition (1) follows from our maintained assumption; condition (2) follows from assumption (a); condition (3) follows from assumption (b) and the fact that the matrix exponential preserves continuity; and condition (4) follows because the transition probability matrix for an irreducible continuous-time Markov jump process has elements strictly bounded between 0 and 1.

To establish condition (ii), note as for (3) above that $\Psi(\theta, \sigma)$ is continuous in θ by condition (b) and the matrix exponential operation preserves continuity. Finally, condition (iii) holds by assumption (e).

Step 2. If $\hat{\theta}^{l-1}$ is consistent, then $\hat{\theta}^l \equiv \operatorname{argmax}_{\theta \in \Theta} L_M(\theta, \hat{\sigma}^{l-1})$ converges a.s. to θ^* .

Let $\hat{\sigma}^l = \Psi(\operatorname{argmax}_{\theta \in \Theta} L_M(\theta, \hat{\sigma}^{l-1}), \hat{\sigma}^{l-1})$. By Property 24.2 in (Gourieroux and Monfort, 1995, p.392), if (i) $L_M(\theta, \sigma)$ converges a.s. and uniformly in θ to $L(\theta, \sigma)$ and (ii) $L(\theta, \sigma)$ has a unique maximum in Θ at θ^* , then $\hat{\theta}^l \equiv \operatorname{argmax}_{\theta \in \Theta} L_M(\theta, \hat{\sigma}^{l-1})$ converges a.s. to θ^* . Condition (i) was established above in Step 1. Condition (ii) follows from the identification assumption (d), which implies that θ^* is the only element of Θ that maximizes $L_M(\theta, \hat{\sigma}^{l-1})$,

and the information inequality.

Step 3. For $l \geq 1$, if $\hat{\sigma}^{l-1} \xrightarrow{a.s.} \sigma^*$, then $\hat{\sigma}^l \xrightarrow{a.s.} \sigma^*$.

By definition, $\hat{\sigma}^l = \Psi(\hat{\theta}^l, \hat{\sigma}^{l-1})$. By Step 2, $\hat{\theta}^l \xrightarrow{a.s.} \theta^*$. Since Ψ is continuous in (θ, σ) , by Slutsky's theorem, $\hat{\sigma}^l \xrightarrow{a.s.} \sigma^*$. By assumption (e), $\hat{\sigma}^0 \xrightarrow{a.s.} \sigma^*$, then by induction on step 2 and 3, we have the result.

Asymptotic normality. Now, we prove the asymptotic normality of $\hat{\theta}$. To apply an induction argument, we show that if $(\hat{\theta}^l, \hat{\sigma}^{l-1})$ is \sqrt{M} -consistent and asymptotic normal, then $\sqrt{M}(\hat{\theta}^l - \theta^*) \xrightarrow{d} N(0, \Omega_{\theta\theta}^{-1})$, and $\hat{\sigma}^l \equiv \Psi(\hat{\theta}^l, \hat{\sigma}^{l-1})$ is also \sqrt{M} -consistent and asymptotic normal.

Step 1. $\sqrt{M}(\hat{\theta}^l - \theta^*) \xrightarrow{d} N(0, V^*)$ and V^* only depends on the upper left $r \times r$ submatrix of Ω where r is the dimension of the parameter vector θ .

First, assume that $\hat{\sigma}^{l-1}$ is a consistent estimator of σ^* such that $[\sqrt{M}\nabla_{\theta}L_M(\theta^*, \sigma^*); \sqrt{M}(\hat{\sigma}^{l-1} - \sigma^*)]' \xrightarrow{d} N(0, \Omega)$. Given assumptions (b) and (d) and the definition of $\hat{\theta}^l$, the first order conditions of optimality imply that $\nabla_{\theta}L_M(\hat{\theta}^l, \hat{\sigma}^{l-1}) = 0$. Since $L_M(\theta, \sigma)$ is twice continuously differentiable, we can apply the mean value theorem:

$$0 = \nabla_{\theta}L_M(\theta^*, \sigma^*) + \nabla_{\theta\theta'}L_M(\theta^*, \sigma^*)(\hat{\theta}^l - \theta^*) + \nabla_{\theta\sigma'}L_M(\theta^*, \sigma^*)(\hat{\sigma}^{l-1} - \sigma^*) + o_p(1)$$

Note that $\nabla_{\theta\theta'}L_M(\theta^*, \sigma^*) \xrightarrow{p} -\Omega_{\theta\theta'}$ and $\nabla_{\theta\sigma'}L_M(\theta^*, \sigma^*) \xrightarrow{p} -\Omega_{\theta\sigma'}$, where $\Omega_{\theta\sigma'} \equiv E[\nabla_{\theta}s_m \nabla_{\sigma'}s_m]$ by Theorems 4.2.1 and 4.1.5 of Amemiya (1985). Rearranging the equation above,

$$\sqrt{M}(\hat{\theta}^l - \theta^*) = \Omega_{\theta\theta'}^{-1} [\sqrt{M}\nabla_{\theta}L_M(\theta^*, \sigma^*) - \Omega_{\theta\sigma'}\sqrt{M}(\hat{\sigma}^{l-1} - \sigma^*)] + o_p(1)$$

which leads to

$$\sqrt{M}(\hat{\theta}^l - \theta^*) \xrightarrow{d} N(0, V^*)$$

where

$$V^* = \Omega_{\theta\theta'}^{-1} [I; \Omega_{\theta\sigma'}] \Omega [I; \Omega_{\theta\sigma'}]' \Omega_{\theta\theta'}^{-1}$$

Note that by the zero Jacobian property of Ψ_{σ} in Proposition 1,

$$\begin{aligned} \nabla_{\sigma}L_M(\theta^*, \Psi(\theta^*, \sigma^*)) &= \frac{\partial L_M(\cdot)}{\partial \Psi(\cdot)} \frac{\partial \Psi(\theta^*, \sigma^*)}{\partial \sigma} \\ &= 0 \end{aligned}$$

