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

Adam Dearing and Jason R. Blevins†
The Ohio State University

December 22, 2019

Abstract

We propose a new sequential Efficient Pseudo-Likelihood (EPL) estimator for structural economic models with an equality constraint, particularly dynamic discrete choice games of incomplete information. Each iteration in the EPL sequence is consistent and asymptotically efficient, and iterating to convergence improves finite sample performance. For dynamic single-agent models, we show that Aguirregabiria and Mira's (2002; 2007) Nested Pseudo-Likelihood (NPL) estimator arises as a special case of EPL. In dynamic games, EPL maintains its efficiency properties, although NPL does not. And a convenient change of variable in the equilibrium fixed point equation ensures EPL iterations have the same computational simplicity as NPL iterations. Furthermore, EPL iterations are stable and locally convergent to the finite-sample maximum likelihood estimator at a nearly-quadratic rate for all regular Markov perfect equilibria, including unstable equilibria where NPL encounters convergence problems. Monte Carlo simulations confirm the theoretical results and demonstrate EPL's good performance in finite samples.

Keywords: dynamic discrete games, dynamic discrete choice, multiple equilibria, pseudo maximum likelihood estimation.

JEL Classification: C57, C63, C73, L13.

*We thank Dan Ackerberg, Victor Aguirregabiria, Chris Conlon, Arvind Magesan, Mathieu Marcoux, Salvador Navarro, and John Rust for comments and insightful discussions. The paper has benefited from comments by seminar participants at Boston College, Georgetown University, University of Toronto, the 2019 International Industrial Organization Conference (IIOC, Boston), and the 2019 University of Calgary Empirical Microeconomics Workshop (Banff).

†Department of Economics, 410 Arps Hall, 1945 N. High St., Columbus OH 43210. Email: dearing.30@osu.edu and blevins.141@osu.edu.

1 Introduction

Estimation of structural models characterized by equality constraints is a topic of considerable interest in economics, particularly for dynamic discrete choice models and dynamic discrete games of incomplete information. Broadly, likelihood-based estimation of these models takes the form

$$\max_{(\theta,Y)\in\Theta\times\mathcal{Y}} Q_N(\theta, Y)$$
$$\text{s.t. } G(\theta, Y) = 0,$$

where Q_N is the log-likelihood function, θ is a finite-dimensional vector of parameters, Y is a vector of important auxiliary parameters, and $G(\theta, Y) = 0$ is the equality constraint. The parameters θ usually consist of the structural parameters of the model. Common examples of auxiliary parameters Y include expected/integrated value functions and conditional choice probabilities. Finally, the equality constraint is often derived from an equilibrium fixed point condition such as $G(\theta, Y) = Y - \Gamma(\theta, Y) = 0$.

One approach to estimating these models is to directly impose the fixed point equation in estimation by solving for Y_θ such that $G(\theta, Y_\theta) = 0$ for each trial value of θ visited by the optimization algorithm. In dynamic discrete choice, this approach was pioneered by Rust (1987) and Pakes (1986) for single-agent models, where the fixed point is unique. Solution algorithms are available for dynamic games (Pakes and McGuire, 1994, 2001), but due to the computational complexity it is infeasible at present to nest those within an estimation routine. Furthermore, in games the model may be incomplete due to multiple solutions, Y_θ (Tamer, 2003).

These dual issues led researchers to extend conditional choice probability (CCP) estimators, first introduced in the seminal work of Hotz and Miller (1993), to the case of dynamic discrete games. 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, as an algorithm for computing 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

[1] Some other examples of CCP estimators are described in Hotz, Miller, Sanders, and Smith (1994); Bajari, Benkard, and Levin (2007); Pakes, Ostrovsky, and Berry (2007); Pesendorfer and Schmidt-Dengler (2008).

converges to the true parameter values with probability approaching one in large samples.

However, these attractive properties of k-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 ∞-NPL estimator outperforms the 1-NPL estimator in efficiency when both are consistent. Pesendorfer and Schmidt-Dengler (2010), Kasahara and Shimotsu (2012), Egesdal, Lai, and Su (2015), and Aguirregabiria and Marcoux (2019) show that the sequence may fail to converge to the equilibrium that generated the data, even with very good starting values, so that ∞-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 single-agent settings to games. To this end, we introduce the Efficient Pseudo-Likelihood (EPL) estimator. The estimator is based on replacing the equality constraint, $G(\theta, Y) = 0$, with a fixed point condition, $Y = \Upsilon(\theta, Y)$, representing a Newton step in the direction of Y holding θ fixed. We show that this has the benefit of imposing the same equilibrium conditions while rendering the second step estimation of θ orthogonal from the first step estimation of Y. The EPL procedure is also stable for all regular Markov perfect equilibria, including those where the NPL mapping is unstable. We discuss several implementations that are asymptotically equivalent. We formulate the estimator in a setting with a general equality constraint and establish asymptotic equivalence to maximum likelihood for any number of iterations $k \geq 1$, as well as convergence to the finite-sample maximum-likelihood estimate.

One of the distinctive and attractive features of k-NPL is that it simplifies each optimization problem in the sequence when flow utility is linear in the parameters of interest. For example, when the private information shocks are i.i.d. Type 1 Extreme Value draws, each optimization problem reduces to static multinomial logit problem. We are able to preserve this attractive feature in k-EPL by introducing a change of variables 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 variables is necessary in games because of two issues that arise when using choice probabilities directly: i) interaction terms between the parameters and the choice probabilities in the expected flow utilities introduce additional nonlinearities in the objective function; and ii) applying Newton steps may result in choice probabilities outside the unit simplex.[2] However, the change of variables is unnecessary in single-agent models, and we show that the k-NPL estimator can be interpreted as an implementation of k-EPL in that setting.

In a related paper, Bugni and Bunting (2019) derive a sequence of consistent and asymp-

[2]It is also common to work with log choice probabilities, but those are still bounded above by zero, resulting in a similar problem with Newton steps.

totically efficient minimum-distance estimators for dynamic discrete choice models—including dynamic games—which they refer to as k-MD. One basic difference is that our approach is likelihood-based, although Pesendorfer and Schmidt-Dengler (2008) showed that likelihood-based estimators can be represented as minimum-distance estimators for dynamic discrete choice models. Another substantive difference is that k-MD updates are based on the NPL mapping, which can be unstable at the true equilibrium as previously described, so the k-MD estimator may suffer from the same convergence problems as NPL as $k \to \infty$.[3] The efficiency gains of k-MD also come at the expense of additional computational time. On the other hand, the stability of the EPL mapping allows us to prove convergence of k-EPL as $k \to \infty$, and k-EPL does not increase computation time.[4]

Recently, Aguirregabiria and Marcoux (2019) studied the convergence properties of NPL and introduced a variation of the NPL algorithm that updates the conditional choice probabilities using spectral methods applied to the NPL mapping. The goal of their algorithm is to improve convergence properties of NPL for unstable fixed points.[5] However, upon convergence, the NPL and spectral NPL algorithms do not produce the maximum likelihood estimator, and convergence can require many iterations. In contrast, our EPL estimator is asymptotically equivalent to the MLE at each iteration and converges locally to the MLE after few iterations in finite samples. We verify these properties in our simulation study below.

The remainder of the paper proceeds as follows. Section 2 describes the k-EPL estimator and its asymptotic properties, along with some example applications. Section 3 describes the generic dynamic discrete choice game of incomplete information. Section 4 provides Monte Carlo simulations. Section 5 concludes. Proofs appear in the Appendix.

[3]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 standard 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 $k \to \infty$. Because Pesendorfer and Schmidt-Dengler (2008) showed that 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.

[4]The computational times reported for the Monte Carlo experiments of Bugni and Bunting (2019) indicate that one MD iteration requires substantially more computational time than one NPL iteration. One EPL iteration, on the other hand, is approximately the same as one NPL iteration in terms of computational time. Additionally, fewer iterations are required for EPL to converge, resulting in reduced overall computational burden when iterating to convergence ($k \to \infty$).

[5]Stability here refers to the spectral radius of the NPL operator. In the population or in finite samples, the NPL operator may have spectral radius larger than one for some equilibria. Conversely, the spectral radius of the EPL operator is zero in the population and near zero in finite samples due to its use of Newton steps.

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. Let w_i for $i = 1, \ldots, N$ denote 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 θ^* and Y^*, with $G(\theta^*, Y^*) = 0$. Note that there may be other values of Y satisfying the constraint at θ^*, but we will assume that the data are generated from only one such value, a common assumption in the literature. Define

$$Q_N(\theta, Y) = N^{-1} \sum_{i=1}^{N} q_i(\theta, Y) = N^{-1} \sum_{i=1}^{N} \ln f(w_i \mid \theta, Y)$$

and let $Q^*(\theta, Y) = E[Q_N(\theta, Y)]$.

