# Efficient and Convergent Sequential Pseudo-Likelihood Estimation of Dynamic Discrete Games

Adam Dearing*

January 15, 2019

[Work in Progress]

**Abstract**

We propose a class of sequential pseudo-likelihood estimators for structural economic models with an equality constraint, with particular focus on dynamic discrete games. We show that under some regularity conditions, all estimates in the sequence are asymptotically efficient and that the sequence converges to the true parameter values with probability approaching one as the sample size grows. The key insight is to apply a single Newton iteration to the fixed point equation in estimation. We show that the nested pseudo-likelihood (NPL) estimator of Aguirregabiria and Mira (2002; 2007) is a special case of our sequential estimator in single-agent models but not in dynamic games. Furthermore, we show that a change of variable in the equilibrium fixed point equation results in an estimator that is no more computationally burdensome than NPL when flow utility is linear in parameters.

## 1 Introduction

Estimation of structural models characterized by an equality constraint is a topic of considerable interest in economics, particularly for dynamic discrete choice models. The equality constraint is often a fixed point constraint, although it can more generally be expressed as $G(\theta, Y) = 0$, where $\theta$ is a finite-dimensional vector of structural parameters and $Y$ is some important parameter (e.g., conditional choice probabilities). Common examples of $Y$ include expected/integrated value functions and conditional choice probabilities. One approach to

*Department of Economics, The Ohio State University, 401 Arps Hall, 1945 N. High St., Columbus OH 43210. Email: dearing.30@osu.edu.

estimating these models is to directly impose the fixed point equation in estimation, solving for every $Y(\theta)$ such that $G(\theta, Y(\theta)) = 0$ at every guess of $\theta$ in the estimation algorithm. In dynamic discrete choice, prominent examples of this approach include Rust (1987) and Pakes (1986) for single-agent models. Pakes and McGuire (1994; 2001) provide solution algorithms dynamic games.

In dynamic models, solving for $Y(\theta)$ for multiple values of $\theta$ can be computationally prohibitive. This lead to the development of sequential estimators, many of which are variants of the conditional choice probability (CCP) estimator first introduced in the seminal work of Hotz and Miller (1993). Of particular interest here is the nested pseudo-likelihood approach of Aguirregabiria and Mira (2002; 2007).[1] They suggest using a $k$-step nested pseudo-likelihood ($k$-NPL) approach, which defines a sequence of estimators. When $k = \infty$ and the sequence converges, one obtains the nested pseudo-likelihood (NPL) estimator. In single-agent models, Aguirregabiria and Mira (2002) show that the $k$-NPL estimator is "efficient" for $k \geq 1$ when initialized with a consistent estimate in the sense that it is asymptotically equivalent to the (partial) maximum likelihood estimator. Furthermore, Kasahara and Shimotsu (2012) showed that the sequence converges to the true parameter values with probability approaching one in large samples. However, these attractive properties of NPL are lost in dynamic games. Aguirregabiria and Mira (2007) show that $k$-NPL estimates are in general not efficient for $k \leq \infty$, although they show that the $\infty$-NPL estimator outperforms the 1-NPL estimator in efficiency when both are consistent. Pesendorfer and Schmidt-Dengler (2010) and Kasahara and Shimotsu (2012) show that the sequence may fail to converge to the "correct" equilibrium, even with a very good initial value, so that $\infty$-NPL may not be consistent.

The primary contribution of this paper is to provide a sequential method that extends the attractive properties of $k$-NPL in a single-agent setting to dynamic discrete games. To this end, we introduce the "$k$-step efficient pseudo likelihood" ($k$-EPL) estimator, with several implementations that are asymptotically equivalent. We formulate the estimator in a setting with a general equality constraint and establish asymptotic equivalence to full maximum-likelihood for $k \geq 1$, as well as convergence to the "correct" value. Intuitively, these properties hold because each iteration of the $k$-EPL is similar to taking a single Newton-Raphson step on the full maximum likelihood problem from a consistent initial estimate. It is well-known that Newton-Raphson steps are locally convergent and that taking a single step from a consistent estimate yields an estimate that is both consistent and asymptotically efficient.

With respect to dynamic discrete choice in particular, we show that in single agent

[1] Some other examples of CCP estimators are described in Hotz et al. (1994); Pakes et al. (2007); Pesendorfer and Schmidt-Dengler (2008).

models the $k$-NPL estimator is an implementation of $k$-EPL. In dynamic games, we show that convergence of $\infty$-NPL is sufficient for convergence of $\infty$-EPL and its variants, so that $\infty$-EPL can be applied to a larger class of models. [This has not actually been added yet!]

One of the distinctive and attractive features of $k$-NPL is that it simplifies each optimization problem in the sequence when utility is linear in the parameters of interest. For example, when the private information shocks are distributed i.i.d. Type 1 Extreme Value, each optimization problem reduces to static multinomial logit optimization. We are able to preserve this attractive feature in $k$-EPL by introducing a change of variable in the equilibrium fixed point equation. Rather than characterizing the equilibrium with choice probabilities (as in $k$-NPL), we instead characterize it with choice-specific value functions. This change of variable is necessary in games because of interaction terms between the parameters and the choice probabilities in the expected flow utilities. We show that attempts to use the $k$-NPL fixed point conditions results in additional non-linearities in $k$-EPL, even in simple static models.

In a related paper, Bugni and Bunting (2018) derive a sequence of asymptotically consistent and efficient minimum-distance estimators for dynamic discrete choice models – including dynamic games – which they call $K$-MD. Pesendorfer and Schmidt-Dengler (2008) showed that likelihood-based estimators can be represented as minimum-distance estimators for dynamic discrete choice models, so it is perhaps unsurprising that an efficient sequence of pseudo-likelihood estimators can also be constructed. Aside from the use of a likelihood objective function, another way our paper differs from Bugni and Bunting (2018) is that we prove convergence of $k$-EPL as $N, k \to \infty$, whereas convergence of $K$-MD is still an open question to the best of our knowledge.[2]

The remainder of the paper proceeds as follows. Section 2 describes the $k$-EPL estimator and its asymptotic properties. Section 3 describes the generic single and multiple-agent discrete choice model. Section 4 provides some example applications, including Monte-Carlo simulations. Section 5 concludes. All formal proofs are relegated to the Appendix.

[2] The analysis of Kasahara and Shimotsu (2012) shows that consistency of $k$-NPL depends on stability of the equilibrium (i.e., that it can be computed using fixed point iteration). When the equilibrium is unstable, the use of fixed point iterates to update the choice probabilities between estimation iterations results in inconsistency of $k$-NPL for $k \geq 2$ and particularly for $k = \infty$. Because Pesendorfer and Schmidt-Dengler (2008) pseudo-likelihood estimates are asymptotically a special case of minimum-distance estimates, we conjecture that the same properties apply to $K$-MD, although a rigorous analysis is beyond the scope of this paper.

# 2 The $k$-EPL Estimator

This section describes the $k$-EPL estimator and provides a simple example application to illustrate its performance relative to $k$-NPL.