By the information matrix equality,

$$\nabla_{\theta\sigma'}L(\theta^*, \sigma^*) = E\left(\frac{\partial L_M(\theta^*, \sigma^*)}{\partial \theta} \frac{\partial L_M(\theta^*, \sigma^*)}{\partial \sigma'}\right) = 0$$

and this leads to $\Omega_{\theta\sigma} = 0$. Then,

$$\begin{aligned} V^* &= \Omega_{\theta\theta'}^{-1} \Omega_{r \times r} \Omega_{\theta\theta'}^{-1} \\ &= \Omega_{\theta\theta'}^{-1} \end{aligned}$$

where $\Omega_{r \times r}$ is the upper left $r \times r$ submatrix of Ω and the second equality comes from the fact that

$$\begin{aligned} \Omega_{r \times r} &= \text{Var}(\sqrt{M} \nabla_{\theta} L_M(\theta^*, \sigma^*)) \\ &= M \text{Var}(\nabla_{\theta} L_M(\theta^*, \sigma^*)) \\ &= M \text{Var}\left(\frac{1}{M} \sum_{m=1}^M \nabla_{\theta} s_m\right) \\ &= \frac{1}{M} \sum_{m=1}^M \text{Var}(\nabla_{\theta} s_m) \\ &= \frac{1}{M} \sum_{m=1}^M \text{E}[\nabla_{\theta} s_m \nabla_{\theta'} s_m] \\ &= \Omega_{\theta\theta'}. \end{aligned}$$

Step 2. $[\sqrt{M} \nabla_{\theta} L_M(\theta^*, \sigma^*); \sqrt{M}(\hat{\sigma}^l - \sigma^*)] \xrightarrow{d} \text{N}(0, \Omega^*)$, and the upper left $r \times r$ submatrices of Ω and Ω^* are identical.

Define

$$\omega_M^l \equiv [\sqrt{M} \nabla_{\theta} L_M(\theta^*, \sigma^*); \sqrt{M}(\hat{\sigma}^l - \sigma^*)]'$$

We know that $\hat{\sigma}^{l+1} = \Psi(\hat{\theta}^{l+1}, \hat{\sigma}^l)$ and $\sigma^* = \Psi(\theta^*, \sigma^*)$. We apply mean value expansion to $(\hat{\sigma}^{l+1} - \sigma^*)$:

$$\begin{aligned} \hat{\sigma}^{l+1} - \sigma^* &= \Psi(\hat{\theta}^l, \hat{\sigma}^l) - \Psi(\theta^*, \sigma^*) \\ &= \Psi_{\theta}(\theta^*, \sigma^*)(\hat{\theta}^{l+1} - \theta^*) + \Psi_{\sigma}(\theta^*, \sigma^*)(\hat{\sigma}^l - \sigma^*) + o_p(1) \end{aligned}$$

Then, from the mean value expansion in step 1, we have $\sqrt{M}(\hat{\theta}^l - \theta^*) = \Omega_{\theta\theta'}^{-1} \sqrt{M} \nabla_{\theta} L_M(\theta^*, \sigma^*) + o_p(1)$, so we substitute into the equation above,

$$\sqrt{M}(\hat{\sigma}^{l+1} - \sigma^*) = \Psi_{\theta}(\theta^*, \sigma^*) \Omega_{\theta\theta'}^{-1} \sqrt{M} \nabla_{\theta} L_M(\theta^*, \sigma^*) + \Psi_{\sigma}(\theta^*, \sigma^*) \sqrt{M}(\hat{\sigma}^l - \sigma^*) + o_p(1).$$

We rewrite in matrix form:

$$\begin{bmatrix} \sqrt{M}\nabla_{\theta}L_M(\theta^*, \sigma^*) \\ \sqrt{M}(\hat{\sigma}^{l+1} - \sigma^*) \end{bmatrix} = \begin{bmatrix} I & 0 \\ \Psi_{\theta}(\theta^*, \sigma^*)\Omega_{\theta\theta'}^{-1} & \Psi_{\sigma}(\theta^*, \sigma^*) \end{bmatrix} \begin{bmatrix} \sqrt{M}\nabla_{\theta}L_M(\theta^*, \sigma^*) \\ \sqrt{M}(\hat{\sigma}^l - \sigma^*) \end{bmatrix} + o_p(1). \quad (17)$$

So, we have $\omega_M^{l+1} = A\omega_M^{l+1}$ where A is the first matrix in the RHS of (17). It follows that if ω_M^l is asymptotically normal, ω_M^{l+1} is also asymptotically normal. From Step 1, we also know that V^* depends on upper $r \times r$ submatrix of Ω and The upper-left $r \times r$ matrix of A is the identity matrix from (17). Therefore, $r \times r$ submatrices of Ω and Ω^* are equal. Since $\sqrt{M}(\hat{\theta}^l - \theta^*) = \Omega_{\theta\theta'}^{-1}\sqrt{M}\nabla_{\theta}L_M(\theta^*, \sigma^*) + o_p(1)$, and denoting $\sqrt{M}(\hat{\theta}^l - \theta^*) \xrightarrow{d} N(0, V^{**})$, we have

$$\begin{aligned} V^{**} &= \Omega_{\theta\theta'}^{-1}\Omega_{r \times r}^*\Omega_{\theta\theta'}^{-1} \\ &= \Omega_{\theta\theta'}^{-1}\Omega_{r \times r}\Omega_{\theta\theta'}^{-1} \\ &= \Omega_{\theta\theta'}^{-1}. \end{aligned}$$

From step 1, we showed that when $\hat{\sigma}^{l-1}$ is a consistent estimator of σ^* s.t. $\omega^l \xrightarrow{d} N(0, \Omega)$, it follows that $\sqrt{M}(\hat{\theta}^l - \theta^*) \xrightarrow{d} N(0, \Omega_{\theta\theta'}^{-1})$. Step 2 proves that $\omega^{l+1} \xrightarrow{d} N(0, \Omega^*)$ which implies $\sqrt{M}(\hat{\theta}^{l+1} - \theta^*) \xrightarrow{d} N(0, \Omega_{\theta\theta'}^{-1})$. Since we start from asymptotic normal estimator $\hat{\sigma}^0$ such that $\omega^0 \xrightarrow{d} N(0, \Omega)$, this is true for all $l \leq L - 1$ by induction.

