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

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

February 8, 2021

Abstract

We propose a new sequential Efficient Pseudo-Likelihood (k-EPL) estimator for dynamic discrete choice games of incomplete information. We show that each iteration in the k-EPL sequence is consistent and asymptotically efficient, so the first-order asymptotic properties do not vary across iterations. Furthermore, we show the sequence achieves higher-order equivalence to the finite-sample maximum likelihood estimator with iteration and that the sequence of estimators converges almost surely to the maximum likelihood estimator at a nearly-superlinear rate when the data are generated by any regular Markov perfect equilibrium, including equilibria that lead to inconsistency of other sequential estimators. When utility is linear in parameters, k-EPL iterations are computationally simple, only requiring that the researcher solve linear systems of equations to generate pseudo-regressors which are used in a static logit/probit regression. Monte Carlo simulations demonstrate the theoretical results and show k-EPL's good performance in finite samples in both small- and large-scale games, even when the game admits spurious equilibria in addition to one that generated the data.

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, Lanier Benkard, Chris Conlon, Arvind Magesan, Mathieu Marcoux, Salvador Navarro, John Rust, and Eduardo Souza-Rodrigues for comments and insightful discussions. The paper has benefited from comments by seminar participants at Boston College, Cornell JCB, Georgetown University, University of Notre Dame, University of Toronto, the 2019 International Industrial Organization Conference (IIOC, Boston), the 2019 University of Calgary Empirical Microeconomics Workshop (Banff), the 2020 ASSA Annual Meeting (San Diego), the Spring 2020 IO^2 Seminar, and the 2020 Econometric Society World Congress (Milan).

†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 dynamic discrete choice models – particularly dynamic discrete games of incomplete information – is a topic of considerable interest in economics. 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 based on N independent markets, θ is a finite-dimensional vector of parameters, Y is a vector of important auxiliary parameters, and $G(\theta, Y) = 0$ is a vector equality constraint that represents equilibrium conditions. The parameters θ usually consist of the structural parameters of the model. Common examples of auxiliary parameters Y include expected/integrated value functions or conditional choice probabilities, since the equality constraint is often derived from an equilibrium fixed point condition of the form $G(\theta, Y) \equiv Y - \Gamma(\theta, Y) = 0$.

One approach to estimating these models is to directly impose the fixed point equation for each trial value of θ visited by the optimization algorithm by solving for Y_θ such that $G(\theta, Y_\theta) = 0$. This approach was pioneered by Rust (1987) and Pakes (1986) for single-agent models, where the fixed point is unique and can be computed via standard value function iteration or backwards induction. Solution algorithms are available for dynamic games (Pakes and McGuire, 1994, 2001), but it is often infeasible to nest those within an estimation routine because the computational burden can be quite large. Furthermore, in games the model may be incomplete due to multiple solutions Y_θ (Tamer, 2003).

These issues with the nested fixed-point approach led researchers to extend conditional choice probability (CCP) estimators – first introduced in the seminal work of Hotz and Miller (1993) for single-agent models – 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, a fixed point of the sequence. 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 CCP estimate, in the sense that it has the same limiting distribution as the (partial) maximum

[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).

likelihood estimator. Furthermore, Kasahara and Shimotsu (2008) showed that the sequence converges to the true parameter values with probability approaching one in large samples. Indeed, the Monte Carlo simulations in Aguirregabiria and Mira (2002) show that single-agent k-NPL reliably converges to the maximum likelihood estimate. The combination of computational simplicity, efficiency, and convergence stability make k-NPL an attractive alternative to other approaches, such as nested fixed point (computationally burdensome) or standard Newton or Fisher scoring steps on the full maximum likelihood problem (often diverges in finite samples in practice).[2]