## 2.1 Description and Properties of $k$-EPL

We begin by describing the model and discussing full maximum likelihood estimation. There are observed outcomes, $s_i \in \mathcal{S}$, where $i = 1, \ldots, N$ indexes the observations. The model is parameterized by a finite-dimensional vector, $\theta \in \Theta \subset \mathbb{R}^{|\Theta|}$, and a constraint $G(\theta, Y) = 0$ where $Y \in \mathcal{Y} \subset \mathbb{R}^{|\mathcal{Y}|}$ and $G : \Theta \times \mathcal{Y} \to \mathbb{R}^{|\mathcal{Y}|}$. The true parameter values are $\theta^*$ and $Y^*$, with $G(\theta^*, Y^*) = 0$. Note that there may be other values of $Y$ satisfying the constraint at $\theta^*$, but we will assume that the data are generated from only one such value. We denote the partial derivatives of $G$ at the true values as $\nabla_\theta G^*$ and $\nabla_Y G^*$. Define

$$
\begin{aligned}
\hat{Q}(\theta, Y) &= N^{-1} \sum_{i=1}^{N} Q_i(\theta, Y) \\
&= N^{-1} \sum_{i=1}^{N} \ln Pr(s_i|\theta, Y)
\end{aligned}
$$

and $\tilde{\theta}(Y) = \arg\max_{\theta\in\Theta} \hat{Q}(\theta, Y)$. Also, define $Q^*(\theta, Y) = E[Q_N(\theta, Y)]$, and $\tilde{\theta}^*(Y) = \arg\max_{\theta\in\Theta} Q(\theta, Y)$.

**Assumption 1.** *(a) The observations $\{s_i : i = 1, \ldots, N\}$ are independent. (b) $\Theta$ and $\mathcal{Y}$ are compact and convex and $(\theta^*, Y^*) \in int(\Theta \times \mathcal{Y})$. (c) $\hat{Q}(\theta, Y)$ and $Q(\theta, Y)$ are twice continuously differentiable. They have unique maxima in $\Theta \times \mathcal{Y}$ subject to $G(\theta, Y) = 0$, and the maximum occurs at $(\theta^*, Y^*)$ for $Q$. (d) $G(\theta, Y)$ is thrice continuously differentiable and $\nabla_Y G^*$ is non-singular.*

Assumptions 1(a)-(c) echo standard identification assumptions. One consequence of Assumption 1(d) is that we can invoke the implicit function theorem to get twice-continuously differentiable $Y(\theta)$ in some neighborhood of $\theta^*$, $\mathcal{B}^* \subset \Theta$. We can define a modified limiting objective function over this region: $Q(\theta) = Q(\theta, Y(\theta))$. We now define $(\hat{\theta}_{MLE}, \hat{Y}_{MLE}) = \arg\max_{(\theta,Y)\in\Theta\times\mathcal{Y}} \hat{Q}(\theta, Y)$ subject to $G(\theta, Y) = 0$. We make a further assumption on the behavior of $G$ at this solution.

**Assumption 2.** *$\nabla_Y G(\hat{\theta}_{MLE}, \hat{Y}_{MLE})$ is non-singular.*

**Proposition 1.** *Under Assumptions 1(a)-(d),* $\sqrt{N}(\hat{\theta}_{MLE} - \theta^*) \xrightarrow{d} \mathcal{N}(0, V)$ *where* $V$ *is the inverse of the information matrix of* $Q(\theta)$ *evaluated at* $\theta^*$*:*

$$V = \Omega_{\theta\theta}^{*-1}.$$

Proposition 1 gives the asymptotic consistency, normality, and efficiency of the full maximum likelihood estimator. Although such a result is quite powerful, there may be substantial practical difficulties in implementing such an estimator. The main difficulty is the need to find all values of $Y$ solving the constraint equation at multiple values of $\theta$. Assume, for a moment, that $\nabla_Y G(\theta, Y)^{-1}$ exists and is non-singular for all $(\theta, Y) \in \Theta \times \mathcal{Y}$. We can solve for $Y_\theta = Y(\theta)$ via Newton-Kantorovich iterations of the form

$$\begin{aligned} Y_{k+1} &= Y_k - (\nabla_Y G(\theta, Y_k))^{-1} G(\theta, Y_k) \\ &= \Upsilon(\theta, Y_k). \end{aligned}$$

When the initial value, $Y_0$, is close enough to the solution, the Newton-Kantorovich iterates will converge locally and at a quadratic rate.[3] This property, along with some others, are detailed in the next proposition.

**Proposition 2.** *Under Assumption 1, if* $\nabla_Y G(\theta, Y_\theta)$ *is non-singular and* $\Upsilon(\theta, Y) = Y - (\nabla_Y G(\theta, Y))^{-1} G(\theta, Y)$*, then*

1. $\Upsilon(\theta, Y_\theta) = Y_\theta$.

2. $\nabla_\theta \Upsilon(\theta, Y_\theta) = Y'(\theta)$.

3. $\nabla_Y \Upsilon(\theta, Y_\theta) = 0$.

4. There exists some $\delta > 0$ such that iterations of the form $Y_{k+1} = \Upsilon(\theta, Y_k)$ converges to $Y_\theta$ when the starting value, $Y_0$, is an element of $\mathcal{B}_\delta = \{Y \in \mathcal{Y} : ||Y - Y_\theta|| \leq \delta\}$. Furthermore, the rate of convergence is quadratic.

Proposition 2 is the key result of this section, which will become apparent as the exposition proceeds. In particular, we are interested in applying the results at $(\theta^*, Y^*)$. For now, we note that Result 3 of Proposition 2 is analogous to the "zero Jacobian" property from Proposition 2 of Aguirregabiria and Mira (2002), which was the key to their efficiency results.

[3] In some applications, Newton-Kantorovich iterations will converge from any initial guess, $Y_0$. One such example is the integrated value function in single-agent dynamic discrete choice with a finite state (and action) space. Aguirregabiria and Mira (2002) and Dearing (2018) provide two different proofs of this result.

In order to motive our sequential estimator, we will consider alternatives to the original maximum likelihood problem. First, consider an alternative maximization problem where the constraint $G(\theta, Y) = 0$ is replaced with the fixed point constraint $Y = \Upsilon(\theta, Y)$, and define

$$(\tilde{\theta}, \tilde{Y}) = \arg\max_{\theta, Y} \quad Q(\theta, Y)$$

$$\text{s.t.} \quad Y = \Upsilon(\theta, Y).$$

More succinctly, $\tilde{\theta} = \arg\max_\theta Q(\theta, \Upsilon(\theta, Y(\theta))$. Suppose that this has a unique solution. Then, we can use Proposition 2 to show that the first order condition is satisfied at $\theta = \theta^*$, implying that $\tilde{\theta} = \theta^*$ and $\tilde{Y} = Y^*$, so that the solution is equivalent to the (limiting) maximum likelihood estimate. This equivalence is the crux of our efficiency results later on.

Now, suppose that $Y^*$ is known but that $\theta^*$ is unknown. Then we redefine $\tilde{\theta}$ as the solution to yet another alternative problem:

$$\tilde{\theta} = \arg\max_{\theta} \quad Q(\theta, \Upsilon(\theta, Y^*)).$$

Again, if the problem has a unique solution, then $\tilde{\theta} = \theta^*$. However, this problem is infeasible for two reasons. First, the form of $Q(\theta, Y)$ will be unknown. Second, $Y^*$ is also unknown. However, we can construct a feasible version of this estimator instead. To do this, we will introduce $\gamma = (\theta, Y)$ and construct estimates using $\hat{Q}$ and $\hat{\gamma}$ instead. This allows us to define our sequential estimation algorithm.