A.3 Proof of Proposition 3

We now extend our results for single agent models to the case of games and prove the large sample properties of the CTNPL estimator stated in Proposition 3.

Proof. Consistency. Before proving the main result, we first establish that for an arbitrary consistent estimator $\tilde{\sigma}$ for σ^* , the estimator $\tilde{\theta} = \arg \max_{\theta \in \Theta} L_M(\theta, \tilde{\sigma})$ is strongly consistent for θ^* . The argument is the same as in Step 1 of the proof of Proposition 2. First, by the strong uniform law of large numbers of Jennrich (1969) we have a.s. convergence of L_M to L uniformly in (θ, σ) . Second, we use Lemma 24.1 of Gourieroux and Monfort (1995) to establish a.s. convergence of $L_M(\theta, \tilde{\sigma})$ to $L(\theta, \sigma^*)$ uniformly in θ . Finally, strong consistency of the extremum estimator $\tilde{\theta}$ follows from Property 24.2 of (Gourieroux and Monfort, 1995, p.392).

Now, we prove the main result $(\hat{\theta}^l, \hat{\sigma}^{l-1}) \xrightarrow{as} (\theta^*, \sigma^*)$ by induction. For $l = 1$, $\hat{\sigma}^0 \xrightarrow{as} \sigma^*$ follows directly from assumption (e), that we begin with a strongly consistent estimator $\hat{\sigma}^0$ for σ^* . To show strong consistency of $\hat{\theta}^1$ starting from $\hat{\sigma}^0$, we take $\hat{\sigma}^0$ to be the consistent estimator $\tilde{\sigma}$ above which yields strong consistency of $\hat{\theta}^1$ for θ^* at iteration $l = 1$.

Now assume that $(\hat{\theta}^l, \hat{\sigma}^{l-1}) \xrightarrow{\text{as}} (\theta^*, \sigma^*)$ for all $l \leq L - 1$. Then, $\hat{\sigma}^l = \Psi(\hat{\theta}^l, \hat{\sigma}^{l-1}) \xrightarrow{\text{as}} \Psi(\theta^*, \sigma^*) = \sigma^*$. Using $\hat{\sigma}^l$ as our consistent estimator $\tilde{\sigma}$, as above, yields $\hat{\theta}^{l+1} \xrightarrow{\text{as}} \theta^*$. Since we start from a strongly consistent estimator $\hat{\sigma}^0$, we have strong consistency of $\hat{\theta}^l$ for all $l \leq L$ by induction.

Asymptotic normality. As with consistency, we prove that $\hat{\sigma}^{l-1} \xrightarrow{d} N(0, \Sigma^{l-1})$ and $\hat{\theta}^l \xrightarrow{d} N(0, \Omega_{\theta\theta'} + \Omega_{\theta\theta'}\Omega_{\theta\sigma}\Sigma^{l-1}\Omega_{\theta\sigma'}\Omega_{\theta\theta'}^{-1})$ by induction. We first show that the result holds for $l = 1$. As in assumption (e), let $\hat{\sigma}^0 = \frac{1}{M} \sum_{m=1}^M r_m$ denote the initial nonparametric estimator for σ^* . Let $\nabla_{\theta} s_m^l = \nabla_{\theta} \ln P_{k_{m,n-1}, k_{m,n}}(\Delta; \Psi(\theta, \hat{\sigma}^{l-1}))$ denote the pseudo score at iteration l based on estimated CCPs from iteration $l-1$. We expand the first order condition $\nabla_{\theta} L_M(\hat{\theta}^1, \hat{\sigma}^0) = 0$ around (θ^*, σ^*) :

$$0 = \nabla_{\theta} L_M(\theta^*, \sigma^*) + \nabla_{\theta\theta'} L_M(\theta^*, \sigma^*)(\hat{\theta}^1 - \theta^*) + \nabla_{\theta\sigma'} L_M(\theta^*, \sigma^*)(\hat{\sigma}^0 - \sigma^*) + o_p(1)$$

By the generalized information matrix equality, $\nabla_{\theta\theta'} L_M(\theta^*, \sigma^*) \xrightarrow{p} -\Omega_{\theta\theta'}$ and $\nabla_{\theta\sigma'} L_M(\theta^*, \sigma^*) \xrightarrow{p} -\Omega_{\theta\sigma'}$. Then,

$$\sqrt{M}(\hat{\theta}^1 - \theta^*) = \Omega_{\theta\theta'}^{-1} \left\{ -\Omega_{\theta\sigma'} \left(\frac{1}{\sqrt{M}} \sum_{m=1}^M \nabla_{\sigma} r_m \right) + \left(\frac{1}{\sqrt{M}} \sum_{m=1}^M \nabla_{\theta} s_m^0 \right) \right\} + o_p(1). \quad (18)$$

From condition (b), $\Psi(\theta, \sigma)$ is continuously differentiable and all elements of $\exp(Q(\theta, \sigma))$ are strictly bounded between 0 and 1 since the Markov chain is irreducible. Then, by the generalized information matrix equality, $E[\nabla_{\sigma} r_m \nabla_{\theta'} s_m^0] = 0$. Then,

$$\frac{1}{\sqrt{M}} \left(\sum_{m=1}^M \nabla_{\theta} s_m^0 \right) - \Omega_{\theta\sigma'} \left(\frac{1}{\sqrt{M}} \sum_{m=1}^M \nabla_{\sigma} r_m \right) \xrightarrow{d} N(0, \Omega_{\theta\theta'} + \Omega_{\theta\sigma'} \Sigma^0 \Omega_{\theta\sigma'}').$$

Substituting into equation (18), we have asymptotic normality for $\hat{\theta}^1$.

$$\sqrt{M}(\hat{\theta}^1 - \theta^*) \xrightarrow{d} N(0, \Omega_{\theta\theta'}^{-1} + \Omega_{\theta\theta'}^{-1} \Omega_{\theta\sigma'} \Sigma^0 \Omega_{\theta\sigma'}' \Omega_{\theta\theta'}^{-1}).$$