Assumption 1. *(a) The observations $\{w_i : i = 1, \ldots, N\}$ are i.i.d. and generated by a single equilibrium (θ^*, Y^*). (b) Θ and $\mathcal{Y}$ are compact and convex and $(\theta^*, Y^*) \in int(\Theta \times \mathcal{Y})$. (c) $Q_N(\theta, Y)$ 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 (θ^*, Y^*) for Q. (d) $G(\theta, Y)$ is thrice continuously differentiable and $\nabla_Y G(\theta^*, Y^*)$ is non-singular.*

Assumptions 1(a)-(c) echo standard identification assumptions. Non-singularity of the Jacobian in (d) is the defining feature of regular Markov Perfect Equilibria in the sense of Doraszelski and Escobar (2010). We now define

$$\left(\hat{\theta}_{MLE}, \hat{Y}_{MLE}\right) = \arg \max_{(\theta, Y) \in \Theta \times \mathcal{Y}} Q_N(\theta, Y)$$
$$\text{s.t. } G(\theta, Y) = 0.$$

Now, suppose that $\nabla_Y G(\theta, Y)$ exists and is non-singular for all $(\theta, Y) \in \Theta \times \mathcal{Y}$ and define $\Upsilon(\theta, Y)$ to be a Newton step on Y holding θ fixed:

$$\Upsilon(\theta, Y) \equiv Y - \nabla G(\theta, Y)^{-1} G(\theta, Y). \tag{1}$$

Notice that the set of fixed points of $Y = \Upsilon(\theta, Y)$ and the set of roots of $G(\theta, Y) = 0$ are equivalent, so we can substitute the fixed point constraint for the equality constraint.

Furthermore, from an initial value Y_0 we could solve for a root Y_θ via Newton-Kantorovich iterations of the form:

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

When Y_0 is within the basin of attraction of a solution, the Newton-Kantorovich iterates will converge and will do so at a quadratic rate.[6] These properties, along with some others, are detailed in the following lemma.

Lemma 1. *Let Υ denote the Newton operator defined in (1). Under Assumption 1, if $\nabla_Y G(\theta, Y_\theta)$ is non-singular, then the following properties hold:*

1. Roots of G and fixed points of Υ are identical: $\Upsilon(\theta, Y_\theta) = Y_\theta$ if and only if $G(\theta, Y_\theta) = 0$.
2. $\nabla_\theta \Upsilon(\theta, Y_\theta) = \nabla_\theta Y_\theta$.
3. $\nabla_Y \Upsilon(\theta, Y_\theta) = 0$ (Zero Jacobian Property).
4. There exists some $\delta > 0$ such that iterations of the form $Y_{k+1} = \Upsilon(\theta, Y_k)$ converges to Y_θ 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.

Lemma 1 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 (θ^*, Y^*). For now, we note that Result 3 of Lemma 1 is analogous to the "zero Jacobian" property from Proposition 2 of Aguirregabiria and Mira (2002), which was the key to their efficiency results.

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

$$\begin{aligned}(\tilde{\theta}, \tilde{Y}) = &\arg\max_{(\theta, Y) \in \Theta \times \mathcal{Y}} Q(\theta, Y) \\ &\text{s.t.} \quad Y = \Upsilon(\theta, Y).\end{aligned}$$

or more succinctly,

[6] 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 (2019) provide two different proofs of this result.

$$\tilde{\theta} = \arg\max_{\theta \in \Theta} Q(\theta, \Upsilon(\theta, Y_\theta)).$$

Suppose that this has a unique solution. Then, we can use Lemma 1 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 population maximum likelihood estimate. This equivalence is the crux of our efficiency results later on.

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

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

Again, if the problem has a unique solution, then $\tilde{\theta} = \theta^*$. However, this problem is infeasible for two reasons. First, the population function $Q(\theta, Y)$ is unknown. Instead, we use the sample analog, Q_N. Second, Y^* is also unknown. Instead, we can substitute a consistent estimate.

Although the form of Υ we have used so far is useful in motivating our estimator, using it in practice would lead to unnecessary computational burden. Instead, Algorithm 1 below defines our sequential estimation procedure with a more practical form of Υ which carries out a quasi-Newton step given a compound parameter vector $\gamma = (\breve{\theta}, Y)$:[7]

$$\Upsilon(\theta, \gamma) \equiv Y - \nabla_Y G(\breve{\theta}, Y)^{-1} G(\theta, Y). \tag{2}$$

Econometrically, this change makes no substantive difference because we still have the zero Jacobian property: $\nabla_\gamma \Upsilon(\theta, \gamma_\theta) = 0$, where $\gamma_\theta = (\theta, Y_\theta)$. We elaborate on other possible choices of Υ later, in Theorem 3.

Algorithm 1. *(k-step Efficient Pseudo-Likelihood, or k-EPL) Let Υ be defined as in (2).*

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

$$\hat{\theta}_k = \arg\max_{\theta \in \Theta} \quad Q_N\left(\theta, \Upsilon(\theta, \hat{\gamma}_{k-1})\right)$$

[7]This definition of Υ is an approximate Newton step, rather than an exact one. 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 econometric results.

and update the auxiliary parameters:

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

- *Step 3: Increment k and repeat Step 2 until desired value of k is reached or until convergence.*

We claimed in the introduction that the k-EPL algorithm gives a sequence of asymptotically efficient estimators that converge in large samples. We now state this result formally in the following Theorem.

Theorem 1. *(Asymptotic Properties of k-EPL) Under Assumption 1, the k-EPL estimates computed with Algorithm 1 satisfy the following for any $k \geq 1$:*

1. *(Consistency) $\hat{\gamma}_k = (\hat{\theta}_k, \hat{Y}_k)$ is a strongly consistent estimator of (θ^*, Y^*).*
2. *(Asymptotic Equivalence to MLE) $\sqrt{N}(\hat{\theta}_k - \theta^*) \xrightarrow{d} \mathcal{N}(0, \Omega^{*-1}_{\theta\theta})$, where $\Omega^*_{\theta\theta}$ is the information matrix evaluated at θ^*.*
3. *(Large Sample Convergence) There exists a neighborhood of $\gamma^* = (\theta^*, Y^*)$, $\mathcal{B}^*$, such that $\lim_{k\to\infty} \hat{\gamma}_k = \hat{\gamma}_{MLE}$ almost surely for any $\hat{\gamma}_0 \in \mathcal{B}^*$.*

The results of Theorem 1 are similar to some results in the literature on iteratively estimating single-agent dynamic discrete choice models and depend heavily on the "zero Jacobian" property (Result 3 of Lemma 1). The zero Jacobian property ensures that $\hat{\gamma}_k = (\hat{\theta}_k, \hat{Y}_k)$ is asymptotically orthogonal to $\hat{\gamma}_{k-1}$. This means that using $\hat{\gamma}_{k-1}$ is asymptotically equivalent to using $\gamma^* = (\theta^*, Y^*)$ at each step. Intuitively, an EPL step is very similar to a Newton step on the full maximum likelihood problem. This drives the consistency (Result 1) and asymptotic equivalence to MLE (Result 2) of each step. Result 3 in the theorem is similar to Kasahara and Shimotsu (2012, Proposition 1) and shows convergence in large samples.

While our large sample results are encouraging, there is also the question of finite sample performance. We obtain favorable local results, as detailed in the next theorem. The only additional requirement is that the Jacobian of the equality constraints, G, with respect to Y is nonsingular at the MLE estimate.

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

Theorem 2. *(Finite Sample Properties) Under Assumptions 1 and 2,*

1. *The MLE is a fixed point of the EPL iterations: if $\hat{\gamma}_{k-1} = \hat{\gamma}_{MLE}$, then $\hat{\gamma}_k = \hat{\gamma}_{MLE}$.*

2. *For all* $k \geq 1$,

$$\hat{\gamma}_k - \hat{\gamma}_{MLE} = O_p(N^{-1/2}||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}|| + ||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}||^2).$$

3. *W.p.a. 1 as* $N \to \infty$*, for any* $\varepsilon > 0$ *there exists some neighborhood of* $\hat{\gamma}_{MLE}$*,* $\mathcal{B}$*, such that the EPL iterations define a contraction mapping on* $\mathcal{B}$ *with Lipschitz constant,* $L < \varepsilon$*.*

The first result of Theorem 2 establishes that the MLE is a fixed point of the EPL iterations in a finite sample, similar to Aguirregabiria and Mira (2002, Proposition 3) for single-agent NPL. The second result gives a theoretical explanation of why we should expect iteration to yield improvements in finite samples. This result is analogous Proposition 2 of Kasahara and Shimotsu (2008), but their result was only for NPL in the single-agent case. Following the same logic, suppose the initial estimates are such that $\hat{\gamma}_0 - \gamma^* = O_p(N^{-b})$ for $b \in (1/4, 1/2]$, so that $||\hat{\gamma}_0 - \hat{\gamma}_{MLE}|| = O_p(N^{-b})$.[8] Repeated substitution gives $||\hat{\gamma}_k - \hat{\gamma}_{MLE}|| = O_p(N^{-(k-1)/2-2b})$. In particular, in the case where the state space is finite and frequency estimates are used, $b = 1/2$ and $||\hat{\gamma}_k - \hat{\gamma}_{MLE}|| = O_p(N^{-(k+1)/2})$, noting that $N^{-(k+1)/2} \to 0$ as $k \to \infty$ for $N > 1$. Our own Monte-Carlo simulations in Section 4 exhibit such improvements.