**Algorithm 1.** *(k-step Efficient Pseudo-Likelihood)*

- *Step 1: Obtain strongly $\sqrt{N}$-consistent initial estimates, $\hat{\gamma}_0 = (\hat{\theta}_0, \hat{Y}_0)$.*
- *Step 2: For $k \geq 1$, define*

$$\Upsilon_k(\theta, \hat{\gamma}_{k-1}) = \hat{Y}_k - (\nabla_Y G(\hat{\theta}_{k-1}, \hat{Y}_{k-1}))^{-1} G(\theta, \hat{Y}_{k-1})$$

*and obtain estimates iteratively:*

$$\hat{\theta}_k = \mathit{arg\ max}_{\theta \in \Theta} \quad \hat{Q}\left(\theta, \Upsilon_k(\theta, \hat{\gamma}_{k-1})\right)$$

*and*

$$\hat{Y}_k = \Upsilon_k(\hat{\theta}_k, \hat{\gamma}_{k-1}).$$

- *Step 3: Iterate on Step 2.*

Algorithm 1 defines the $k$-EPL procedure. Notice that $\Upsilon_k(\cdot)$ in the algorithm is defined somewhat differently than $\Upsilon(\cdot)$ in the preceding analysis, so that it gives approximate instead of exact Newton steps. We have found this definition to be more useful, although the researcher could still use full Newton steps if desired. The intuition from full Newton steps also applies to approximate Newton steps, so the change is made for computational reasons and does not affect the theoretical results. We will elaborate more on various possible definitions of $\Upsilon_k(\cdot)$ shortly.

We claimed in the introduction that this gives a sequence of asymptotically efficient estimators that converge in large samples. We now state this result formally.

**Theorem 1.** *Under Assumptions 1(a)-(d) and some additional regularity assumptions, the $k$-EPL estimates computed with Algorithm 1 satisfy the following for all $k \geq 1$:*

1. $\hat{\gamma}_k = (\hat{\theta}_k, \hat{Y}_k)$ *is a strongly consistent estimator of* $(\theta^*, Y^*)$.
2. $\sqrt{N}(\hat{\theta}_k - \theta^*) \xrightarrow{d} \mathcal{N}(0, \Omega_{\theta\theta}^{*-1})$.
3. *As $N \to \infty$, the sequence of estimators, $\{\hat{\theta}_k, \hat{Y}_k\}_{k=1}^{\infty}$, converges to $(\hat{\theta}_{MLE}, \hat{Y}_{MLE})$ almost surely. Furthermore, the sequence $\{\hat{Y}_k\}_{k=0}^{\infty}$ converges at a superlinear rate.*

The results of Theorem 1 follow from Proposition 2. In particular, the "zero Jacobian" property (Result 3 of Proposition 2) ensures that $\hat{\theta}_k$ and $(\hat{\theta}_{k-1}, \hat{Y}_{k-1})$ are asymptotically orthogonal, so using $(\hat{\theta}_{k-1}, \hat{Y}_{k-1})$ is asymptotically equivalent to using $(\theta^*, Y^*)$ at each step. Asymptotic equivalence to MLE is then obtained using reasoning similar to the analysis preceding the description of Algorithm 1, explaining the asymptotic distribution. Convergence of the sequence follows from local convergence of the Newton-Kantorovich iterations (Result 4 of Proposition 2). In large samples, both $\hat{Y}_{k-1}$ and $\hat{\theta}_k$ are very close to $Y^*$ and $\theta^*$, respectively. Because $\hat{Y}_k$ is obtained with a Newton-Kantorovich step, an iteration of the algorithm is very similar to performing a Newton-Raphson iteration on the full MLE problem, which is a locally convergent operation. The convergence result is important, as it implies that the inconsistency result for $\infty$-NPL from Pesendorfer and Schmidt-Dengler (2010) does not apply to $\infty$-EPL.

We have already mentioned that the specification of $\Upsilon_k(\cdot)$ used in the algorithm could be replaced with full Newton steps without affecting the asymptotic results. These are only two of several choices that yield the same asymptotic results, as shown in the next theorem.

**Theorem 2.** *The results of Theorem 1 hold when $\Upsilon_k(\theta)$ is defined as any of the following:*

1. $\Upsilon_k(\theta, \hat{\gamma}_{k-1}) = \hat{Y}_{k-1} - (\nabla_Y G(\hat{\theta}_{k-1}, \hat{Y}_{k-1}))^{-1} G(\theta, \hat{Y}_{k-1})$.

*2.* $\Upsilon_k(\theta, \hat{\gamma}_{k-1}) = \hat{Y}_{k-1} - Z(\hat{\theta}_{k-1}, \hat{Y}_{k-1})^{-1} G(\theta, \hat{Y}_{k-1})$*, where* $Z(\cdot)$ *is a continuously differentiable function and* $Z(\theta, Y_\theta) = \nabla_Y G(\theta, Y_\theta)$ *for all* $\theta$*.*

*3.* $\Upsilon_k(\theta, \hat{\gamma}_{k-1}) = \hat{Y}_{k-1} - (\nabla_Y G(\theta, \hat{Y}_{k-1}))^{-1} G(\theta, \hat{Y}_{k-1})$*.*

The first definition of $\Upsilon_k(\theta)$ in the theorem is the one we have worked with so far. The second definition is a generalization of the first and can allow researchers to circumvent the need for an initial $\hat{\theta}_0$ if they can find some $Z(\hat{\theta}_{k-1}, \hat{Y}_{k-1}) = Z(\hat{Y}_{k-1})$ or even $Z(\hat{\theta}_{k-1}, \hat{Y}_{k-1}) = A$ that has the required properties.[4] We will show later on that this can be used in single-agent dynamic discrete choice models. The third definition is likely the least useful, as it requires inverting $G_Y(\theta, \hat{Y}_{k-1})$ at multiple values of $\theta$, which can be computationally burdensome and also will introduce additional nonlinearities in the objective function for optimization.[5] However, we include it for completeness. For all of the definitions of $\Upsilon_k(\theta)$ in the theorem, the results of Proposition 2 hold when all appropriate terms are replaced with $(\theta^*, Y^*)$ and or $(\hat{\theta}_{MLE}, \hat{Y}_{MLE})$. So, the proof techniques from Theorem 1 can be used to prove Theorem 2. We prove them simultaneously in the Appendix.

One potential concern is that the parametric model may be misspecified. In this case, we will lose asymptotic equivalence of the $k$-EPL iterates and MLE for finite $k$. However, we can still use the EPL algorithm to compute the MLE by iterating to convergence. Furthermore, the iterates will converge superlinearly from a "good enough" starting guess, $(\tilde{\theta}_0, \tilde{Y}_0)$.