Similarly, we can expand the first order condition $\nabla_{\sigma} L_M(\hat{\theta}^1, \hat{\sigma}^0) = 0$ around (θ^*, σ^*) :

$$0 = \nabla_{\theta} L_M(\theta^*, \sigma^*) + \nabla_{\theta\theta'} L_M(\theta^*, \sigma^*)(\hat{\theta}^1 - \theta^*) + \nabla_{\theta\sigma'} L_M(\theta^*, \sigma^*)(\hat{\sigma}^0 - \sigma^*) + o_p(1)$$

Now we assume that $\sqrt{M}(\hat{\sigma}^{l-1} - \sigma^*) \xrightarrow{d} N(0, \Sigma^{l-1})$ and $\sqrt{M}(\hat{\theta}^l - \theta^*) \xrightarrow{d} N(0, \Omega_{\theta\theta'}^{-1} + \Omega_{\theta\theta'}^{-1} \Omega_{\theta\sigma'} \Sigma^{l-1} \Omega_{\theta\sigma'}' \Omega_{\theta\theta'}^{-1})$ for some $l \leq L - 1$. We can consider $\hat{\theta}^{l+1}$ as a two step estimator based on the moment equation $E[\nabla_{\theta} L_M(\theta, \hat{\sigma}^l)] = 0$. We use the preliminary estimate $\hat{\sigma}^l$ of σ^* from the previous iteration based on the moment equation, $E[\Psi(\hat{\theta}^l, \hat{\sigma}^{l-1}) - \sigma] = 0$. We

can write the moment conditions as $\tilde{g}(\theta, \sigma)$.

$$\tilde{g}(\theta, \sigma) = \begin{bmatrix} g(\theta, \sigma) \\ h(\sigma) \end{bmatrix} = \begin{bmatrix} \nabla_{\theta} L_M(\theta, \sigma) \\ \Psi(\hat{\theta}^l, \hat{\sigma}^{l-1}) - \sigma \end{bmatrix}.$$

Applying the generalized information matrix inequality,

$$\mathbb{E}(\nabla_{\theta} \tilde{g}(\theta, \sigma)) = -\mathbb{E}[\tilde{g}(\theta, \sigma) \nabla_{\theta} s_m^{l+1}]$$

This results in $\mathbb{E}[\nabla_{\theta} h(\sigma)] = -\mathbb{E}[\nabla_{\theta}(\Psi_m(\hat{\theta}^l, \sigma) - \sigma^*) \nabla_{\theta} s_m^{l+1}] = 0$. As before, we expand the first moment condition $\nabla_{\theta} L_M(\hat{\theta}^{l+1}, \hat{\sigma}^l) = 0$ using the intermediate value theorem.

$$0 = \nabla_{\theta} L_M(\theta^*, \sigma^*) + \nabla_{\theta\theta'} L_M(\theta^*, \sigma^*)(\hat{\theta}^{l+1} - \theta^*) + \nabla_{\theta\sigma'} L_M(\theta^*, \sigma^*)(\hat{\sigma}^l - \sigma^*) + o_p(1). \quad (19)$$

Then, $\sqrt{M}(\hat{\theta}^{l+1} - \theta^*)$ can be written as

$$\sqrt{M}(\hat{\theta}^{l+1} - \theta^*) = \Omega_{\theta\theta'}^{-1} \left\{ -\Omega_{\theta\sigma'} \sqrt{M}(\hat{\sigma}^l - \sigma^*) + \left(\frac{1}{\sqrt{M}} \sum_{m=1}^M \nabla_{\theta} s_m^{l+1} \right) \right\} + o_p(1) \quad (20)$$

where $\sqrt{M}(\hat{\sigma}^l - \sigma^*) \xrightarrow{d} N(0, \Sigma^l)$. Then, from equation (20),

$$\sqrt{M}(\hat{\theta}^{l+1} - \theta^*) \xrightarrow{d} N(0, \Omega_{\theta\theta'}^{-1} + \Omega_{\theta\theta'}^{-1} \Omega_{\theta\sigma'} \Sigma^l \Omega_{\theta\sigma'}' \Omega_{\theta\theta'}^{-1}).$$

The recursive expression for Σ^l stated in the proposition follows by expanding $\hat{\sigma}^l - \Psi(\hat{\theta}^l, \hat{\sigma}^{l-1})$ around (θ^*, σ^*) to obtain

$$\hat{\sigma}^l - \sigma^* = \nabla_{\theta} \Psi^*(\hat{\theta}^l - \theta^*) + \nabla_{\sigma} \Psi^*(\hat{\sigma}^{l-1} - \sigma^*) + o_p(1),$$

and using the asymptotic normality of both $\hat{\theta}^l$ and $\hat{\sigma}^l$ from the inductive hypothesis. ■

A.4 Proof of Proposition 4

We omit the detailed proof of Proposition 4 as it is essentially the same as the proof in the Appendix of Aguirregabiria and Mira (2007), since the CTNPL operator has the same required properties as the NPL operator. Changing $Q_0(\theta, P)$ to $L(\theta, \sigma)$ and P to σ will be sufficient to prove an analogous proposition for the CTNPL estimator.

A.5 Proof of Proposition 5

We first prove an auxiliary proposition needed for convergence condition. This is analogous to Proposition 7 in Kasahara and Shimotsu (2012).

Proposition 6. *Suppose that Assumption 8 holds. Then, there exists a neighborhood \mathcal{N}_1 of σ^* such that $\hat{\theta}^l - \hat{\theta} = O(\|\hat{\theta}^{l-1} - \hat{\theta}\|)$ a.s. and $\hat{\sigma}^l - \hat{\sigma} = M_{\Psi_\theta} \Psi_\sigma^* (\hat{\sigma}^{l-1} - \hat{\sigma}) + O(M^{-1/2} \|\hat{\sigma}^{l-1} - \hat{\sigma}\| + \|\hat{\sigma}^{l-1} - \hat{\sigma}\|^2)$ a.s. uniformly in $\hat{\sigma}^{l-1} \in \mathcal{N}_1$.*

Proof. For $\varepsilon > 0$, define a neighborhood $\mathcal{N}(\varepsilon) = \{(\theta, \sigma) : \|\theta - \theta^*\| + \|\sigma - \sigma^*\| < \varepsilon\}$. Then, there exists $\varepsilon_1 > 0$ such that $\mathcal{N}(\varepsilon_1) \subset \mathcal{N}$ and $\sup_{(\theta, \sigma) \in \mathcal{N}(\varepsilon_1)} \|\nabla_{\theta\theta'} L(\theta, \sigma)^{-1}\| < \infty$ since $\nabla_{\theta\theta'} L(\theta, \sigma)$ is continuous and $\nabla_{\theta\theta'} L(\theta^*, \sigma^*)$ is nonsingular.