Unfortunately, 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 (taking $k \to \infty$ or until convergence) estimator outperforms the 1-NPL estimator in efficiency when both are consistent. But Pesendorfer and Schmidt-Dengler (2010) 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 (see also, Kasahara and Shimotsu (2012), Egesdal, Lai, and Su (2015), and Aguirregabiria and Marcoux (2019)). Kasahara and Shimotsu (2012) show that inconsistency occurs when the NPL mapping is unstable at the data-generating equilibrium, which is essentially equivalent to best-response instability of the equilibrium.[3]

Another type of CCP estimator is the minimum-distance estimator proposed by Pesendorfer and Schmidt-Dengler (2008). This estimator is both consistent and efficient in dynamic games, and Bugni and Bunting (2020) develop a sequential version of this estimator, which we refer to as k-MD. However, Monte Carlo simulations show that these econometric properties come at the expense of greatly increased computational burden, taking 12 to 26 times longer per iteration than k-NPL in even the simplest small-scale dynamic games.[4] This difference in computation time is likely to grow with the size of the game, which would be a serious concern for empirical applications. This leaves the researcher with an undesirable tradeoff between the computationally simple k-NPL sequence and the more burdensome k-MD sequence with better efficiency properties.

But there is another concern shared by the k-NPL and k-MD sequences: finite sample performance of successive iterations. The k-MD estimator uses the NPL mapping to update

[2]This is a well-known limitation of standard Newton steps, and further regularization is often needed to ensure convergence.

[3]The NPL mapping is first-order equivalent to the best-response mapping at the equilibrium.

[4]Bugni and Bunting (2020) perform Monte Carlo simulations with a very small-scale game (two players, two actions, four states), and even in that setting they remark, "... computing the optimally weighted 1-MD, 2-MD, and 3-MD estimators takes us roughly 33%, 75%, and 80% more time than computing the 20-[NPL] estimator, respectively." These translate to roughly 26, 17.5, and 12 times longer per iteration for each respective case.

choice probabilities between iterations, so there is reason to be concerned that it may mimic k-NPL's finite sample properties when the data-generating equilibrium is NPL-unstable. While the first-order asymptotic analysis of Bugni and Bunting (2020) implies that both sequences are consistent for *finite* k, this asymptotic consistency does not necessarily lead to good performance in finite samples. Indeed, the finite-sample performance of k-NPL deteriorates even for finite k when the equilibrium is NPL-unstable. For example, Pesendorfer and Schmidt-Dengler (2008) consider a finite number of iterations for k-NPL and find that it is severely biased in large-but-finite samples when the equilibrium is NPL-unstable. In our own Monte Carlo simulations presented later, we find that substantial bias appears rather quickly, even for low values of k. This issue may also be a concern for k-MD, since it is also consistent for *finite* k but has unknown stability properties as $k \to \infty$, which may lead to deterioration in finite-sample performance even for fixed k, similar to k-NPL. This concern arises because Kasahara and Shimotsu (2012) show that instability/inconsistency of the k-NPL sequence arises from instability of the NPL mapping used to update the choice probabilities between iterations.[5]

With these concerns about various sequential methods in mind, an important question arises: is there a CCP-based sequence that achieves a balance between computational simplicity, asymptotic efficiency, and good finite-sample properties with any number of iterations – including as $k \to \infty$ (iterating to convergence) – in dynamic games? The primary contribution of this paper is to provide such a method, which we name the k-step Efficient Pseudo-Likelihood (k-EPL) estimator. We show that k-EPL estimates are first-order asymptotically equivalent to the maximum likelihood estimate for any number of iterations, $k \geq 1$. Thus, every estimate in the sequence is efficient. Furthermore, we also show that higher-order improvements are achieved with iteration, so the k-EPL sequence converges to the finite-sample maximum likelihood estimator almost surely. The convergence rate is fast, approaching super-linear as $N \to \infty$. This convergence result for k-EPL holds even when the data-generating equilibrium is best-response unstable, rendering k-NPL inconsistent.