**Proposition 3.** *(Convergence in Finite Samples) Suppose Assumptions 1 and 2 hold and that* $\tilde{\theta}(\gamma)$ *is twice continuously differentiable. Then there exists some neighborhood of* $\hat{\gamma}_{MLE}$*,* $\mathcal{B}$*, such that the iterates converge to* $(\hat{\theta}_{MLE}, \hat{Y}_{MLE})$ *from any starting value* $\tilde{\gamma}_0 \in \mathcal{B}$*. Furthermore, the convergence rate for* $\tilde{Y}_k$ *is superlinear.*

## 2.2 Example: Static Game with an Unstable Equilibrium

In order to illustrate the performance of $k$-EPL, we consider estimating the static game of Pesendorfer and Schmidt-Dengler (2010). This example is particularly interesting, as it was constructed as an example where converged NPL is inconsistent. We discuss only some relevant details model and refer the reader to the original paper for a full description.

There are two agents (players), $j \in \{1, 2\}$, and two possible actions, $a \in \{0, 1\}$. The structural parameter is a scalar: $\theta \in [-10, 1]$. The choice probabilities are $Pr(a_j = 1|\theta, P_{-j}) = 1 - F_\alpha(-\theta P_{-j})$, where $0 < \alpha < 1$. $F_\alpha(x) = x$ for $x \in [\alpha, 1-\alpha)$ and has a more complicated

[4] Of course, $Z(\hat{\theta}_{k-1}, \hat{Y}_{k-1}) = \nabla_Y G(\hat{\theta}_{k-1}, \hat{Y}_{k-1})$ is an option.

[5] In truth, it only requires that one solve the linear system $G_Y(\theta, \hat{Y}_{k-1})b = G(\theta, \hat{Y}_{k-1})$ for $b$. However, the point about computational burden remains.

form for $x \in \mathbb{R}\backslash[\alpha, 1-\alpha)$. The probability mass in the uniform region can be made arbitrarily close to 1 by taking $\alpha \to 0$. Given a value of $\theta$, the model has three equilibria for $\alpha$sufficiently close to zero. The equilibrium generating the data is described by the following fixed point equation:

$$\begin{bmatrix} P_1 \\ P_2 \end{bmatrix} = \begin{bmatrix} 1+\theta P_2 \\ 1+\theta P_1 \end{bmatrix}$$

$$= \begin{bmatrix} 1 \\ 1 \end{bmatrix} + \begin{bmatrix} 0 & \theta \\ \theta & 0 \end{bmatrix} \begin{bmatrix} P_1 \\ P_2 \end{bmatrix}$$

or more compactly, $P = \Psi(\theta, P)$. This linear system has a solution if and only if $\theta \neq -1$, and the solution is $P_1 = P_1 = \frac{1}{1-\theta}$. The $k$-NPL iterates are defined by

$$\hat{\theta}_k^{NPL} = \underset{\theta \in [-10,-1]}{\arg\max} \quad \hat{Q}(\theta, \hat{P}_{k-1})$$

$$\hat{P}_{k-1} = \Psi(\hat{\theta}_k^{NPL}, \hat{P}_{k-1}).$$

We note that $\hat{Q}_i(\theta, \hat{P}_{k-1})$ is the log of a linear function of $\theta$, a point we will return to later. Pesendorfer and Schmidt-Dengler (2010) consider the case where $\theta^* = -2$, implying $P_1^* = P_2^* = \frac{1}{3}$. They show that as $N \to \infty$, $\hat{\theta}_\infty^{NPL} \xrightarrow{p} -1$. Rather than repeat their full explanation of this result, we instead focus on explaining why the sequence does not converge to $\theta^* = -2$. The reason is, essentially, because the equilibrium is unstable. Notice that

$$\nabla_P \Psi(\theta^*, P^*) = \begin{bmatrix} 0 & -2 \\ -2 & 0 \end{bmatrix},$$

which has eigenvalues $\lambda = \pm 2$, implying that the equilibrium is unstable. Kasahara and Shimotsu (2012) show that the non-convergence issue in $k$-NPL can be rectified by estimating separate parameters for each player. However, such an adjustment may not always be practical in more realistic applications.

Consider, instead, estimating $\theta^*$ with $k$-EPL. We still want the objective function to be linear in $\theta$, which will require a change of variable in the equilibrium fixed point equation. Define $v_j = \theta P_{-j}$ and consider the following re-characterization of the fixed point equation:

$$\begin{bmatrix} P_1 \\ P_2 \\ v_1 \\ v_2 \end{bmatrix} = \begin{bmatrix} 1+v_1 \\ 1+v_2 \\ \theta P_2 \\ \theta P_1 \end{bmatrix}$$

$$\Rightarrow \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \begin{bmatrix} (1+v_2)\theta \\ (1+v_1)\theta \end{bmatrix}.$$

So, we can define $Y = (v_1, v_2)$ and therefore

$$\begin{aligned} G(\theta, Y) &= Y - \left( \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} Y + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right) \theta \\ &= Y - (AY + b)\,\theta \end{aligned}$$

Because $\theta$ is a scalar, $G(\theta, Y)$ is bilinear in $(\theta, Y)$. Linearity in $Y$ is important because $G(\theta, Y) = 0$ can be solved with a single Newton-Kantorovich iteration from any starting value. So, we expect that this global one-step convergence will result in very good behavior of $k$-EPL. Additionally, we see that $\nabla_Y G(\theta, Y) = I - A\theta$ and we can easily verify that this is invertible if and only if $\theta \neq -1$. And since $v_j = \theta P_{-j}$, we can also define $Q_i(\theta, Y) = Q_i(v)$, so that $\theta$ only influences $\hat{Q}(\cdot)$ through $Y(\theta)$. This modification is made without loss of generality in full MLE subject to the (expanded) equilibrium constraint, so it is also valid here.

We define $\Upsilon_k(\theta) = \hat{Y}_{k-1} - (\nabla_Y G(\hat{\theta}_{k-1}, \hat{Y}_{k-1}))^{-1} G(\theta, \hat{Y}_{k-1})$. Because $G(\theta, Y)$ is linear in $\theta$, $\Upsilon_k(\theta)$ is also linear in $\theta$. This means that $\hat{Q}_i(\theta, \Upsilon_k(\theta)) = \hat{Q}_i(\Upsilon_k(\theta))$ is the log of a linear function of $\theta$, so there are no additional nonlinearities in the objective function, relative to $k$-NPL. If we instead tried to use $Y = (P_1, P_2)$, then we would still get $\Upsilon(\theta, Y)$ linear in $\theta$, but $Q_i(\cdot)$ would then depend on $\theta\Upsilon(\theta, Y)$, which is non-linear in $\theta$.