We assume that $(\hat{\theta}^l, \hat{\sigma}^{l-1}) \in \mathcal{N}(\varepsilon_1)$ which we will show in the end of this proof.

Step 1. We first prove the first statement: $\hat{\theta}^l - \hat{\theta} = O(\|\hat{\theta}^{l-1} - \hat{\theta}\|)$ a.s.

From the first step in NPL iteration, we know that $\nabla_\theta L_M(\hat{\theta}^l, \hat{\sigma}^{l-1}) = 0$. Applying the mean value theorem around $(\hat{\theta}, \hat{\sigma})$, we have

$$0 = \nabla_{\theta\theta'} L_M(\bar{\theta}, \bar{\sigma})(\hat{\theta}^l - \hat{\theta}) + \nabla_{\theta\sigma'} L_M(\bar{\theta}, \bar{\sigma})(\hat{\sigma}^{l-1} - \hat{\sigma}) + o_p(1). \quad (21)$$

where $(\bar{\theta}, \bar{\sigma})$ lie between $(\hat{\theta}^l, \hat{\sigma}^{l-1})$ and $(\hat{\theta}, \hat{\sigma})$.

Then, we can rewrite

$$\hat{\theta}^l - \hat{\theta} = -\nabla_{\theta\theta'} L_M(\bar{\theta}, \bar{\sigma})^{-1} \nabla_{\theta\sigma'} L_M(\bar{\theta}, \bar{\sigma})(\hat{\sigma}^{l-1} - \hat{\sigma}) + o_p(1). \quad (22)$$

Note that (i) $(\bar{\theta}, \bar{\sigma}) \in \mathcal{N}(\varepsilon_1)$ from assumption and $(\bar{\theta}, \bar{\sigma}) \in \mathcal{N}$ since $\mathcal{N}(\varepsilon_1) \subset \mathcal{N}$, and (ii) $\sup_{(\theta, \sigma) \in \mathcal{N}(\varepsilon_1)} \|\nabla_{\theta\theta'} L_M(\theta, \sigma)^{-1} \nabla_{\theta\sigma'} L_M(\theta, \sigma)\| = O(1)$ a.s. since $\sup_{(\theta, \sigma) \in \mathcal{N}(\varepsilon_1)} \|\nabla_{\theta\theta'} L(\theta, \sigma)^{-1}\| < \infty$ and $\sup_{(\theta, \sigma) \in \mathcal{N}} \|\nabla^2 L_M(\theta, \sigma) - \nabla^2 L(\theta, \sigma)\| = o(1)$ by general uniform convergence.⁷

Then, we have the first result: $\hat{\theta}^l - \hat{\theta} = O(\|\hat{\sigma}^{l-1} - \hat{\sigma}\|)$ a.s.

Step 2. We now prove $\hat{\sigma}^l - \hat{\sigma} = M_{\Psi_\theta} \Psi_\sigma^* (\hat{\sigma}^{l-1} - \hat{\sigma}) + O(M^{-1/2} \|\hat{\sigma}^{l-1} - \hat{\sigma}\| - \|\hat{\sigma}^{l-1} - \hat{\sigma}\|^2)$ a.s. uniformly in $\hat{\sigma}^{l-1} \in \mathcal{N}_1$.

We use the Taylor expansion, root-M consistency of $(\hat{\theta}, \hat{\sigma})$, and the information matrix equality for the results below:

⁷(i) Θ is compact implies total boundedness. (ii) $\nabla^2 L_M(\theta, \sigma) \rightarrow \nabla^2 L(\theta, \sigma)$ a.s. $\forall \theta \in \Theta$. (iii) $\nabla^2 L(\theta, \sigma)$ is a nonrandom function that is uniformly continuous in $\theta \in \Theta$, and $|\nabla^2 L_M(\theta, \sigma) - \nabla^2 L_M(\theta', \sigma')| \leq |\nabla^2 L_M(\theta, \sigma) - \nabla^2 L_M(\theta^*, \sigma^*)| + |\nabla^2 L_M(\theta', \sigma') - \nabla^2 L_M(\theta^*, \sigma^*)| \leq 2\delta \leq \delta' d((\theta, \sigma), (\theta', \sigma'))$ for some small δ and δ' , $\forall \theta', \theta \in \mathcal{N}$ a.s. since $d(\cdot) < \varepsilon$ and $L_M(\theta, \sigma)$ is continuous. Then, by Theorem 2 and Lemma 1 in Andrews (1992), the result follows.

$$\begin{aligned}
\nabla_{\theta\theta'} L_M(\hat{\theta}, \hat{\sigma}) &= -\Omega_{\theta\theta'} + O(M^{-1/2}) \\
\nabla_{\theta\sigma'} L_M(\hat{\theta}, \hat{\sigma}) &= -\Omega_{\theta\sigma'} + O(M^{-1/2}) \\
\nabla_{\theta} \Psi(\hat{\theta}, \hat{\sigma}) &= \Psi_{\theta}^* + O(M^{-1/2}) \\
\nabla_{\sigma} \Psi(\hat{\theta}, \hat{\sigma}) &= \Psi_{\sigma}^* + O(M^{-1/2})
\end{aligned} \tag{23}$$

We use the second step of NPL estimation $\hat{\sigma}^l = \Psi(\hat{\theta}^l, \hat{\sigma}^{l-1})$ and expand twice around $(\hat{\theta}, \hat{\sigma})$.