The third result in Theorem 2 allows us to consider EPL iterations as a computationally attractive algorithm for computing the MLE. It establishes that we can expect the EPL iterations to be a local contraction around the MLE with a very fast convergence rate. For the population objective function, the convergence rate is quadratic. However, we only have samples, so we should expect the convergence rate to be linear with a small Lipschitz constant, implying that we'll need only a few iterations to achieve convergence. So, we can use EPL iterations to compute the MLE even when a consistent $\hat{\gamma}_0$ is unavailable. We can simply use multiple starting values, iterate to convergence, and use the converged estimate that provides the highest log-likelihood. We illustrate this usage of EPL in our Monte Carlo experiments in Section 4.

Aside from EPL, there are two potential alternative algorithms for computing the MLE: the nested fixed-point (NFXP) algorithm á la Rust (1987) and the MPEC approach proposed by Su and Judd (2012) and extended to dynamic games by Egesdal et al. (2015). The NFXP algorithm searches over θ in an outer loop and finds $Y(\theta)$ such that $G(\theta, Y) = 0$ in an inner loop. MPEC leverages modern optimization packages to search over θ and Y simultaneously, only imposing that $G(\theta, Y) = 0$ at the solution. The algorithm of choice will depend on the structure of the model.

[8] For $b \in (1/4, 1/2]$, we have $\hat{\gamma}_0 - \hat{\gamma}_{MLE} = \hat{\gamma}_0 - \gamma^* - (\hat{\gamma}_{MLE} - \gamma^*) = O_p(N^{-b}) + O_p(N^{-1/2}) = O_p(N^{-b})$.

While this section discusses the EPL algorithm in the context of a general constrained maximum-likelihood problem, we are ultimately focused on estimating dynamic discrete choice games of incomplete information. As discussed in the introduction, NFXP is computationally unattractive—or even infeasible—in such games.[9] MPEC, however, remains feasible and performs well, as demonstrated by Egesdal et al. (2015). The key difference here between MPEC and EPL is that EPL will be able to heavily exploit the structure of the problem.[10] In Section 3, we show that—much like NPL in single-agent models—common modeling assumptions lead to EPL iterations that break down into two easily-computed parts: solving a (potentially large-scale) linear system, followed by solving an unconstrained, globally concave maximization problem. Neither of these operations require sophisticated optimization software; and repeating them just a few times may ultimately be more computationally attractive than using MPEC to simultaneously solve for all variables in a non-concave, large-scale, constrained maximization problem.

We have already mentioned that the choice of Υ 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 3. *(Asymptotically Equivalent Definitions of Υ) The results of Theorems 1 and 2 hold when Υ is defined as any of the following, where $\hat{\gamma}_{k-1} = (\hat{\theta}_{k-1}, \hat{Y}_{k-1})$:*

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

The first definition of Υ in the theorem is the one we have worked with so far, defined earlier in (2). 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.[11] We will show later on that this definition can be used in single-agent dynamic discrete choice models. The third definition, which is an exact Newton step from (1), is likely the least useful because it requires inverting $G_Y(\theta, \hat{Y}_{k-1})$ at multiple values of θ, which can be computationally burdensome and also will

[9] We note that NFXP still performs well in single-agent dynamic models. See Doraszelski and Judd (2012) and Arcidiacono, Bayer, Blevins, and Ellickson (2016) for details on the computational burden of computing equilibria in discrete-time dynamic discrete games.

[10] Egesdal et al. (2015) exploit sparsity patterns in their MPEC implementation but do not further exploit other features of the problem structure.

[11] 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, so definition 2 includes definition 1.

introduce additional nonlinearities in the objective function for optimization.[12] However, we include it for completeness. For all of the definitions of Υ in the theorem, the results of Lemma 1 hold when all appropriate terms are replaced with (θ^*, Y^*) or $(\hat{\theta}_{\text{MLE}}, \hat{Y}_{\text{MLE}})$. So, the proof techniques from Theorems 1 and 2 can be used to prove Theorem 3.

2.2 Example: A Static Game with an Unstable Equilibrium

As a simple illustration of the performance of k-EPL, we consider estimating the static game of incomplete information from Pesendorfer and Schmidt-Dengler (2010). This example is particularly interesting because it was constructed as an example where converged NPL ($k \to \infty$) is inconsistent. We discuss only some relevant details of the model and refer the reader to Pesendorfer and Schmidt-Dengler (2010) 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 $\Pr(a_j = 1 \mid \theta, P_{-j}) = 1 - F_\alpha(-\theta P_{-j})$, where $0 < \alpha < 1$. F_α is an approximate uniform distribution with $F_\alpha(x) = x$ for $x \in [\alpha, 1-\alpha)$ and a more complicated form for $x \in \mathbb{R} \setminus [\alpha, 1-\alpha)$ to guarantee that it is a proper distribution function with full support. The probability mass in the uniform region can be made arbitrarily close to 1 by taking $\alpha \to 0$. Given a value of θ, the model has three equilibria for α sufficiently close to zero. The equilibrium generating the data is described by the following fixed point equation:

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

or more compactly, $P = \Psi(\theta, P)$. This linear system has a unique solution if and only if $\theta \neq -1$, and the solution is $P_1 = P_2 = \frac{1}{1-\theta}$. The k-NPL iterates are defined by

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

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

We note that $q_i(\theta, \hat{P}_{k-1})$ is the log of a linear function of θ, a point we will return to later.

[12] More precisely, it only requires solving 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.

Pesendorfer and Schmidt-Dengler (2010) consider the case where $\theta^* = -2$, implying $P_1^* = P_2^* = \frac{1}{3}$. They show that as $N \to \infty$, $\hat{\theta}_\infty^{\text{NPL}} \xrightarrow{p} -1$. 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 induce convergence in more general settings.

Consider, instead, estimating θ^* with k-EPL. We still want q_i to be the natural log of a function that is linear in θ, and this will require a change of variables 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 θ is a scalar, $G(\theta, Y)$ is linear in θ and Y separately (holding the other fixed) but not jointly. 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 θ only influences q_i through $Y(\theta) = v(\theta)$. This modification is made without loss of generality in full MLE subject to the equilibrium constraint, so it is also valid here.

We define $\Upsilon(\theta, \hat{\gamma}_{k-1}) = \hat{Y}_{k-1} - \nabla_Y G(\hat{\theta}_{k-1}, \hat{Y}_{k-1})^{-1} G(\theta, \hat{Y}_{k-1})$, the "default" Definition 1 in Theorem 3. Because $G(\theta, Y)$ is linear in θ, Υ is also linear in θ. This means that $\hat{q}_i(\theta, \Upsilon(\theta, \hat{\gamma}_{k-1})) = \hat{q}_i(\Upsilon(\theta, \hat{\gamma}_{k-1}))$ is the log of a linear function of θ, 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 Υ linear in θ, but q_i would then depend on $\theta\Upsilon(\theta, Y)$, which is non-linear in θ.

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 $\hat{P}_{1,0}$ 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 5,000 observations. We estimate the model using MLE, ∞-EPL, and ∞-NPL.[13] The results are summarized in Table 1. The MLE and ∞-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 unsurprising: with so many observations and only two players/actions, we get *very* precise initial estimates, so iteration converges very quickly. 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.

Table 1: Pesendorfer and Schmidt-Dengler (2010) Monte Carlo Results

Estimator	Mean	MSE
MLE	-2.0017	0.0017
∞-EPL	-2.0014	0.0017
∞-NPL	-1.0342	0.9651

On the other hand, ∞-NPL performs poorly, as is to be expected since this model was constructed as an example where ∞-NPL is inconsistent. The estimate has a mean of -1.0342 and MSE of 0.9651. Almost all of the MSE is due to the asymptotic bias, so the estimate is reliably converging to the wrong number.[14]

[13] By ∞-EPL and ∞-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.

[14] There were 17 samples for which NPL converged in 3 or fewer iterations. The mean and MSE for these

In addition to demonstrating the good performance of k-EPL, this example also introduces the change of variables that will be used in dynamic games. Notice that v_j is essentially a choice-specific value 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.