All that remains now is to obtain $\hat{Y}_0$ and $\hat{\theta}_0$. Notice that the best response equations imply $\theta = \frac{P_j - 1}{P_{-j}}$ for $j \in \{1, 2\}$. So first, we obtain frequency estimators for $\hat{P}_{1,0}$ and $\hat{P}_{2,0}$. We then use these to construct

$$\hat{\theta}_0 = \frac{\frac{\hat{P}_{1,0}-1}{\hat{P}_{2,0}} + \frac{\hat{P}_{2,0}-1}{\hat{P}_{1,0}}}{2}$$

$$\hat{v}_{j,0} = \hat{\theta}_0 \hat{P}_{-j,0}.$$

We run Monte-Carlo simulations of this model to illustrate the performance of the estimators. We simulate 500 samples, each with 10,000 observations (5,000 per player). We estimate the model using MLE, converged $\infty$-EPL, and converged $\infty$-NPL.[6] Both the MLE and $\infty$-EPL estimates achieve mean -2.0017 and -2.0014, respectively, and mean squared error (MSE) 0.0017 and 0.0017. The two-sample Kolmogorov-Smirnov p-value is equal to 1. Furthermore, $k$-EPL obtained convergence at $k = 2$ in all 500 datasets. This is unsur-

[6]By $\infty$-EPL and $\infty$-NPL, we mean that we iterate until $||\hat{\theta}_k - \hat{\theta}_{k-1}||_\infty < 10^{-6}$. We allow for up to 20 iterations. Estimation was performed with MATLAB R2017a using 'fmincon.' The default tolerance of $10^{-6}$ is used for the solver.

prising: with so many observations and only two players/actions, we get *very* precise initial estimates and iteration converges superlinearly. The slight difference in means and MSE are likely due to a combination of the tolerance used in estimation and non-linearity in the full MLE objective function. On the other hand, converged $\infty$-NPL performs poorly, as expected. The estimate has a mean of -1.0342 and MSE of 0.9651. Almost all of the MSE due to the asymptotic bias, so the estimate is reliably converging to the wrong number.[7]

In addition to demonstrating the good performance of $k$-EPL, this example also introduces the change of variable that will be used in dynamic games. Notice that $v_j$ is essentially a choice-specific utility function; we only need to know $v_j$ to calculate the choice probability for player $j$. We showed that the change of variable is needed to avoid additional nonlinearities in the objective function relative to $k$-NPL here, and the same will be true in dynamic games.

# 3 Dynamic Discrete Choice

[Authors' note: much of this section is still to be written, including many the preliminaries. It will essentially follow Aguirregabiria and Mira (2007) up until this point.]

Now consider the $|\mathcal{X}|\times 1$ vector of player $j$'s (expected) choice-specific value functions, $v^j$, and define the corresponding choice probabilities as $\Lambda^j(v^j)$. In equilibrium, the choice probabilities will be $P_a^j = \Lambda_a^j(v^j)$. And let $\Lambda^{-j}(v^{-j}) = (\Lambda^1(v^1), \dots \Lambda^{j-1}(v^{j-1}), \Lambda^{j+1}(v^{j+1}), \dots, \Lambda^{|\mathcal{J}|}(v^{|\mathcal{J}|}))$, so that in equilibrium $P^{-j} = \Lambda^{-j}(v^{-j})$. Furthermore, define the function

$$\Phi_a^j(\theta, v^j, v^{-j}) = u_a^j(\theta, \Lambda^{-j}(v^{-j})) + \beta F_a^j(\theta, \Lambda^{-j}(v^{-j}))S(v^j),$$

where $\Phi : \Theta \times \mathbb{R}^{|\mathcal{J}|\times|\mathcal{X}|\times|\mathcal{A}|} \to \mathbb{R}^{|\mathcal{J}|\times|\mathcal{X}|\times|\mathcal{A}|}$ and $S(\cdot)$ is the social surplus function. This function allows us to characterize the equilibrium fixed point equation with a convenient change of variable.

**Lemma 1.** *(Representation Lemma) Choice-specific value functions characterize an equilibrium for $\theta$ if and only if $v^j = \Phi^j(\theta, v^j, v^{-j})$ for all $j \in \mathcal{J}$. More compactly,*

$$v = \Phi(\theta, v).$$

Lemma 1 describes our change of variable that will be used to implement Algorithm 1 in estimation. This is in contrast to the representation lemma from Aguirregabiria and

[7] There were 17 samples for which NPL converged in 3 or fewer iterations. The mean and MSE for these samples were -1.9932 and 0.0012, respectively. For the other 483 samples, convergence took at least 12 iterations. These had a mean and MSE of 0.9991 and 0.9991, respectively.

Mira (2007), which uses choice probabilities to characterize the equilibrium. In short, there is a one-to-one mapping between choice probabilities and (expected) choice-specific value functions, conditional on the (conditional) flow utilities and transition probabilities. We use this representation to define our constraints with $Y = v$ and

$$G(\theta, Y) = Y - \Phi(\theta, Y).$$

So, why do we choose this particular representation of the equilibrium? There are two reasons. The first is to preserve a nice property of the pseudo-likelihood problem in $k$-NPL: linearity in $\theta$ of the (pseudo-)choice-specific value functions used to compute the pseudo-likelihood function. That is, when $u_a^j(\theta, P^{-j})$ is linear in $\theta$ and $F_a^j$ does not depend on $\theta$, the computational burden of each pseudo-likelihood estimation reduces to that of a linear-in-parameters static model.[8] In Section 2.2, we showed that this property won't hold for $k$-EPL when we use a fixed point in probability space, even in a static model. However, it will hold when we use the representation in choice-specific value function space due to linearity of $\Phi(\theta, v)$ in $\theta$ under these conditions. By defining $Y = v$ and $G(\theta, Y) = Y - \Phi(\theta, Y)$, we get linearity of $G(\theta, Y)$ in $\theta$. And finally, we have $\Upsilon_k(\theta, \hat{\gamma}_{k-1}) = \hat{Y}_{k-1} - (\nabla_Y G(\hat{\theta}_{k-1}, \hat{Y}_{k-1}))^{-1} G(\theta, \hat{Y}_{k-1})$, which is also linear in $\theta$. The function $\Upsilon_k(\theta, \hat{\gamma}_{k-1})$ determines the choice-specific value functions used to compute the pseudo-likelihood, so these will be linear in $\theta$.

The second reason we use this particular representation of the choice-specific value functions is because of concerns with validity of the (quasi-)Newton steps computed with $\Upsilon_k(\theta, \hat{\gamma}_{k-1})$. If we use the fixed point in probability space to define $Y$ and $G(\theta, Y)$, then we will encounter problems if $\Upsilon_k(\theta, \hat{\gamma}_{k-1})$ does not map into the interior of the simplex over which the probabilities are well-defined, especially when we iterate. Using choice-specific value functions alleviates this concern because they can reside anywhere in $\mathbb{R}^{|\mathcal{J}| \times |\mathcal{X}| \times |\mathcal{A}|}$.

[8] Usually, we have $\theta = (\theta_u, \theta_f)$, $F_a^j(\theta_f, P^{-j})$, and linear-in-parameters $u_a^j(\theta_u, P^{-j})$. We then implement a "two-step" estimator where $\theta_f^*$ is estimated in a first stage and $\theta_u^*$ is estimated via $k$-NPL. Computational burden of each $k$-NPL estimation step then reduces to that of estimating a static model plus an additional matrix inversion.