$$\begin{aligned}
\hat{\sigma}^l - \hat{\sigma} &= \nabla_{\theta} \Psi(\hat{\theta}, \hat{\sigma})(\hat{\theta}^l - \hat{\theta}) + \nabla_{\sigma} \Psi(\hat{\theta}, \hat{\sigma})(\hat{\sigma}^{l-1} - \hat{\sigma}) + O(\|\hat{\sigma}^{l-1} - \hat{\sigma}\|^2) \\
&= \Psi_{\theta}^*(\hat{\theta}^l - \hat{\theta}) + \Psi_{\sigma}^*(\hat{\sigma}^{l-1} - \hat{\sigma}) + O(\|\hat{\sigma}^{l-1} - \hat{\sigma}\|^2) + O(M^{-1/2}\|\hat{\sigma}^{l-1} - \hat{\sigma}\|).
\end{aligned} \tag{24}$$

where the first equality follows from Step 1 and $\sup_{(\theta, \sigma) \in \mathcal{N}(\varepsilon_1)} \nabla^3 \Psi(\theta, \sigma) < \infty$ in assumption (b) and the second equality comes from (23).

We expand $\nabla_{\theta\theta'} L_M(\bar{\theta}, \bar{\sigma})$ around $(\hat{\theta}, \hat{\sigma})$. Since $\|\bar{\theta} - \hat{\theta}\| \leq \|\hat{\theta}^l - \hat{\theta}\|$, $\|\bar{\sigma} - \hat{\sigma}\| \leq \|\hat{\sigma}^{l-1} - \hat{\sigma}\|$ and $\hat{\theta}^{l-1} - \hat{\theta} = O(\|\hat{\sigma}^{l-1} - \hat{\sigma}\|)$ from Step 1, we have $\nabla_{\theta\theta'} L_M(\bar{\theta}, \bar{\sigma}) = \nabla_{\theta\theta'} L_M(\hat{\theta}, \hat{\sigma}) + O(\|\hat{\sigma}^{l-1} - \hat{\sigma}\|)$ a.s. Using (23) and repeating the similar process for $\nabla_{\theta\sigma'} L_M(\bar{\theta}, \bar{\sigma})$,

$$\begin{aligned}
\nabla_{\theta\theta'} L_M(\bar{\theta}, \bar{\sigma}) &= -\Omega_{\theta\theta'} + O(M^{-1/2}) + O(\|\hat{\sigma}^{l-1} - \hat{\sigma}\|) \\
\nabla_{\theta\sigma'} L_M(\bar{\theta}, \bar{\sigma}) &= -\Omega_{\theta\sigma'} + O(M^{-1/2}) + O(\|\hat{\sigma}^{l-1} - \hat{\sigma}\|)
\end{aligned}$$

Then, applying (22),

$$\hat{\theta}^l - \hat{\theta} = -\Omega_{\theta\theta'}^{-1} \Omega_{\theta\sigma'} (\hat{\sigma}^{l-1} - \hat{\sigma}) + O(M^{-1/2} \|\hat{\sigma}^{l-1} - \hat{\sigma}\| + \|\hat{\sigma}^{l-1} - \hat{\sigma}\|^2).$$

Substituting the equation above into (24), we have the result.

Step 3. Now we need to show that $(\hat{\theta}^l, \hat{\sigma}^{l-1}) \in \mathcal{N}(\varepsilon_1)$. We first prove that $\|\hat{\theta}^l - \hat{\theta}^0\| < \varepsilon/2$ and then, show that $\|\hat{\sigma}^{l-1} - \hat{\sigma}\| < \varepsilon/2$ if we set \mathcal{N}_1 sufficiently small.

Let $\mathcal{N}_{\theta} \equiv \{\theta : \|\theta - \theta^*\| < \varepsilon_1/2\}$ and define $\Delta = L(\theta^*, \sigma^*) - \sup_{\theta \in \mathcal{N}_{\theta}^c \cap \Theta} L(\theta, \sigma^*) > 0$. The inequality follows from information inequality, compactness of $\mathcal{N}_{\theta}^c \cap \Theta$, and continuity

of $L(\theta, \sigma)$. Then, if $\hat{\theta}^l \notin \mathcal{N}_\theta$, then $L(\theta^*, \sigma^*) - L(\hat{\theta}^l, \sigma^*) \geq \Delta$. We can also derive that

$$\begin{aligned}
& L(\theta^*, \sigma^*) - L(\hat{\theta}^l, \sigma^*) \\
& \leq L_M(\theta^*, \sigma^*) - L_M(\hat{\theta}^l, \sigma^*) + 2 \sup_{(\theta, \sigma) \in \Theta \times \Xi} |L_M(\theta, \sigma) - L(\theta, \sigma)| \\
& \leq L_M(\theta^*, \hat{\sigma}^{l-1}) - L_M(\hat{\theta}^l, \hat{\sigma}^{l-1}) + 2 \sup_{(\theta, \sigma) \in \Theta \times \Xi} |L_M(\theta, \sigma) - L(\theta, \sigma)| + 2 \sup_{\theta \in \Theta} |L(\theta, \sigma^*) - L(\theta, \hat{\sigma}^{l-1})| \\
& \leq 2 \sup_{(\theta, \sigma) \in \Theta \times \Xi} |L_M(\theta, \sigma) - L(\theta, \sigma)| + 2 \sup_{\theta \in \Theta} |L(\theta, \sigma^*) - L(\theta, \hat{\sigma}^{l-1})|.
\end{aligned}$$

The last equality follows from $\hat{\theta}^l = \operatorname{argmax}_{\theta \in \Theta} L_M(\theta, \hat{\sigma}^{l-1})$. We know that $2 \sup_{(\theta, \sigma) \in \Theta \times \Xi} |L_M(\theta, \sigma) - L(\theta, \sigma)| = o(1)$ a.s. from Step 1 of Proof for consistency in Proposition 4. Since $L(\theta, \sigma)$ is continuous, there exists ε_Δ such that $2 \sup_{\theta \in \Theta} |L(\theta, \sigma^*) - L(\theta, \hat{\sigma}^{l-1})| < \Delta/2$ if $\|\sigma^* - \hat{\sigma}^{l-1}\| \leq \varepsilon_\Delta$. It follows that if $\|\sigma^* - \hat{\sigma}^{l-1}\| \geq \varepsilon_\Delta$, $L(\theta^*, \sigma^*) - L(\hat{\theta}^l, \sigma^*) \leq \Delta$ which means $\hat{\theta}^l \in \mathcal{N}_\theta$.