2.3 Single-Agent Dynamic Discrete Choice

We'll now show that k-NPL in a single-agent dynamic discrete choice model is a special case of k-EPL. Here, we can work directly in probability space using the fixed point condition,

$$P = \Psi(\theta, P),$$

with $\Psi(\theta, P)$ defined as in Aguirregabiria and Mira (2002). Define $Y \equiv P$, so that $G(\theta, Y) = Y - \Psi(\theta, Y)$ and

$$\nabla_Y G(\theta, Y) = I - \nabla_Y \Psi(\theta, Y).$$

Now, consider implementing k-EPL with Definition 2 from Theorem 3: $\Upsilon(\theta, \hat{\gamma}_{k-1}) = \hat{Y}_{k-1} - Z(\hat{\theta}_{k-1}, \hat{Y}_{k-1})^{-1} G(\theta, \hat{Y}_{k-1})$, where $Z(\theta, Y_\theta) = \nabla_Y G(\theta, Y_\theta)$ for all θ. Proposition 2 from Aguirregabiria and Mira (2002) shows that $\nabla_P \Psi(\theta, P_\theta) = 0$. Thus, $\nabla_Y G(\theta, Y_\theta) = I$ for all θ since $Y \equiv P$. So, we set $Z(\theta, Y) = I$ and obtain

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

Finally, we have that $q_i(\theta, \Upsilon(\theta, \hat{\gamma}_{k-1})) = \ln \Psi(\theta, \hat{Y}_{k-1})(w_i) = \ln \Psi(\theta, \hat{P}_{k-1})(w_i)$ because because $Y \equiv P$. So, this implementation of k-EPL is identical to k-NPL.

This equivalence of k-NPL to k-EPL in single agent models 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, Proposition 1(c)) show that their policy iterations are equivalent to Newton iterations on the value function in single-agent models. Since k-EPL is built around Newton iterations, such an equivalence is again suggestive of the relationship shown in this section.

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. Aguirregabiria and Marcoux (2019) explain why the estimates converge to "good" values in some samples even though the equilibrium generating the data is unstable.

3 Dynamic Discrete Games of Incomplete Information

Here we describe a canonical stationary dynamic discrete game of complete information in the style of Aguirregabiria and Mira (2007) and Pesendorfer and Schmidt-Dengler (2008), although our notation is closer to that of Bugni and Bunting (2019). Time is discrete, indexed by $t \in \{1, 2, \ldots, \infty\}$. In a given market, there are J firms operating, indexed by $j \in \mathcal{J} = \{1, 2, \ldots, |\mathcal{J}|\}$. Given a vector of state variables observable to all agents and the econometrician, x_t, and its own private information ε_t^j, each firm chooses an action, $a_t^j \in \mathcal{A} = \{0, 1, 2, \ldots, |\mathcal{A}|\}$. Action zero is the outside option when applicable. All players choose their actions simultaneously.

Agents have flow utilities (profits), $\tilde{u}_{at}^j(\theta_u, x_t, a_t^{-j}, \varepsilon_{at}^j)$, where a_t^{-j} are the actions of the other players. States transition according to $p(x_{t+1}, \varepsilon_{t+1} \mid a_t, x_t, \varepsilon_t; \theta_f)$, and the discount factor is $\beta \in (0, 1)$. Agents choose actions to maximize expected discounted utility,

$$E\left\{\sum_{s=0}^{\infty} \beta^{s-t} \tilde{u}_{as}^j(\theta, x_s, a_s^{-j}, \varepsilon_s^j) \,\middle|\, x_t, \varepsilon_t^j\right\}.$$

The primary parameter of interest is $\theta = (\theta_u, \theta_f)$. Furthermore, we impose the following standard assumptions on the primitives.

Assumption 3. *(Additive Separability)* $\tilde{u}_{at}^j(\theta_u, x_t, a_t^{-j}, \varepsilon_{at}^j) = \bar{u}(\theta_u, x_t, a_t^{-j}) + \varepsilon_{at}^j$.

Assumption 4. *(Conditional Independence)* $p(x_{t+1}, \varepsilon_{t+1} \mid a_t, x_t, \varepsilon_t; \theta_f) = g(\varepsilon_{t+1}) f(x_{t+1} \mid a_t, x_t; \theta_f)$*, where* $g(\varepsilon_{t+1})$ *has full support and is thrice-continuously differentiable.*

Assumption 5. *(Independent Private Values) Private values are independently distributed across players.*

Assumption 6. *(Finite Observed State Space)* $x_t \in \mathcal{X} = \{1, 2, \ldots, |\mathcal{X}|\}$.

Assumptions 3–6 here correspond to Assumptions 1–4 in Aguirregabiria and Mira (2007), although our conditional independence assumption also includes an additional full support assumption, along with some differentiability assumptions.

The operative equilibrium concept here will be that of a Markov-perfect Nash equilibrium. We will consider stationary equilibria only, so from here we drop the time subscript. Because moves are simultaneous, the actions of player j do not depend directly on $a^{-j} \in \mathcal{A}^{J-1}$, but rather on $P^{-j} \in \Delta^{J-1}$, where P^{-j} is its belief about the other players' probability of playing the corresponding actions and Δ is the unit simplex in $\mathbb{R}^{|\mathcal{A}|-1}$. So, from here on out we will

work with the following utility function:

$$u_a^j(\theta_u, x, P^{-j}) = \sum_{a^{-j} \in \mathcal{A}^{J-1}} P_{a^{-j}}^{-j} \bar{u}(\theta_u, x, a^{-j}).$$

Now consider the vector of player j's (expected) choice-specific value functions, $v^j \in \mathbb{R}^{|\mathcal{X}|\times|\mathcal{A}|}$, 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 variables as $v = \Phi(\theta, v)$.

Lemma 2. *(Representation Lemma) Choice-specific value functions characterize an equilibrium for θ 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 2 describes our change of variables that will be used to implement Algorithm 1 in estimation. This is in contrast to the representation lemma from Aguirregabiria and 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

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

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 θ 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 θ and F_a^j does not depend on θ, the computational burden of the pseudo-likelihood maximization in each k-NPL iteration reduces

to that of a linear-in-parameters static model.[15] 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. The key here is the linearity of $\Phi(\theta, v)$ in θ under these conditions. By defining $Y = v$ and $G(\theta, v) = v - \Phi(\theta, v)$, we get linearity of $G(\theta, v)$ in θ. And finally, we have $\Upsilon(\theta, \hat{\gamma}_{k-1}) = \hat{Y}_{k-1} - \nabla G(\hat{\theta}_{k-1}, \hat{v}_{k-1})^{-1} G(\theta, \hat{v}_{k-1})$, which is also linear in θ. The function $\Upsilon(\theta, \hat{\gamma}_{k-1})$ determines the choice-specific value functions used to compute the pseudo-likelihood, so these will be linear in θ. Furthermore, only choice-specific value functions are needed to compute choice probabilities, so that $Q(\theta, Y) = Q(v)$. Thus, each step of our estimator uses $Q(\Upsilon(\theta, \hat{\gamma}_{k-1}))$, so that θ only enters the objective function indirectly through $\Upsilon(\theta, \hat{\gamma}_{k-1})$.

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 $\Upsilon(\theta, \hat{\gamma}_{k-1})$. If we use the fixed point in probability space to define Y and $G(\theta, Y)$, then we will encounter problems if $\Upsilon(\theta, \hat{\gamma}_{k-1})$ does not map into the interior of the unit 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 $\mathbb{R}^{|\mathcal{J}| \times |\mathcal{X}| \times |\mathcal{A}|}$.

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

We now turn our attention a series of Monte Carlo experiments where we estimate the model from Pesendorfer and Schmidt-Dengler (2008). There are two firms indexed by $j \in \{1, 2\}$ who choose an action in each market i, denoted $a_{ij} \in \{0, 1\}$, where 1 is entry and 0 is exit. The observed state variable $x_i = (x_{i1}, x_{i2}) \in \{(0,0), (0,1), (1,0), (1,1)\}$ represents the incumbency status of firms 1 and 2, respectively. Flow utilities are period profits:

$$\tilde{u}^j(x_i, a_{ij} = 1) = \theta^*_{\text{M}} + \theta^*_{\text{C}} a_{i,-j} + \theta^*_{\text{EC}}(1 - x_{ij}) + \varepsilon_{j1},$$

$$\tilde{u}^j(x_i, a_{ij} = 0) = \theta^*_{\text{SV}} x_{ij} + \varepsilon_{j0},$$

where θ^*_{EC} represents the entry cost, θ^*_{SV} is the scrap value, θ^*_{M} is the monopoly profit, and θ^*_{C} is the effect of competition on profit. The discount factor is $\beta^* \in (0, 1)$. The data are generated using the parameter values $(\theta^*_{\text{M}}, \theta^*_{\text{C}}, \theta^*_{\text{EC}}, \theta^*_{\text{SV}}, \beta^*) = (1.2, -2.4, -0.2, 0.1, 0.9)$.