Our k-EPL estimator also has a simple computational structure, similar to k-NPL. When utility is linear in the parameters of interest, both k-EPL and k-NPL iterations proceed in two stages: (i) solve a set of linear systems to generate pseudo-regressors; and (ii) use the pseudo-regressors in a static logit/probit maximum likelihood problem. The linear systems only need to be solved once per iteration in the sequence, and the static logit/probit problem is a low-dimensional, strictly concave problem that has a unique solution and is easy to solve with out-of-the-box optimization software. A difference between k-EPL and k-NPL does arise

[5]However, a rigorous econometric analysis of the the k-MD estimator's behavior as $k \to \infty$ is beyond the scope of this paper.

because the linear systems in k-EPL have larger dimension. We discuss this in detail below, but in short k-NPL proceeds on a player-by-player basis while the k-EPL system incorporates all players. While we do find that this creates an increase in relative computation time in large-dimensional problems, the increase appears much less severe than the computation time required for k-MD. We also find that iterating k-EPL to convergence can be faster than doing so with k-NPL when the game is not too large, since the need for fewer iterations results in lower overall computation time.

One interesting implication of our higher-order analysis is that iterating k-EPL can also provide a convenient algorithm for computing the maximum likelihood estimator. We explore this in some of our Monte Carlo simulations and find that it performs quite well, even with random starting values. However, our primary focus is on the entire k-EPL sequence, beginning with consistent initial estimates. We find that much of the practical improvement from iteration is achieved with just a few iterations in our Monte Carlo experiments, suggesting that low values of k can be quite effective when the initial CCP estimates are consistent.

In recent related work, Aguirregabiria and Marcoux (2019) studied the finite-sample properties of ∞-NPL (k-NPL with $k \to \infty$) and introduced a variant of the k-NPL algorithm that updates the conditional choice probabilities by applying spectral methods to the CCP updates. The goal of their algorithm is to improve convergence properties of ∞-NPL for unstable fixed points.[6] However, there are several differences between their work and ours. First, they limit their analysis of the spectral algorithm to computing the best fixed point of the k-NPL sequence, whereas we provide analysis for all the iterations in the k-EPL sequence. Second, upon convergence, the k-NPL and spectral k-NPL algorithms do not produce the maximum likelihood estimator, and convergence can require many iterations. In contrast, our k-EPL estimator has the same limiting distribution as the MLE at each iteration and usually converges locally to the MLE after few iterations in finite samples. We verify these properties in our simulation studies.

The remainder of the paper proceeds as follows. Section 2 describes a generic dynamic discrete choice game of incomplete information. Section 3 describes the k-EPL estimator, its asymptotic and finite-sample properties, and its numerical implementation. Section 4 provides Monte Carlo simulations. Section 5 concludes. All proofs appear in the Appendix.

[6] In the population or in finite samples, the NPL operator may have spectral radius larger than one for some equilibria, rendering it unstable. Conversely, the spectral radius of the EPL operator is zero in the population and near zero in finite samples.

2 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). Time is discrete, indexed by $t = 1, 2, 3.\ldots$. In a given market, there are J firms 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, $a_t^j \in \mathcal{A} = \{0, 1, 2, \ldots, |\mathcal{A}| - 1\}$. Action zero is the "outside option" when applicable. All players choose their actions simultaneously.

Agents have flow utilities (profits), $U^j(x_t, a_t^j, a_t^{-j}, \varepsilon_t^j; \theta_u)$, 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,

$$\mathrm{E}\left\{\sum_{s=0}^{\infty} \beta^{s-t} U^j(x_s, a_s^j, a_s^{-j}, \varepsilon_s^j; \theta_u) \,\middle|\, x_t, \varepsilon_t^j\right\}.$$

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

Assumption 1. *(Additive Separability)* $U^j(x_t, a_t^j, a_t^{-j}, \varepsilon_{at}^j; \theta_u) = \bar{u}(x_t, a_t^j, a_t^{-j}; \theta_u) + \varepsilon_t^j(a_t^j)$.