Then, if we set $\mathcal{N}_1 = \{\sigma : \|\sigma - \sigma^*\| \leq \min\{\varepsilon_1/2, \varepsilon_\Delta\}\}$, we have $\|\sigma - \sigma^*\| < \varepsilon/2$ which gives $(\hat{\theta}^l, \hat{\sigma}^{l-1}) \in \mathcal{N}(\varepsilon_1)$ a.s. \blacksquare

We now present a proof for Proposition 5. This is analogous to Proposition 2 in Kasahara and Shimotsu (2012).

Proof. Let $b > 0$ be a constant such that

$$\rho(M_{\Psi_\theta} \Psi_\sigma) + 2b < 1.$$

From Lemma 5.6.10 of Horn (1995), there is a matrix norm $\|\cdot\|_\alpha$ such that

$$\|M_{\Psi_\theta} \Psi_\sigma\|_\alpha \leq \rho(M_{\Psi_\theta} \Psi_\sigma) + b.$$

Define a vector norm $\|\cdot\|_\beta$ for $x \in \mathbb{R}^{NJK}$ as $\|x\|_\beta \equiv \|[x \ 0 \ \dots \ 0]\|_\alpha$, then a direct calculation gives

$$\|Ax\|_\beta = \|A[x \ 0 \ \dots \ 0]\|_\alpha \leq \|A\|_\alpha \|x\|_\beta$$

for any matrix A .

From the equivalence of vector norms in \mathbb{R}^{NJK} , we can restate Proposition 6 in terms of β : there exists $c > 0$ such that

$$\hat{\sigma}^l - \hat{\sigma} = M_{\Psi_\theta} \Psi_\sigma (\hat{\sigma}^{l-1} - \hat{\sigma}) + O(M^{-1/2} \|\hat{\sigma}^{l-1} - \hat{\sigma}\|_\beta + \|\hat{\sigma}^{l-1} - \hat{\sigma}\|_\beta^2)$$

a.s. holds uniformly in $\hat{\sigma}^{l-1} \in \{\sigma : \|\sigma - \sigma^*\|_\beta < c\}$.

We rewrite this statement further so that it is amenable to recursive substitution.

- (i) $\|M_{\Psi_\theta} \Psi_\sigma (\hat{\sigma}^{l-1} - \hat{\sigma})\|_\beta \leq \|M_{\Psi_\theta} \Psi_\sigma\|_\alpha \|\hat{\sigma}^{l-1} - \hat{\sigma}\|_\beta \leq (\rho(M_{\Psi_\theta} \Psi_\sigma) + b) \|\hat{\sigma}^{l-1} - \hat{\sigma}\|_\beta$.
- (ii) The remainder term can be written as $O(M^{-1/2} + \|\hat{\sigma}^{l-1} - \hat{\sigma}\|_\beta) \|\hat{\sigma}^{l-1} - \hat{\sigma}\|_\beta$. Then setting

$c < b$ and using consistency of $\hat{\sigma}$, this term is smaller than $b\|\hat{\sigma}^{l-1} - \hat{\sigma}\|_\beta$ a.s.

(iii) Since $\hat{\sigma}$ is consistent, $\{\sigma : \|\hat{\sigma} - \sigma\|_\beta < c/2\} \subset \{\sigma : \|\sigma - \sigma^*\|_\beta < c\}$ a.s.

From (i) to (iii),

$$\|\hat{\sigma}^{l-1} - \hat{\sigma}\|_\beta \leq (\rho(M_{\Psi_\theta} \Psi_\sigma) + 2b)\|\hat{\sigma}^{l-1} - \hat{\sigma}\|_\beta$$

holds a.s. for all $\hat{\sigma}^{l-1} \in \{\sigma : \|\hat{\sigma}^{l-1} - \hat{\sigma}\|_\beta < c/2\}$. Because each NPL updating of (θ, σ) uses the same pseudo-likelihood function, we may recursively substitute for the $\hat{\sigma}_j$'s.

$$\begin{aligned} \|\hat{\sigma}^{l-1} - \hat{\sigma}\|_\beta &\leq (\rho(M_{\Psi_\theta} \Psi_\sigma) + 2b)\|\hat{\sigma}^{l-1} - \hat{\sigma}\|_\beta \\ &\leq (\rho(M_{\Psi_\theta} \Psi_\sigma) + 2b)^2\|\hat{\sigma}^{l-2} - \hat{\sigma}\|_\beta \\ &\vdots \\ &\leq (\rho(M_{\Psi_\theta} \Psi_\sigma) + 2b)^l\|\hat{\sigma}^0 - \hat{\sigma}\|_\beta \end{aligned}$$

Then, $\lim_{l \rightarrow \infty} \hat{\sigma}^l = \hat{\sigma}$ a.s. if $\|\hat{\sigma}^{l-1} - \hat{\sigma}\|_\beta < c/2$. Applying the equivalence of vector norms in \mathcal{R}^L to $\|\hat{\sigma}^0 - \hat{\sigma}\|_\beta$ and $\|\hat{\sigma}^0 - \hat{\sigma}\|$ and consistency of $\hat{\sigma}$, the result follows. \blacksquare

B Implementation Details

In this section, we provide more details on estimation process. Our goal is to estimate $\theta = (\theta_{\text{FC},1}, \dots, \theta_{\text{FC},5}, \theta_{\text{RS}}, \theta_{\text{RN}}, \theta_{\text{EC}})$ using the NPL algorithm. We first review two steps needed for estimating:

1. Given $\hat{\sigma}^{l-1}$, update $\hat{\theta}$ by

$$\hat{\theta}^l = \operatorname{argmax}_{\theta \in \Theta} L_M(\theta, \hat{\sigma}^{l-1}) = \operatorname{argmax}_{\theta \in \Theta} \frac{1}{M} \sum_{m=1}^M \sum_{n=1}^T \ln P_{k_m, n-1, k_{mn}}(\Delta; \Psi(\theta, \hat{\sigma}^{l-1})). \quad (12)$$

2. Update $\hat{\sigma}$ using the equilibrium condition, i.e.

$$\hat{\sigma}^l = \Psi(\hat{\theta}^l, \hat{\sigma}^{l-1}) \quad (13)$$

where

$$\hat{\sigma}_{jk} = \frac{\exp(\psi_{jk} + V_{l(j,k)})}{\sum_{j'} \exp(\psi_{j'k} + V_{l(j',k)})}$$

We need to express the likelihood function in terms of θ to maximize it with regard to θ .