[15] Usually, we have $\theta = (\theta_u, \theta_f)$, $F^j_a(\theta_f, P^{-j})$, and linear-in-parameters $u^j_a(\theta_u, P^{-j})$. We then implement a "two-step" estimator where θ^*_f is estimated in a first stage and θ^*_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.

The private shocks have distribution $\varepsilon_{ja} \sim \mathrm{N}(0, 0.5)$. We note that this is a slightly different parameterization of the model than the one used by Pesendorfer and Schmidt-Dengler (2008), but it is straightforward to show that the resulting flow utilities and hence the equilibria are the same.

There are multiple equilibria in the game, and we generate data from equilibria (i), (ii), and (iii) from Pesendorfer and Schmidt-Dengler (2008). The NPL mapping is unstable for two of the three equilibria, but the EPL mapping, having spectral radius zero, is stable for all three equilibria. Specifically, equilibrium (i) is asymmetric and NPL-stable, equilibrium (ii) is asymmetric and NPL-unstable, and equilibrium (iii) is symmetric and NPL-unstable. In their Monte-Carlo simulations, Pesendorfer and Schmidt-Dengler (2008) find that k-NPL performs well for equilibrium (i) but becomes severely biased for equilibria (ii) and (iii) as k grows, which is in line with the convergence analysis of Kasahara and Shimotsu (2012).[16] However, we expect that EPL will perform well for all three equilibria.

We estimate $(\theta^*_{\mathrm{M}}, \theta^*_{\mathrm{C}}, \theta^*_{\mathrm{EC}})$ and assume the other parameter values are known and held fixed at $\theta^*_{SV} = 0.1$ and $\beta^* = 0.9$. We present results for estimation using NPL and EPL. The initial estimates of the conditional choice probabilities are sample frequencies, $\hat{P}^j_{x,a}$. We generate the data by first taking $N \in \{250, 1000\}$ i.i.d. draws from the stationary distribution of the observed state, x, for each equilibrium. One interpretation of this sampling procedure is that each of the draws from the stationary distribution of x represents an independent market. For each of these N draws we then sample actions for each player using the equilibrium choice probabilities. We carry out 1000 replications for each sample size. For ∞-NPL and ∞-EPL we terminate the algorithm when $|\hat{\theta}_k - \hat{\theta}_{k-1}| < 10^{-6}$ or after 100 iterations. Computational times reported are minutes of "wall clock" time required to carry out the full set of replications.[17]

In order to implement EPL in this context, we use the fixed-point constraint in choice-specific value function space, defined in Section 3. Computing $\nabla_v G(\hat{\theta}_{k-1}, \hat{v}_{k-1})$ for $k = 1$ requires initial estimates $(\hat{\theta}_0, \hat{v}_0)$. We use $\hat{\theta}_0 = \hat{\theta}_{\text{1-NPL}}$, the estimate from 1-NPL, which is similar to the way Pesendorfer and Schmidt-Dengler (2008) use $\hat{\theta}_{\text{1-NPL}}$ to obtain an estimate of the efficient weighting matrix used in their minimum-distance estimator. For each player, we then set $\hat{v}^j_{0,x,a} = v^j_a = u^j_a(\hat{\theta}_0, \hat{P}^{-j}) + \beta F^j_a(\hat{P}^{-j})\Gamma^j(\hat{\theta}_0, \hat{P})$, where

$$\Gamma^j(\theta, P) = \big(I - \beta F^j(\theta, P)\big)^{-1} \sum_a P^j(a) * \big(u^j_a(\theta, P^{-j}) + e_a(P^j)\big)$$

[16] Pesendorfer and Schmidt-Dengler (2008) use the terminology "k-PML" for k-NPL and iterate until $k = 20$.

[17] Experiments were carried out using MATLAB R2018a on a 2017 iMac Pro in parallel using 18 Intel Xeon 2.3 GHz cores.

maps (θ, P) into an *ex-ante* value function for player j, as in Aguirregabiria and Mira (2007, p. 10).

4.1 Results for NPL-Stable Equilibrium (i)

Table 2 shows results for equilibrium (i), for which NPL is stable and consistent. We consider both the one-step ($k = 1$) NPL and EPL estimators as well as the converged estimators ($k = \infty$). For $k = \infty$, we report the total estimation time across all datasets, as well as the median and interquartile range (IQR) of the number of iterations. For the large sample experiments, we obtained convergence in fewer than 100 iterations for all algorithms in almost all datasets, with ∞-NPL and ∞-EPL failing to converge in only 5 and 1 out of 1000 datasets, respectively. Convergence rates were somewhat lower, especially for ∞-NPL, with the smaller sample size. Our reported results include all datasets, including those where we obtain non-convergence.

Table 2: Monte Carlo Results for Pesendorfer and Schmidt-Dengler (2008) NPL-Stable Equilibrium (i)

Obs.	Parameter	Statistic	1-NPL	1-EPL	∞-NPL	∞-EPL
$N = 250$	$\theta_{\text{M}} = 1.2$	Mean Bias	-0.0579	-0.0277	-0.0158	0.0277
		MSE	0.0461	0.0376	0.0368	0.0312
	$\theta_{\text{C}} = -2.4$	Mean Bias	0.1120	0.0425	0.0294	-0.0482
		MSE	0.1642	0.1061	0.0585	0.0512
	$\theta_{\text{EC}} = -0.2$	Mean Bias	-0.0393	-0.0205	-0.0270	-0.0039
		MSE	0.0494	0.0338	0.0116	0.0045
	Converged	%			92.6%	97.5%
	Iterations	Median			70	8
		IQR			28	2
	Time (min.)	Total			0.4481	0.1047
$N = 1000$	$\theta_{\text{M}} = 1.2$	Mean Bias	-0.0165	-0.0050	-0.0044	0.0033
		MSE	0.0116	0.0107	0.0083	0.0059
	$\theta_{\text{C}} = -2.4$	Mean Bias	0.0340	0.0119	0.0076	-0.0052
		MSE	0.0423	0.0320	0.0106	0.0076
	$\theta_{\text{EC}} = -0.2$	Mean Bias	-0.0127	-0.0061	-0.0059	-0.0012
		MSE	0.0123	0.0086	0.0018	0.0008
	Converged	%			99.5%	99.9%
	Iterations	Median			70	6
		IQR			19	1
	Time (min.)	Total			0.5847	0.0743

Comparing 1-NPL to 1-EPL in Table 2, we see that 1-EPL has lower mean bias and MSE for all three parameters of interest. However, both of these are outperformed by estimators

iterated to convergence, illustrating the gains from such iterations in finite samples. For the larger sample size ∞-EPL has the lowest bias, MSE, number of iterations, and computation time. Even for this equilibrium where NPL is expected to perform well, the efficiency of EPL yields improvements. Time per iteration is similar for NPL and EPL by design, so the lower computational times are driven by the significant reduction in the number of iterations to convergence.[18] In this case for the smaller sample size, the results are more mixed. The mean bias of EPL is higher (in absolute value) for two parameter values, but the MSE is lower for all three. However, as before, convergence is much faster.

4.2 Results for NPL-Unstable Equilibrium (ii)

Table 3: Monte Carlo Results for Pesendorfer and Schmidt-Dengler (2008) Equilibrium (ii)

Obs.	Parameter	Statistic	1-NPL	1-EPL	∞-NPL	∞-EPL
$N = 250$	$\theta_{\text{M}} = 1.2$	Mean Bias	-0.1322	-0.1461	-0.2099	-0.0309
		MSE	0.0902	0.0988	0.0622	0.0740
	$\theta_{\text{C}} = -2.4$	Mean Bias	0.2793	0.2617	0.6719	0.0717
		MSE	0.4643	0.5121	0.4804	0.4106
	$\theta_{\text{EC}} = -0.2$	Mean Bias	-0.0777	-0.0764	-0.3110	-0.0441
		MSE	0.1058	0.1270	0.1117	0.1076
	Converged	%			96.1%	100%
	Iterations	Median			34	9
		IQR			10	3
	Time (min.)	Total			0.3011	0.0785
$N = 1000$	$\theta_{\text{M}} = 1.2$	Mean Bias	-0.0432	-0.0385	-0.2093	-0.0013
		MSE	0.0210	0.0205	0.0480	0.0155
	$\theta_{\text{C}} = -2.4$	Mean Bias	0.0952	0.0612	0.6636	0.0047
		MSE	0.1162	0.1078	0.4459	0.0829
	$\theta_{\text{EC}} = -0.2$	Mean Bias	-0.0269	-0.0122	-0.2983	-0.0039
		MSE	0.0283	0.0279	0.0923	0.0222
	Converged	%			99.7%	100%
	Iterations	Median			32	7
		IQR			6	1
	Time (min.)	Total			0.3585	0.0946

Turning to equilibrium (ii), for which the NPL fixed point is unstable, we have a very different picture. The results are presented in Table 3. For ∞-NPL there is substantial bias in all parameters and seemingly little variation around those biased values. For example, there is attenuation bias in the competitive effect θ_{C}, making it seem less negative. This bias

[18] Each iteration reduces to solving a linear system and then estimating a static binary probit model.

does not appear to decrease with the sample size. The bias for ∞-EPL is lower by an order of magnitude in all cases for $N = 250$ and by two orders of magnitude for $N = 1000$. Relative to 1-NPL and 1-EPL, iterating in the finite sample improves the estimates for ∞-EPL but worsens the estimates for ∞-NPL.