# 4 Examples and Monte-Carlo Simulations

## 4.1 Single-Agent Dynamic Discrete Choice

Here, we show that $k$-NPL in a single-agent dynamic discrete choice model is a special case of $k$-EPL. Then the model is described by a fixed point constraint,

$$\begin{bmatrix} P \\ V \end{bmatrix} = \begin{bmatrix} \Psi(\theta, P) \\ \varphi(\theta, P) \end{bmatrix}$$

where $\Psi(\theta, P) = \Lambda(\varphi(\theta, P))$, with $\varphi(\cdot)$ and $\Lambda(\cdot)$ defined as in Aguirregabiria and Mira (2002). That is,

$$\varphi(\theta, P) = \left(I - \beta F^U(\theta, P)\right)^{-1} \sum_d P(d) * (u_d(\theta) + e_d(P))$$

is the pseudo-value function, and the function $\Lambda(\theta, V)$ gives choice probabilities. Define $Y = (P, V)$. Then we can redefine the fixed point as

$$Y = \Xi(\theta, Y),$$

and define $G(\theta, Y) = Y - \Xi(Y)$. Then,

$$\nabla_Y G(\theta, Y) = \begin{bmatrix} I - \Psi_P(\theta, P) & 0 \\ -\varphi_P(\theta, P) & I \end{bmatrix}.$$

Now, consider implementing $k$-EPL with $\Upsilon_k(\theta) = \hat{Y}_{k-1} - Z(\hat{\theta}_{k-1}, \hat{Y}_{k-1})^{-1} G(\theta, \hat{Y}_{k-1})$ where $Z(\theta, Y_\theta) = \nabla_Y G(\theta, Y_\theta)$ for all $\theta$. Proposition 2 from Aguirregabiria and Mira (2002) shows that $\nabla_P \Psi(\theta, P_\theta) = 0$ and $\nabla_P \varphi(\theta, P_\theta) = 0$. Thus, $\nabla_Y G(\theta, Y_\theta) = I$ for all $\theta$. So, $Z(\theta, Y) = I$ and

$$\Upsilon_k(\theta, \hat{\gamma}_{k-1}) = \Xi(\theta, \hat{Y}_{k-1}).$$

Finally, we have that $Q_i(\theta, Y) = \ln \Lambda(\theta, V)(s_i)$, implying $Q_i(\theta, \Upsilon_k(\theta)) = \ln \Psi(\theta, \hat{P}_{k-1})(s_i)$. So, this implementation of $k$-EPL is equivalent to $k$-NPL. Note that it is only necessary to supply an initial estimate of $P^*$ because $\hat{\theta}_{k-1}$ and $\hat{V}_{k-1}$ do not directly affect $(\hat{\theta}_k, \hat{Y}_k)$, whereas $\hat{P}_{k-1}$ does.

This equivalence of $k$-NPL to $k$-EPL is unsurprising for a couple reasons. First, we stated in the introduction that the motivation for $k$-EPL is to extend the nice properties of $k$-NPL from single-agent models to dynamic games. So, there should be, at the very least, substantial conceptual overlap between the techniques. Second, Aguirregabiria and Mira (2002) noted an equivalence between their updates and Newton iterations on the value function in single-agent models.[9] Since $k$-EPL is built around Newton iterations, such an equivalence is again suggestive of the relationship shown in this section.

## 4.2 Monte-Carlo Simulations: Dynamic Game from Pesendorfer and Schmidt-Dengler (2008)

[Authors' note: this section is also incomplete. Simulations coming soon.]

We now turn our attention to Monte-Carlo simulations, estimating the model from Pesendorfer and Schmidt-Dengler (2008). There are two firms indexed by $j \in \{1, 2\}$ who choose an action in each time period, $a_j \in \{0, 1\}$, where 1 is entry and 0 is exit. The observed state variable, $x \in \{1, 2, 3, 4\} = \{(0, 0), (0, 1), (1, 0), (1, 1)\}$, indicates the firms' actions in the previous period (with firm 1 listed first). Flow utilities are period profits:

$$u_x^j(a_j = 1) = \theta_M^* + \theta_C^* a_{-j,t} + \theta_{EC}^*(1 - a_{j,t-1}) + \varepsilon_{j,1}$$

$$u_x^j(a_j = 0) = \theta_{SV}^* a_{j,t-1} + \varepsilon_{j,0},$$

where $\theta_{EC}^*$ represents the entry cost, $\theta_{SV}^*$ is the scrap value, $\theta_M^*$ is the monopoly profit, and $\theta_C^*$ is the effect of competition on profit. The discount factor is $\beta^* \in (0, 1)$. The data are generated using the parameter values $(\theta_M^*, \theta_C^*, \theta_{EC}^*, \theta_{SV}^*, \beta^*) = (1.2, -2, -0.2, 0.1, 0.9)$.

There are multiple equilibria in the game, and we generate data from equilibria (i) and (ii) from Pesendorfer and Schmidt-Dengler (2008), both of which are asymmetric. In their Monte-Carlo simulations, Pesendorfer and Schmidt-Dengler (2008) find that $k$-NPL performs well for equilibrium (i) but becomes severely biased for equilibrium (ii) as $k$ grows, which suggests that equilibrium (ii) is unstable.[10]

We estimate $(\theta_M^*, \theta_C^*, \theta_{EC}^*)$ and assume the other parameter values are known.

# 5 Conclusion

To be written . . .

[9] See Proposition 1(c) in Aguirregabiria and Mira (2002).

[10] They use the terminology "$k$-PML" and iterate until $k = 20$.

# References


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

_ **and** _ , "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 2007, *75*, 1–53.

**Bugni, F.A. and J. Bunting**, "On the Iterated Estimation of Dynamic Discrete Choice Games," 2018. Working Paper.

**Hotz, V.J. and R.A. Miller**, "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 1993, *60*, 497–529.

_ **,** _ **, S. Sanders, and J. Smith**, "A Simulation Estimator for Dynamic Models of Discrete Choice," *Review of Economic Studies*, 1994, *61*, 265–289.

**Kasahara, H. and K. Shimotsu**, "Sequential Estimation of Structural Models with a Fixed Point Constraint," *Econometrica*, 2012, *80* (5), 2303–2319.

**McFadden, D.L. and W. Newey**, "Large Sample Estimation and Hypothesis Testing," in R.E. Engle and D.L. McFadden, eds., *Handbook of Econometrics*, Vol. 4, Amsterdam: North-Holland, 1994.

**Ortega, J.M. and W.C. Rheinboldt**, *Iterative Solution of Nonlinear Equations in Several Variables*, SIAM, January 1970.

**Pakes, A.**, "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, 1986, *54*, 1027–1057.

_ **and P. McGuire**, "Computing Markov-Perfect Nash Equilibria: Numberical Implications of a Dynamic Differentiated Product Model," *RAND Journal of Economics*, Winter 1994, *25* (4), 555–589.

_ **and** _ , "Stochastic Approximation for Dynamic Models: Markov Perfect Equilibrium and the 'Curse' of Dimensionality," *Econometrica*, September 2001, *69* (5), 1261–1281.

_ **, M. Ostrovsky, and S. Berry**, "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)," *RAND Journal of Economics*, Summer 2007, *38* (2), 373–399.

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

_ **and** _ , "Sequential Estimation of Dynamic Discrete Games: A Comment," *Econometrica*, 2010, *78* (2), 833–842.