Assumption 2. *(Conditional Independence)* $p(x_{t+1}, \varepsilon_{t+1} \mid x_t, a_t, \varepsilon_t; \theta_f) = g(\varepsilon_{t+1}) f(x_{t+1} \mid x_t, a_t; \theta_f)$*, where* $g(\varepsilon_{t+1})$ *has full support and is thrice-continuously differentiable.*

Assumption 3. *(Independent Private Values) Private values are independently distributed across players.*

Assumption 4. *(Finite Observed State Space)* $x_t \in \mathcal{X} = \{1, 2, \ldots, |\mathcal{X}|\}$.

Assumptions 1–4 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 $a^{-j} \in \mathcal{A}^{|\mathcal{J}|-1}$, but rather on $P^{-j}(x) \in \Delta^{|\mathcal{J}|-1}$, where $P^{-j}(x)$ is player j's belief about the other players' probability of playing the corresponding actions in state x 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 and transition probabilities:

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

$$f^j(x' \mid x, a^j; P^{-j}, \theta_f) = \sum_{a^{-j} \in \mathcal{A}^{|\mathcal{J}|-1}} P^{-j}(x, a^{-j}) f(x' \mid x, a^j, a^{-j}; \theta_f)$$

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 $P^j(a) = \Lambda^j(a; v^j)$. And let

$$\Lambda^{-j}(v^{-j}) = (\Lambda^1(v^1), \ldots, \Lambda^{j-1}(v^{j-1}), \Lambda^{j+1}(v^{j+1}), \ldots, \Lambda^{|\mathcal{J}|}(v^{|\mathcal{J}|})),$$

so that in equilibrium $P^{-j} = \Lambda^{-j}(v^{-j})$. Aguirregabiria and Mira (2007) show that in equilibrium, the choice-specific value functions are equal to

$$v^j(x, a^j) = u^j(x, a^j; P^{-j}, \theta_u) + \beta \sum_{x'} f^j(x' \mid x, a^j; P^{-j}, \theta_f) \Gamma^j(x'; \theta, P),$$

where

$$\Gamma^j(\theta, P) = \big(I - \beta F^j(\theta_f, P)\big)^{-1} \sum_{a^j} P^j(a^j) * \big(u^j(a^j; P^{-j}, \theta_u) + e(a^j; P^j)\big)$$

maps into *ex-ante* (or integrated) value function space. They then define the NPL operator, $\Psi(\theta, P)$, such that $\Psi^j(\theta, P) = \Lambda^j(\Gamma^j(\theta, P))$ and combining all players yields the following fixed-point condition that describes any Markov perfect equilibrium (Aguirregabiria and Mira 2007, Lemma 1):

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

While this equilibrium representation based on the NPL operator is often useful, we will ultimately want to work with an alternative representation when implementing our new estimator. This alternative arises due to a change of variables from P space to v space. (See Section 3.3 for a detailed discussion of the importance of this change.) Define the function

$$\Phi^j(x, a^j; v^j, v^{-j}, \theta) = u^j(x, a; \Lambda^{-j}(v^{-j}), \theta_u) + \beta \sum_{x'} f^j(x' \mid x, a^j; \Lambda^{-j}(v^{-j}), \theta_f) S(v^j(x')),$$

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 McFadden's social surplus function.[7] This $\Phi(\cdot)$ function allows us to characterize the equilibrium with an alternative fixed-point equation, as described in the following Lemma.

[7]For example, $S(v^j(x)) = \ln(\sum_{a^j} \exp(v^j(x, a^j)))$ when the private values are i.i.d. and follow the type 1 extreme value distribution.