The results for equilibrium (iii) are qualitatively similar to equilibrium (ii) and are presented in Appendix B. Overall, the results illustrate the good performance of EPL and in particular ∞-EPL. We see that ∞-EPL is generally more efficient, is robust to unstable equilibria, and converges in fewer iterations than ∞-NPL, resulting in substantial time savings.

4.3 Effects of Noisy, Inconsistent Starting Values

Next, we consider robustness to starting values. Like convergence results for Newton's method, our convergence results are local. That is, the starting values (initial estimates) must be in a neighborhood of the maximum likelihood estimates to guarantee convergence. We do not claim, nor should we expect, global convergence results in models with multiple equilibria. This underscores the importance of good initial estimates, i.e., either consistent estimates or multiple starting values, or both. First we explore using a single, polluted version of the consistent estimates as the starting value.

Table 4: Monte Carlo Results for Pesendorfer and Schmidt-Dengler (2008) with Noisy Starting Values ($N = 250$)

		Equilibrium (i)		Equilibrium (ii)	
Parameter	Statistic	∞-NPL	∞-EPL	∞-NPL	∞-EPL
$\theta_{\text{M}} = 1.2$	Mean Bias	-0.0827	0.0266	-0.2116	-0.0855
	MSE	0.0648	0.0756	0.0630	0.1054
$\theta_{\text{C}} = -2.4$	Mean Bias	0.1286	-0.0431	0.6738	0.2055
	MSE	0.1282	0.2234	0.4832	0.5810
$\theta_{\text{EC}} = -0.2$	Mean Bias	-0.0672	0.0020	-0.3107	-0.0965
	MSE	0.0251	0.0073	0.1118	0.1275
Converged	%	89.2%	96.6%	99.3%	99.9%
Iterations	Median	72	9	34	10
	IQR	36	3	11	3
Time (min.)	Total	0.4024	0.1217	0.2533	0.0889

In Table 4 we used initial choice probabilities that were an equally weighted average of the consistent CCP estimates and Uniform$(0, 1)$ noise. We then re-computed each of the converged estimates — ∞-NPL and ∞-EPL — using these noisy starting values with the small sample size $N = 250$. For equilibrium (i), comparing with the consistent starting values from the top panel of Table 2, we see that the added noise increases the MSE values

and decreases convergence rates for both estimators, but the increase in bias is smallest for ∞-EPL. Furthermore, the convergence rate of ∞-EPL decreases less than the convergence rate for ∞-NPL.

For equilibrium (ii) we can compare with the consistent starting values from the top panel of Table 3. In this case the bias and MSE for ∞-NPL only changed slightly because the estimates were previously also biased. There is only a slight increase in bias and MSE as a result of the noisy starting values, but the results are largely the same as before. The convergence percentage for ∞-NPL actually increased with the added noise, but the iterations still converge to inconsistent estimates. On the other hand, the bias and MSE values for ∞-EPL increased—especially for θ_C—while the convergence percentage only decreased from 100% to 99.9%.

Overall, while good starting values are important, these results show that EPL is also somewhat robust starting values with fairly severe estimation errors. Note that we do not actually recommend using only a single starting value if estimation accuracy is a concern. With that in mind, in the next section we consider moving away from consistent starting values entirely.

4.4 EPL as a Computational Procedure (Without Consistent Starting Values)

Rather than rely solely on a single consistent estimate, one can use EPL as a computational procedure to compute the MLE using multiple starting values (possibly with a consistent estimate among them). A similar procedure was suggested by Aguirregabiria and Mira (2007) to compute the NPL estimator by attempting to use the NPL algorithm, with multiple starting values, to compute all NPL fixed points, and taking the estimate that maximizes the likelihood. However, for datasets generated by equilibria for which the NPL mapping is unstable, the initial guess may need to be exactly correct to reach those fixed points.[19] EPL, however, is stable and has a faster rate of convergence, with the maximum likelihood estimator as a fixed point. So, in this section we consider using this approach with EPL to compute the MLE.

Using the same model as before and focusing on equilibria (i) and (ii), we proceed in the following way for each of 1000 replications. First, we generated five completely random starting values for the choice probabilities P. For each, we compute and store the corresponding 1-NPL estimate for θ. Then we compute the ∞-NPL and ∞-EPL estimates for each start-

[19] Aguirregabiria and Marcoux (2019) show that this issue can even arise when the data are generated from a stable equilibrium.

Table 5: Monte Carlo Results for Pesendorfer and Schmidt-Dengler (2008) Without Consistent Starting Values

			Equilibrium (i)		Equilibrium (ii)	
Obs.	Parameter	Statistic	∞-NPL	∞-EPL	∞-NPL	∞-EPL
$N = 250$	$\theta_{\text{M}} = 1.2$	Mean Bias	-0.0127	0.0295	-0.2035	-0.0083
		MSE	0.0359	0.0325	0.0601	0.0812
	$\theta_{\text{C}} = -2.4$	Mean Bias	0.0252	-0.0516	0.6634	0.0088
		MSE	0.0566	0.0552	0.4699	0.4649
	$\theta_{\text{EC}} = -0.2$	Mean Bias	-0.0255	-0.0034	-0.3078	-0.0122
		MSE	0.0113	0.0046	0.1100	0.1228
	Converged	%	89.8%	97.2%	95.2%	100%
	Iterations	Median	347	53	176	53
		IQR	117	18	46	10
	Time (min.)	Total	1.9699	0.5606	1.2573	0.4207
$N = 1000$	$\theta_{\text{M}} = 1.2$	Mean Bias	-0.0044	0.0019	-0.2002	-0.0100
		MSE	0.0083	0.0057	0.0444	0.0173
	$\theta_{\text{C}} = -2.4$	Mean Bias	0.0076	-0.0027	0.6517	0.0264
		MSE	0.0106	0.0065	0.4303	0.0962
	$\theta_{\text{EC}} = -0.2$	Mean Bias	-0.0059	-0.0016	-0.2915	-0.0146
		MSE	0.0018	0.0008	0.0884	0.0255
	Converged	%	95.1%	98.7%	99.6%	100%
	Iterations	Median	362	50	161	52
		IQR	83	12	18	8
	Time (min.)	Total	2.8453	0.6559	1.5548	0.5956

ing value.[20] Finally, for both ∞-NPL and ∞-EPL we select from among the five estimates the one that maximizes the likelihood function. The results in Table 5 are calculated using the best estimates for each of the 1000 replications. Reported iteration counts, convergence percentages, and computational times include all 5 starting values over all replications.

Overall, the comparisons between EPL and NPL are qualitatively similar to the case before with initial consistent estimates. For the NPL-stable equilibrium (i), the small sample results for bias are mixed, but EPL is faster, converges more often, and has smaller MSE. In the large sample, EPL always has lower bias and MSE in addition to being more stable and computationally lighter. Also as before, EPL is robust to the NPL-unstable equilibrium. These results show that converged EPL can be used as a computational procedure to compute the MLE without initial consistent estimates, by using multiple starting values and choosing the best estimates. This is true even for equilibrium data generating processes where using the same procedure with converged NPL would yield inconsistent estimates.

[20] As in the previous simulations, we allow for up to 100 iterations per starting value but terminate early if convergence is achieved.

5 Conclusion

We proposed an iterative efficient pseudo-likelihood (EPL) estimation algorithm that extends the attractive econometric and computational properties of the single-agent NPL algorithm to games. The nice econometric properties arise because EPL uses (quasi-)Newton steps on the fixed point constraint at each iteration. As a result, EPL is stable for all regular Markov perfect equilibria, each EPL iteration is asymptotically equivalent to the MLE, and the iterates quickly converge to the finite-sample MLE with high probability. Computational advantages follow from defining the equilibrium conditions with choice-specific value functions. Standard modeling assumptions then reduce each EPL iteration to solving a linear system, followed by solving a globally concave, unconstrained maximization problem. Our Monte-Carlo simulations show that EPL performs favorably in finite samples and is robust to data-generating processes where standard NPL encounters serious problems.