**Rust, J.**, "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 1987, *55*, 999–1013.

# A Proofs

## A.1 Proof of Proposition 2

Result 1 follows from $G(\theta, Y_\theta) = 0$. For Results 2 and 3, first consider partial differentiation of $\Upsilon(\theta, Y)$:

$$\Upsilon_\theta(\theta, Y) = -(\nabla_Y G(\theta, Y))^{-1} \nabla_\theta G(\theta, Y) - \frac{\partial \nabla_Y G(\theta, Y)^{-1}}{\partial \theta} G(\theta, Y)$$

$$\begin{aligned}\Upsilon_Y(\theta, Y) &= I - (\nabla_Y G(\theta, Y))^{-1} \nabla_Y G(\theta, Y) - \frac{\partial \nabla_Y G(\theta, Y)^{-1}}{\partial Y} G(\theta, Y) \\ &= -\frac{\partial \nabla_Y G(\theta, Y)^{-1}}{\partial Y} G(\theta, Y).\end{aligned}$$

Now consider implicit differentiation on $G(\theta, Y(\theta)) = 0$ yields

$$\begin{aligned}Y'(\theta) &= -(\nabla_Y G(\theta, Y(\theta)))^{-1} \nabla_\theta G(\theta, Y(\theta)) \\ &= -(\nabla_Y G(\theta, Y_\theta))^{-1} \nabla_\theta G(\theta, Y_\theta).\end{aligned}$$

Results 2 and 3 then follow from $G(\theta, Y_\theta) = 0$. Result 4 is a widely-known property of Newton-Kantorovich iterations.

## A.2 Proof of Proposition 3

This is similar to proving that Newton-Kantorovich iterations achieve superlinear local convergence. We will first prove the result for $\tilde{Y}_k$ and then use that to establish the result for $\hat{\theta}_k$. First, note that $\hat{\theta}_{MLE} = \tilde{\theta}(\hat{\gamma}_{MLE})$ and $\hat{Y}_{MLE} = \Upsilon_k(\tilde{\theta}(\hat{\gamma}_{MLE}), \hat{\gamma}_{MLE})$ so that $(\hat{\theta}_{MLE}, \hat{Y}_{MLE}) = H(\hat{\gamma}_{MLE})$. Additionally,

$$\nabla_\gamma \tilde{\theta}(\hat{\gamma}_{MLE}) = 0$$

$$\nabla_\gamma \Upsilon_k(\tilde{\theta}(\hat{Y}_{MLE}), \hat{\gamma}_{MLE}) = 0$$

so that $\nabla_\gamma H(\hat{\gamma}_{MLE}) = 0$, which is central to the result.

We now establish convergence of the $\tilde{Y}_k$ iterates. Because we have $\hat{Y}_{MLE} = H_2(\hat{\gamma}_{MLE})$ and $\nabla_\gamma H_2(\hat{\gamma}_{MLE}) = 0$, we immediately obtain local and superlinear convergence on some neighborhood, $\mathcal{B}$, of $\hat{\gamma}_{MLE}$ (Result 10.1.6 in Ortega and Rheinboldt (1970)).

## A.3 Proof of Theorems 1 and 2

The proofs of Results 1 and 2 adapt the proofs of consistency and asymptotic normality for the 1-NPL estimator from Aguirregabiria and Mira (2007) to an inductive proof for $k$-EPL.[11] We do this by showing that strong $\sqrt{N}$-consistency of $\hat{\gamma}_{k-1} = (\hat{\theta}_{k-1}, \hat{Y}_{k-1})$ implies the results for $\hat{\gamma}_k = (\hat{\theta}_k, \hat{Y}_k)$. The proof of Result 3 follows the arguments used in the proof of Proposition 7 in the supplementary material for Kasahara and Shimotsu (2012). Throughout, we rely heavily on analysis similar to that from the proof of Proposition 2.

It is helpful up-front to define $\tilde{Q}(\theta,\gamma) = \hat{Q}(\theta, \Upsilon_k(\theta,\gamma))$ and $\tilde{\theta}(\gamma) = \arg\max_\theta \tilde{Q}(\theta,\gamma)$. Similarly, $\tilde{Q}^*(\theta,\gamma) = E[\tilde{Q}(\theta,\gamma)]$ and $\tilde{\theta}^*(\gamma) = \arg\max_\theta \tilde{Q}^*(\theta,\gamma)$. Then, $\hat{\theta}_k = \tilde{\theta}(\hat{\gamma}_{k-1})$ and $\hat{Y}_k = \Upsilon(\hat{\theta}_k, \hat{\gamma}_{k-1})$.

### A.3.1 Strong consistency of $\hat{\theta}_k$ and $\hat{Y}_k$

We have uniform continuity of $\tilde{Q}^*(\theta,\gamma)$ and that $\tilde{Q}(\theta,\gamma)$ converges almost surely and uniformly in $(\theta,\gamma) \in \Theta \times (\Theta \times \mathcal{Y})$ to $\tilde{Q}^*(\theta,\gamma)$ . Also, $\hat{\gamma}_{k-1}$ converges almost surely to $\gamma^*$. Together, these imply that $\tilde{Q}(\theta, \hat{\gamma}_{k-1})$ converges almost surely and uniformly in $\theta \in \Theta$ to $\tilde{Q}^*(\theta,\gamma^*)$. By assumption, $\theta^*$ uniquely maximizes $\tilde{Q}^*(\theta,\gamma^*)$ on $\Theta$. So, $\hat{\theta}_k$ converges almost surely to $\theta^*$. Continuity of $\Upsilon_k(\theta,\gamma)$ and the Mann-Wald theorem then give almost sure convergence of $\hat{Y}_k$ to $Y^*$.

### A.3.2 Asymptotic Distribution of $\hat{\theta}_k$ and $\hat{Y}_k$

We will show that strong $\sqrt{N}$-consistency of $\hat{\gamma}_{k-1}$ leads to asymptotic normality of $\hat{\theta}_k$ and $\hat{Y}_k$ with their asymptotic variance the same as the MLE. Using any of the definitions of $\Upsilon_k(\cdot)$ in Theorem 2, we can use the chain rule and analysis similar to the proof of Proposition 2 to obtain

$$\nabla_\gamma \tilde{Q}^*(\theta^*,\gamma^*) = 0$$

$$\nabla_{\gamma\gamma} \tilde{Q}^*(\theta^*,\gamma^*) = 0$$

$$\nabla_{\theta\gamma} \tilde{Q}^*(\theta^*,\gamma^*) = 0.$$

So, invoking the Mann-Wald theorem and the generalized information matrix equality (McFadden and Newey (1994, p. 2163)) we obtain

$$\nabla_\theta \tilde{Q}(\theta^*,\gamma^*) \xrightarrow{p} 0$$

[11] See the proofs of Propositions 1 and 2 in the Appendix of Aguirregabiria and Mira (2007).

$$\nabla_{\theta\theta}\tilde{Q}(\theta^*, \gamma^*) \xrightarrow{p} -\Omega^*_{\theta\theta}$$

$$\nabla_{\theta\gamma}\tilde{Q}(\theta^*, \gamma^*) \xrightarrow{p} 0.$$

Consistency of $(\hat{\theta}_k, \hat{\gamma}_{k-1})$ and a stochastic mean value theorem imply that

$$\begin{aligned} 0 = \nabla_\theta \tilde{Q}(\theta^*, \gamma^*) + \nabla_{\theta\theta}\tilde{Q}(\theta^*, \gamma^*)(\hat{\theta}_k - \theta^*) \\ + \nabla_{\theta\gamma}\tilde{Q}(\theta^*, \gamma^*)(\hat{\gamma}_{k-1} - \gamma^*) + o_p(1). \end{aligned}$$