Lemma 1. *(Representation Lemma) Under Assumptions 1–4, choice-specific value functions characterize a Markov perfect equilibrium for θ if and only if $v^j(x, a^j) = \Phi^j(x, a^j; \theta, v^j, v^{-j})$ for all $(j, x, a) \in \mathcal{J} \times \mathcal{X} \times \mathcal{A}$. Or more succinctly,*

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

3 The k-EPL Estimator

This section describes the k-EPL estimator and discusses its asymptotic and finite-sample properties, as well as computational aspects of its implementation in dynamic discrete choice games. We begin by discussing maximum likelihood estimation, subject to an equilibrium constraint based on some nuisance parameter, Y. The alternative statements of the equilibrium conditions in Section 2 yield two potential choices for Y and the corresponding constraint. We will discuss the actual choice of nuisance parameter later in Section 3.3.

Let w_i for $i = 1, \ldots, N$ denote the observations from N independent markets. 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

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

and let $Q^*(\theta, Y) \equiv \mathrm{E}[Q_N(\theta, Y)]$.

Assumption 5. *(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. Q^* has a unique maximum in $\Theta \times \mathcal{Y}$ subject to $G(\theta, Y) = 0$, and the maximum occurs at (θ^*, Y^*). (d) $G(\theta, Y)$ is thrice continuously differentiable and $\nabla_Y G(\theta^*, Y^*)$ is non-singular.*

Assumptions 5(a)-(c) echo standard identification assumptions. We note that assuming that Q^* has a unique maximum does not rule out games with multiple equilibria. Non-singularity of the Jacobian in (d) is the defining feature of regular Markov perfect equilibria in the sense of Doraszelski and Escobar (2010). Regularity essentially means that the equilibrium is locally isolated and we can apply the implicit function theorem to obtain $Y(\theta)$

locally.[8]

One method to estimate these models is via constrained maximum likelihood:

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

An equivalent statement is $\hat{\theta}_{\text{MLE}} = \arg\max_\theta \; Q_N(\theta, Y(\theta))$, where $G(\theta, Y(\theta)) = 0$ and $\hat{Y}_{\text{MLE}} = Y(\hat{\theta}_{\text{MLE}})$. Pseudo-likelihood estimation replaces $Y(\theta)$ with some other mapping. Aguirregabiria and Mira (2007) define $Y \equiv P$ and replace $Y(\theta) \equiv P(\theta)$ with their $\Psi(\theta, \hat{P}_{k-1})$ for the k-th iteration in the k-NPL sequence. However, this procedure suffers from the issues discussed in Section 1.

Our k-step Efficient Pseudo-Likelihood (k-EPL) sequence instead uses a "Newton-like" step, which provides a good approximation to the full Newton step but uses a fixed value of $\nabla_Y G(\theta, Y)^{-1}$; this value varies between steps but does not vary as the optimizer searches over different values θ within each step. Algorithm 1 below defines our sequential estimation procedure. It uses our Newton-like mapping, $\Upsilon(\cdot)$, which is a function of an initial compound parameter vector, $\gamma = (\theta, Y)$, and an additional (possibly different) value of θ.

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

- *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:*

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

where

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

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 numerical convergence.*

[8] Aguirregabiria and Mira (2007) directly assume the local existence of $Y(\theta)$, instead of appealing to the implicit function theorem.

This procedure enjoys some nice econometric properties, both asymptotically and in finite samples. These properties arise from some convenient features of the $\Upsilon(\cdot)$ function, which are detailed in the following lemma.

Lemma 2. *Let $\Upsilon(\theta, \gamma)$ denote the operator defined in Algorithm 1 and define $Y_\theta \equiv Y(\theta)$ and $\gamma_\theta \equiv (\theta, Y_\theta)$. Under Assumption 5, 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: $\Upsilon(\theta, \gamma_\theta) = Y(\theta) \iff G(\theta, Y_\theta) = 0$.
2. $\nabla_\theta \Upsilon(\theta, \gamma_\theta) = \nabla_\theta Y(\theta)$.
3. $\nabla_\gamma \Upsilon(\theta, \gamma_\theta) = 0$ (Zero Jacobian Property).