One limitation of our analysis is that we did not consider time-invariant unobserved heterogeneity in estimating dynamic discrete games. Given the similarities between the procedures, it seems likely that one could modify the EPL algorithm to include such heterogeneity using techniques similar to Aguirregabiria and Mira (2007, Section 3.5) or Arcidiacono and Miller (2011). We leave such an extension to future work.

References

Aguirregabiria, V. and M. Marcoux (2019). 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*, 1519–1543.

Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica 75*, 1–53.

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. and R. Miller (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica 7*, 1823–2868.

Bajari, P., C. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica 75*(5), 1331–1370.

Bugni, F. and J. Bunting (2019). On the iterated estimation of dynamic discrete choice games. Working Paper. arXiv:1802.06665v3.

Dearing, A. (2019). Pseudo-value functions and closed-form ccp estimation of dynamic discrete choice models. Working Paper.

Doraszelski, U. and J. F. Escobar (2010). A theory of regular Markov perfect equilibria in dynamic stochastic games: Genericity, stability, and purification. *Theoretical Economics 5*, 369–402.

Doraszelski, U. and K. L. Judd (2012). Avoiding the curse of dimensionality in dynamic stochastic games. *Quantitative Economics 3*, 53–93.

Egesdal, 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 2. Cambridge University Press.

Hotz, V. and R. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies 60*, 497–529.

Hotz, V., R. Miller, S. Sanders, and J. Smith (1994). A simulation estimator for dynamic models of discrete choice. *Review of Economic Studies 61*, 265–289.

Kasahara, H. and K. Shimotsu (2008). Pseudo-likelihood estimation and bootstrap inference for structural discrete markov decision models. *Journal of Econometrics 146*, 92–106.

Kasahara, H. and K. Shimotsu (2012). Sequential estimation of structural models with a fixed point constraint. *Econometrica 80*(5), 2303–2319.

Kelley, C. T. (1995). *Iterative Methods for Linear and Nonlinear Equations*. Philadelphia: SIAM.

Newey, W. and D. McFadden (1994). Large sample estimation and hypothesis testing. In R. Engle and D. McFadden (Eds.), *Handbook of Econometrics*, Volume 4. Amsterdam: North-Holland.

Pakes, A. (1986). Patents as options: Some estimates of the value of holding european patent stocks. *Econometrica 54*, 1027–1057.

Pakes, A. and P. McGuire (1994). Computing markov-perfect nash equilibria: Numberical implications of a dynamic differentiated product model. *RAND Journal of Economics 25*(4), 555–589.

Pakes, A. and P. McGuire (2001). Stochastic approximation for dynamic models: Markov perfect equilibrium and the 'curse' of dimensionality. *Econometrica 69*(5), 1261–1281.

Pakes, A., M. Ostrovsky, and S. Berry (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *RAND Journal of Economics 38*(2), 373–399.

Pesendorfer, M. and P. Schmidt-Dengler (2008). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies 75*, 901–928.

Pesendorfer, M. and P. Schmidt-Dengler (2010). Sequential estimation of dynamic discrete games: A comment. *Econometrica 78*(2), 833–842.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica 55*, 999–1013.

Su, C.-L. and K. Judd (2012). Constrained optimization approaches to estimation of structural models. *Econometrica 80*(5), 2213–2230.

Tamer, E. (2003). Incomplete simultaneous discrete response model with multiple equilibria. *Review of Economic Studies 70*, 147–165.

A Proofs

Proofs are presented in order of appearance in the main text.

A.1 Proof of Lemma 1

Result 1 follows by using $G(\theta, Y_\theta) = 0$ in the expression for $\Upsilon(\theta, Y_\theta)$. For Results 2 and 3, first consider partial differentiation of $\Upsilon(\theta, Y)$ as defined in (1):

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

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

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

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

Results 2 and 3 then follow from $G(\theta, Y_\theta) = 0$. Result 4 is a widely-known property of Newton-Kantorovich iterations. See, for example, Theorem 5.1.2 of Kelley (1995).

A.2 Proof of Theorem 1

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.[21] 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 very similar to those used in the proofs of Proposition 2 of Kasahara and Shimotsu (2008) and Proposition 7 in the supplementary material for Kasahara and Shimotsu (2012). Throughout, we rely heavily on analysis similar to that from the proof of Lemma 1.

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

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

A.2.1 Result 1 (Strong consistency of $\hat{\theta}_k$ and $\hat{Y}_k$)

We have uniform continuity of $\tilde{Q}^*(\theta,\gamma)$ and that $\tilde{Q}_N(\theta,\gamma)$ 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 γ^*. Appealing to Lemma 24.1 of Gourieroux and Monfort (1995), these imply that $\tilde{Q}_N(\theta,\hat{\gamma}_{k-1})$ converges almost surely and uniformly in $\theta \in \Theta$ to $\tilde{Q}^*(\theta,\gamma^*)$. Then since θ^* uniquely maximizes $\tilde{Q}^*(\theta,\gamma^*)$ on Θ, $\hat{\theta}_k$ converges almost surely to θ^* (Gourieroux and Monfort, 1995, Property 24.2). Continuity of $\Upsilon_k(\theta,\gamma)$ and the Mann-Wald theorem then give almost sure convergence of $\hat{Y}_k$ to Y^*.

A.2.2 Result 2 (Asymptotic Distribution of $\hat{\theta}_k$ and $\hat{Y}_k$)

We will show that 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 the properties of Υ defined in (2), the chain rule, and the generalized information matrix equality (Newey and McFadden, 1994, p. 2163) we obtain the following population equalities:

$$\begin{aligned}\nabla_{\theta\theta'}\tilde{Q}^*(\theta^*,\gamma^*) &= -\,\Omega^*_{\theta\theta},\\ \nabla_{\theta\gamma'}\tilde{Q}^*(\theta^*,\gamma^*) &= 0.\end{aligned}$$

To establish these, first recall that $\tilde{Q}^*(\theta,\gamma) = \mathrm{E}[\ln f(w \mid \theta, \Upsilon(\theta,\gamma))]$ and let $\gamma = (\breve{\theta}, Y)$ denote the components of γ (to explicitly distinguish θ from $\breve{\theta}$). Then by the generalized information matrix equality we have

$$\begin{aligned}\nabla_{\theta\theta'}\tilde{Q}^*(\theta^*,\gamma^*) &= E[\nabla_{\theta\theta'}\ln f(w \mid \theta^*, \Upsilon(\theta^*,\gamma^*))]\\ &= -E[\nabla_\theta \ln f(w \mid \theta^*, \Upsilon(\theta^*,\gamma^*))\nabla_\theta \ln f(w \mid \theta^*, \Upsilon(\theta^*,\gamma^*))']\\ &= -\Omega^*_{\theta\theta}\end{aligned}$$

and

$$\begin{aligned}\nabla_{\theta\gamma'}\tilde{Q}^*(\theta^*,\gamma^*) &= E[\nabla_{\theta\gamma'}\ln f(w \mid \theta^*, \Upsilon(\theta^*,\gamma^*))]\\ &= -E[\nabla_\theta \ln f(w \mid \theta^*, \Upsilon(\theta^*,\gamma^*))\nabla_\gamma \ln f(w \mid \theta^*, \Upsilon(\theta^*,\gamma^*))']\\ &= 0.\end{aligned}$$

The last equality follows from the chain rule, noting that $\nabla_\gamma \ln f(w \mid \theta, \Upsilon(\theta,\gamma)) = \frac{1}{f(w|\theta,\Upsilon(\theta,\gamma))}\nabla_Y f(w \mid$

$\theta, \Upsilon(\theta,\gamma))\nabla_\gamma \Upsilon(\theta,\gamma)$ with

$$\nabla_\gamma \Upsilon(\theta,\gamma) = I - \nabla_Y G(\breve{\theta}, Y)^{-1} \nabla_Y G(\theta, Y) - \left[\begin{array}{c} \frac{\partial \nabla_Y G(\breve{\theta},Y)^{-1}}{\partial \breve{\theta}} \\ \frac{\partial \nabla_Y G(\breve{\theta},Y)^{-1}}{\partial Y} \end{array} \right] G(\theta, Y). \tag{3}$$

Since $\breve{\theta} = \theta = \theta^*$ and $G(\theta^*, Y^*) = 0$ at the true parameters, we have $\nabla_\gamma \ln f(w \mid \theta^*, \Upsilon(\theta^*, \gamma^*)) = 0$.