First, we start from calculating the value function in (5) using value function below:

$$V_i(\theta, \sigma) = \left[(\rho + N\lambda)I - \lambda \sum_{m=1}^N \Sigma_m(\sigma_m) - Q_0 \right]^{-1} [u_i(\theta) + \lambda_i E_i(\theta, \sigma)] \quad (3)$$

where $\Sigma_m(\sigma_m)$ is the $K \times K$ state transition matrix induced by the actions of player m given the choice probabilities σ_m and where $E_i(\theta, \sigma)$ is a $K \times 1$ vector where each element k is the ex-ante expected value of the choice-specific payoff in state k , $\sum_j \sigma_{ijk} [\psi_{ijk} + e_{ijk}(\theta, \sigma)]$.

We can define the first parenthesis in (3) as Ξ . We assume $\rho_i = 1$, $\lambda_{ik} = 1$ for all $i = 1, \dots, N$ and $k = 1, \dots, K$, and that Q_0 is known. We can also calculate $\Sigma_m(\sigma_m)$ using $\hat{\sigma}_m^{l-1}$. We start from true probabilities σ^* for 2S-True estimator and random draws for $\hat{\sigma}^0$ for NPL-random estimator.

Then, we rewrite the second parenthesis in (3), in terms of θ . The flow payoff u_{ik} can be rewritten as:

$$\begin{aligned} u_{ik} &= \theta_{RS} \ln(s_k) - \theta_{RN} \ln\left(1 + \sum_{m \neq i} a_{mk}\right) - \theta_{FC,i} \\ &= z_{ik}^u \theta^u \end{aligned}$$

where $z_{ik}^u = [I_{ik} \ln(s_k) \ln(1 + \sum_{m \neq i} a_{mk})]$ where I_{ik} is a 1×5 vector with 1 in i -th position and zero elsewhere, and $\theta^u = (\theta_{FC,1}, \dots, \theta_{FC,5}, \theta_{RS}, \theta_{RN})'$.

Now we express $E_{ik}(\theta, \sigma)$, corresponding element of $E(\theta, \sigma)$ in term of θ . The determinant part of instantaneous payoff has the structure:

$$\psi_{ijk} = z_{ijk}^\psi \theta_{EC}$$

where

$$z_{ijk}^\psi = \begin{cases} -1 & \text{if } j = 1 \\ 0 & \text{otherwise} \end{cases}$$

Expected value of deterministic instantaneous payoff z_{ijk}^ψ differs by the choice, so we can write

$$\begin{aligned} \mathbb{E} z_{ik}^\psi &= \sum_{j \in \mathcal{A}} \sigma_{ijk} z_{ijk}^\psi \\ &= \lambda [\sigma_{0k} \times 0 + \sigma_{1k} \times (-1)] \end{aligned}$$

The stochastic part ε_{ijk} has the expected value e_{ijk} :

$$\begin{aligned} e_{ijk} &= \mathbb{E}\varepsilon_{ijk} \\ &= \lambda[\sigma_{0k}(\gamma - \log(\sigma_{0k})) + \sigma_{1k}(\gamma - \log(\sigma_{1k})) + \sigma_{-1,k}(\gamma - \log(\sigma_{-1,k}))] \end{aligned}$$

using T1EV distribution of ε_{ijk} .

Collecting all terms for the value function, we can rewrite the value function as

$$\begin{aligned} V_{ij}(\theta, \sigma) &= \Xi^{-1}[u(\theta) + \lambda E(\theta, \sigma)] \\ &= \Xi^{-1}[z_{ij}^u \theta^u + \lambda \mathbb{E}(z_{ij}^\psi \theta_{\text{EC}} + \varepsilon_{ij})] \\ &= \Xi^{-1} \mathbb{E}[z_{ij} \theta + \varepsilon_{ij}] \\ &= W_{ij} \theta' + \tilde{e}_{ij}. \end{aligned}$$

where z_{ij} is a $K \times |\theta|$ vector with each k th element (z_{ik}^u, z_{ijk}^ψ) , W_{ijk} is a $K \times 1$ vector which equals $\Xi^{-1} \mathbb{E} z_{ij}$, and \tilde{e}_{ij} is $\Xi^{-1} e_{ij}$. Now we can write conditional choice probabilities in terms of θ ,

$$\hat{\sigma}_{ijk} = \frac{\exp(W_{ijk} \theta' + \tilde{e}_{ijk})}{\sum_{j'} \exp(W_{ij'k} \theta' + \tilde{e}_{ij'k})} \quad (25)$$

Since we can calculate W_{ijk} and \tilde{e}_{ijk} from known parameters or conditional choice probabilities, the only unknown term in $\hat{\sigma}$, hence the likelihood function, is θ . Then, we can form Q in terms of θ using Q_0 and $\hat{\sigma}$, and maximize the likelihood with regard to θ . We use starting values of $\theta^0 = (1, 1, 1, 1, 1, 1, 1, 1)$. We use the solver `fminunc` in MATLAB.

Once $\hat{\theta}^l$ is estimated, the conditional choice probability $\hat{\sigma}^l$ is updated. The entire process is repeated 20 times ($L=20$). For Monte Carlo experiments, we generate 100 different data set using same parameters and estimate separately for each data set. Then, we report the mean of estimates as estimates for each estimator and the standard deviation from 100 estimates as standard error.