From strong $\sqrt{N}$-consistency of $\hat{\gamma}_{k-1}$ we have $\sqrt{N}(\hat{\gamma}_{k-1} - \gamma^*) = O_p(1)$. By the central limit theorem, we then have that

$$\sqrt{N}(\hat{\theta}_k - \theta^*) = \Omega^{*-1}_{\theta\theta}\frac{1}{\sqrt{N}}\nabla_\theta \tilde{Q}(\theta^*, \gamma^*) + o_p(N^{-1/2}).$$

We already have that $(1/\sqrt{N})\nabla_\theta \tilde{Q}(\theta^*, \gamma^*) \rightarrow_d \mathcal{N}(0, \Omega^*_{\theta\theta})$ and so by the Mann-Wald theorem,

$$\sqrt{N}(\hat{\theta}_k - \theta^*) \xrightarrow{d} \mathcal{N}(0, \Omega^{*-1}_{\theta\theta}).$$

And because $\hat{Y}_k = \Upsilon_k(\hat{\theta}_k, \hat{\gamma}_{k-1})$, with $\Upsilon_k(\cdot)$ twice continuously differentiable in a neighborhood of $(\theta^*, Y^*)$, consistency and asymptotic normality of $\hat{Y}_k$ follow immediately. Asymptotic equivalence of $\hat{Y}_k$ and $\hat{Y}_{MLE}$ follow from asymptotic equivalence of $\hat{\theta}_k$ and $\hat{\theta}_{MLE}$ and from Proposition 2. Strong $\sqrt{N}$-consistency of $\hat{\gamma}_0$ completes the proof by induction.

### A.3.3 Convergence of $\{\hat{\theta}_k, \hat{Y}_k\}_{k=1}^\infty$

First, note the following results that arise from Propositions 2-like derivations, $\sqrt{N}$-consistency of $\hat{\gamma}_{MLE} = (\hat{\theta}_{MLE}, \hat{Y}_{MLE})$, Taylor expansion around $(\theta^*, \gamma^*)$, and the generalized information matrix equality:

$$\nabla_\theta \tilde{Q}(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE}) = 0$$

$$\nabla_{\theta\theta}\tilde{Q}(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE}) = -\Omega^*_{\theta\theta} + O_p(N^{-1/2})$$

$$\nabla_{\theta\gamma}\tilde{Q}(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE}) = 0 + O_p(N^{-1/2})$$

$$\nabla_\theta \Upsilon_k(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE}) = Y'(\theta^*) + O_p(N^{-1/2})$$

$$\nabla_\gamma \Upsilon_k(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE}) = 0.$$

almost surely.

We will first show that $||\hat{\theta}_k - \hat{\theta}_{MLE}|| = O_p(||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}||)$ and then show that $||\hat{Y}_k - \hat{Y}_{MLE}|| = O_p(N^{-1/2}||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}|| + ||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}||^2)$. We have $(\hat{\theta}_k - \hat{\theta}_{MLE}, \hat{Y}_{k-1} - \hat{Y}_{MLE}) = o_p(1)$ almost surely, so consider a Taylor expansion of $\nabla_\theta \tilde{Q}(\theta, \gamma)$ around $(\hat{\theta}_{MLE}, \hat{Y}_{MLE})$:

$$\begin{aligned}
0 = &\nabla_{\theta\theta} \tilde{Q}(\hat{\theta}_{MLE}, \hat{Y}_{MLE})(\hat{\theta}_k - \hat{\theta}_{MLE}) + \nabla_{\theta\gamma} \tilde{Q}(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE})(\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}) \\
&+ o_p(1).
\end{aligned}$$

This can be rearranged to obtain

$$\hat{\theta}_k - \hat{\theta}_{MLE} = -\nabla_{\theta\theta} \tilde{Q}(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE})^{-1} \left[ \nabla_{\theta\gamma} \tilde{Q}(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE})(\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}) + o_p(1) \right].$$

Invoking some strong laws of large numbers, $\nabla_{\theta\theta} \tilde{Q}(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE})^{-1} \nabla_{\theta Y} \tilde{Q}(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE}) = O_p(1)$. So, $||\hat{\theta}_k - \hat{\theta}_{MLE}|| = O_p(||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}||)$. In order to derive the next bound, note that $\hat{Y}_{MLE} = \Upsilon_k(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE})$. Taking a Taylor expansion then gives

$$\begin{aligned}
\hat{Y}_k - \hat{Y}_{MLE} &= \nabla_\theta \Upsilon_k(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE})(\hat{\theta}_k - \hat{\theta}_{MLE}) + \nabla_\gamma \Upsilon_k(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE})(\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}) \\
&\quad + O_p(||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}||^2) \\
&= (Y'(\theta^*) + O_p(N^{-1/2}))O_p(||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}||) + O_p(||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}||^2) \\
&= O_p(N^{-1/2}||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}|| + ||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}||^2),
\end{aligned}$$

where the second equality follows from substituting previous results.

## A.4 Proof of Lemma 1

First, consider player $j$ and fix $(\theta, v^{-j})$. Then let $\phi^j(v^j) = \Phi^j(\theta, v^j, v^{-j})$. We will show that $v^j = \phi(v^j)$ has a unique solution. First, consider $\nabla_v S(v^j) = \boldsymbol{P}(v^j)$, where the rows of $\boldsymbol{P}(v^j)$ are $P_x(v^j)'$ with entries

$$P_x(v^j)\{k\} = \begin{cases} Pr(a|x, v^j), & k = (y, a) \; s.t. \; y = x \\ 0, & k = (y, a) \; s.t. \; y \neq x. \end{cases}$$

Note that $\boldsymbol{P}(v^j)$ is a row-stochastic matrix. Now define

$$F = \begin{bmatrix} F_{a=1}(\theta, \Lambda^{-j}(v^{-j})) \\ F_{a=2}(\theta, \Lambda^{-j}(v^{-j})) \\ \vdots \\ F_{a=|\mathcal{A}|}(\theta, \Lambda^{-j}(v^{-j})) \end{bmatrix},$$

so that $F$ is also row-stochastic. These matrices will allow us to form the derivative,

$$\nabla_v \phi^j(v^j) = \beta F \boldsymbol{P}(v^j),$$

where $F\boldsymbol{P}$ is row-stochastic because $F$ and $\boldsymbol{P}(v^j)$ are both row-stochastic. Now define $g(v^j) = v^j - \phi^j(v^j)$. We have $\nabla_v g(v^j) = I - \beta F \boldsymbol{P}(v^j)$ and can easily show that

$$\left\| \nabla_v g(v^j)^{-1} \right\|_\infty = \frac{1}{1-\beta}.$$

So, we can apply Hadamard's global inverse function theorem to establish that $g(v^j)$ is a homeomorphism on $\mathbb{R}^{|\mathcal{X}| \times |\mathcal{A}|}$ and therefore has a unique root.[12]

[12] Alternatively, $||\nabla_v \phi^j(v^j)||_\infty = ||\beta F \boldsymbol{P}||_\infty = \beta < 1$, so $\phi^j(v^j)$ is a contraction with a unique fixed point.