Turning to the sample objective function, a Taylor expansion of the first-order condition gives

$$0 = \nabla_\theta \tilde{Q}_N(\theta^*, \gamma^*) + \nabla_{\theta\theta'} \tilde{Q}_N(\bar{\theta}, \bar{\gamma})(\hat{\theta}_k - \theta^*) + \nabla_{\theta\gamma'} \tilde{Q}_N(\bar{\theta}, \bar{\gamma})(\hat{\gamma}_{k-1} - \gamma^*).$$

Solving and scaling then yields

$$\sqrt{N}(\hat{\theta}_k - \theta^*) = -\nabla_{\theta\theta'} \tilde{Q}_N(\bar{\theta}, \bar{\gamma})^{-1} \left[\sqrt{N} \nabla_\theta \tilde{Q}_N(\theta^*, \gamma^*) + \nabla_{\theta\gamma'} \tilde{Q}_N(\bar{\theta}, \bar{\gamma}) \sqrt{N}(\hat{\gamma}_{k-1} - \gamma^*) \right].$$

By consistency of $\hat{\gamma}_{k-1}$ and $\hat{\theta}_k$ and the Mann-Wald theorem we have $\nabla_{\theta\theta'} \tilde{Q}_N(\bar{\theta}, \bar{\gamma}) \xrightarrow{p} -\Omega^*_{\theta\theta}$ and by the central limit theorem, $\sqrt{N} \nabla_\theta \tilde{Q}_N(\theta^*, \gamma^*) \xrightarrow{d} \mathrm{N}(0, \Omega^*_{\theta\theta})$. For the last term in square brackets we have $\sqrt{N}(\hat{\gamma}_{k-1} - \gamma^*) = O_p(1)$ and $\nabla_{\theta\gamma'} \tilde{Q}_N(\bar{\theta}, \bar{\gamma}) = o_p(1)$. Therefore,

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

Furthermore, because $\hat{Y}_k = \Upsilon(\hat{\theta}_k, \hat{\gamma}_{k-1})$, with Υ twice continuously differentiable in a neighborhood of (θ^*, 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 the properties of Υ. Strong $\sqrt{N}$-consistency of $\hat{\gamma}_0$ completes the proof by induction.

A.2.3 Result 3 (Large Sample Convergence)

This result follows from Kasahara and Shimotsu (2012, Proposition 1). The zero Jacobian property ensures that the required spectral radius is equal to zero: as established in Result 2 above $\nabla_{\theta\gamma} \tilde{Q}^*(\theta^*, \gamma^*) = 0$, which is strictly less than 1.

A.3 Proof of Theorem 2

By examining the first-order conditions, we can see that $\hat{\theta}_{MLE} = \tilde{\theta}_N(\hat{\gamma}_{MLE})$, so that $\hat{Y}_{MLE} = \Upsilon(\tilde{\theta}_N(\hat{\gamma}_{MLE}), \hat{\gamma}_{MLE})$. This proves Result 1: the MLE is a fixed point of the EPL iterations.

Now let H_N denote the EPL iteration mapping by stacking the updating equations for θ

and Y so that $\hat{\gamma}_k = H_N(\hat{\gamma}_{k-1})$:

$$H_N(\gamma) = \begin{bmatrix} H_{1,N}(\gamma) \\ H_{2,N}(\gamma) \end{bmatrix} \equiv \begin{bmatrix} \tilde{\theta}_N(\gamma) \\ \Upsilon(\tilde{\theta}_N(\gamma), \gamma) \end{bmatrix}.$$

We then consider the first-order conditions evaluated at $\hat{\gamma}_{MLE}$. First we have

$$\nabla_\gamma H_{1,N}(\hat{\gamma}_{MLE}) = \nabla_\gamma \tilde{\theta}_N(\hat{\gamma}_{MLE}) \xrightarrow{p} \nabla_\gamma \tilde{\theta}^*(\gamma^*) = 0 \tag{4}$$

because $\nabla_{\theta\gamma'}\tilde{Q}^*(\theta^*, \gamma^*) = 0$, as shown in the proof of Theorem 1 (Result 2). Second, we have

$$\nabla_\gamma H_{2,N}(\hat{\gamma}_{MLE}) = \nabla_\theta \Upsilon(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE}) \nabla_\gamma \tilde{\theta}_N(\hat{\gamma}_{MLE}) + \nabla_\gamma \Upsilon(\hat{\theta}_{MLE}, \hat{\gamma}_{MLE}).$$

recalling that $\tilde{\theta}_N(\hat{\gamma}_{MLE}) = \hat{\theta}_{MLE}$. The first term converges to zero in probability as in (4). The second term is zero due to the zero Jacobian property of Υ, as can be seen by evaluating (3).

The above analysis implies that $\hat{\gamma}_{MLE} = H_N(\hat{\gamma}_{MLE})$ and $\nabla_\gamma H_N(\hat{\gamma}_{MLE}) \xrightarrow{p} 0$, which are key to Results 2 and 3. To obtain Result 2, note that the $\sqrt{N}$-consistency of $\hat{\gamma}_{MLE}$ implies $\nabla_\gamma H(\hat{\gamma}_{MLE}) = 0 + O_p(N^{-1/2})$. So, by a first-order expansion around $\hat{\gamma}_{MLE}$,

$$\begin{aligned} \hat{\gamma}_k &= H_N(\hat{\gamma}_{k-1}) \\ &= H_N(\hat{\gamma}_{MLE}) + \nabla_\gamma H_N(\hat{\gamma}_{MLE})(\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}) + O_p(||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}||^2) \\ &= \hat{\gamma}_{MLE} + (0 + O_p(N^{-1/2}))(\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}) + O_p(||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}||^2). \end{aligned}$$

It follows that

$$\hat{\gamma}_k - \hat{\gamma}_{MLE} = O_p(N^{-1/2}||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}|| + ||\hat{\gamma}_{k-1} - \hat{\gamma}_{MLE}||^2).$$

For Result 3, we appeal to continuity of $\nabla_\gamma H_N(\cdot)$ and $||\cdot||$. For any $\varepsilon > 0$, if $||\nabla_\gamma H_N(\hat{\gamma}_{MLE})|| < \varepsilon$, then there exists some neighborhood around $\hat{\gamma}_{MLE}$, $\mathcal{B}$, such that $H_N(\cdot)$ is a contraction mapping on $\mathcal{B}$ with Lipschitz constant, $L < \varepsilon$, and fixed point $\hat{\gamma}_{MLE}$. We have $\nabla_\gamma H_N(\hat{\gamma}_{MLE}) \xrightarrow{p} 0$, so that $||\nabla_\gamma H_N(\hat{\gamma}_{MLE})|| < \varepsilon$ w.p.a. 1 as $N \to \infty$. Result 3 follows immediately.

A.4 Proof of Theorem 3

All of the listed Υ functions satisfy the zero Jacobian property, so the results from the previous proofs carry through.

A.5 Proof of Lemma 2

We show that $v = \Phi(\theta, v)$ is a Bellman-like representation of the best-response equilibrium conditions, $P = \sigma(\theta, P)$. First, note that $P^j = \Lambda^j(v^j)$ for all j. And we have

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

We see that v^{-j} only influences $\Phi^j(\cdot)$ through its effect on $P^{-j} = \Lambda^{-j}(v^{-j})$ and we can then define $\phi^j(\theta, v^j, P^{-j}) = \Phi_a^j(\theta, v^j, v^{-j})$. It is straightforward to show that $||\nabla_{v^j}\phi^j||_\infty = \beta < 1$, so that $\phi^j(\cdot)$ is a Bellman-like contraction in v^j (fixing θ and P^{-j}) with a unique fixed point. Player j's best response is $\sigma^j(\theta, P^{-j}) = \Lambda^j(v^j)$ where $v^j = \phi^j(\theta, v^j, P^{-j})$. Imposing $v = \Phi(\theta, v)$ is therefore equivalent to imposing $P = \sigma(\theta, P)$; their sets of fixed points for a given θ are isomorphic.

B Additional Monte Carlo Results

In Section 4, we report Monte Carlo results for the model of Pesendorfer and Schmidt-Dengler (2008) for equilibrium (i), which is NPL-stable, and equilibrium (ii), which is NPL-unstable. Table 6 reports the some results for equilibrium (iii), which is also NPL-unstable, and therefore the results are qualitatively similar to those for equilibrium (ii) presented in Table 3.

Table 6: Monte Carlo Results for Pesendorfer and Schmidt-Dengler (2008) Equilibrium (iii) with $N = 1000$

Parameter	Statistic	1-NPL	1-EPL	∞-NPL	∞-EPL
$\theta_{\mathrm{M}} = 1.2$	Mean Bias	-0.0419	-0.0383	-0.2099	-0.0003
	MSE	0.0204	0.0194	0.0480	0.0174
$\theta_{\mathrm{C}} = -2.4$	Mean Bias	0.0948	0.0625	0.6806	0.0043
	MSE	0.1127	0.1000	0.4683	0.0987
$\theta_{\mathrm{EC}} = -0.2$	Mean Bias	-0.0277	-0.0133	-0.3146	-0.0044
	MSE	0.0286	0.0265	0.1023	0.0277
Converged	%			99.8%	100%
Iterations	Median			30	8
	IQR			5	2
Time (min.)	Total			0.3486	0.1026