Lemma 2 is the key to most of the results in this section, which arise from applying the lemma at (θ^*, γ^*) and $(\hat{\theta}^{MLE}, \hat{\gamma}^{MLE})$. For now, we note that Result 3 of Lemma 2 is analogous to the "zero Jacobian" property from Proposition 2 of Aguirregabiria and Mira (2002), which was the key to both their efficiency results and the finite-sample convergence results of Kasahara and Shimotsu (2008) for single-agent k-NPL. By utilizing Newton-like steps on the equilibrium constraint, the k-EPL algorithm restores this zero Jacobian property in dynamic games.

3.1 Asymptotic Properties of k-EPL

One implication of Lemma 2 is that k-EPL then gives a sequence of asymptotically efficient estimators that converges almost surely in large samples. We state this result formally in the following Theorem.

Theorem 1. *(Asymptotic Properties of k-EPL) Under Assumption 5, 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. *(Efficiency) $\sqrt{N}(\hat{\theta}_k - \theta^*) \xrightarrow{d} \mathrm{N}(0, \Omega^{*-1}_{\theta\theta})$, 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 for k-EPL in games are shared by k-NPL in single-agent models, but not games. In short, the zero Jacobian property ensures that $\hat{\gamma}_k = (\hat{\theta}_k, \hat{Y}_k)$ is asymptotically orthogonal to $\hat{\gamma}_{k-1}$, so that using $\hat{\gamma}_{k-1}$ is asymptotically equivalent to using

$\gamma^* = (\theta^*, Y^*)$ at each step. This drives the consistency (Result 1) and asymptotic equivalence to MLE (Result 2) of each step. Intuitively, an EPL step is similar to a Newton step on the full maximum likelihood problem, although iterating on that procedure is notoriously unstable unless properly regularized. Our Monte Carlo simulations in Section (4) show that k-EPL is stable without further regularization.

3.2 Iteration: Finite-Sample Properties of k-EPL

While the asymptotic distribution is insensitive to iteration, we can obtain substantial finite-sample improvements by iterating and can even compute the finite-sample MLE by iterating to convergence. We proceed with a formal econometric analysis of these favorable finite-sample results here. Later on, they are illustrated in the Monte Carlo simulations in Section 4.

Our finite-sample results are similar to those of the single-agent version of k-NPL in Aguirregabiria and Mira (2002), which also has the zero Jacobian property. In their Monte Carlo simulations, single-agent k-NPL iterations exhibit rapid finite-sample improvements and reliably converged to the finite-sample MLE. Kasahara and Shimotsu (2008) then provided a formal econometric explanation for these results. The analysis of k-EPL's finite sample properties in this section is similar but also applies to games.

The only additional requirement for our finite-sample results is that the Jacobian of the equality constraints, G, with respect to Y is nonsingular at the finite-sample MLE.

Assumption 6. $\nabla_Y G(\hat{\theta}_{MLE}, \hat{Y}_{MLE})$ *is non-singular.*

Theorem 2. *(Finite Sample Properties) Suppose Assumptions 5 and 6 hold and that the optimization problem in Step 2 of Algorithm 1 has a unique solution for all $k \geq 1$. Then,*

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 k-EPL iterations in a finite sample, similar to Aguirregabiria and Mira (2002, Proposition 3) for

single-agent k-NPL. The second result gives a theoretical explanation of why we should expect iteration to yield improvements in finite samples. This result is analogous to Proposition 2 of Kasahara and Shimotsu (2008), although their result only applies in the single-agent case. In short, even though iteration on EPL provides no improvement up to $O_p(N^{-1/2})$, it still yields higher-order improvements. To see why, suppose the initial estimates are such that $\hat{\gamma}_0 - \gamma^* = O_p(N^{-b})$ for $b \in (1/4, 1/2]$, so that $||\hat{\gamma}_0 - \hat{\gamma}_{\text{MLE}}|| = O_p(N^{-b})$.[9] Repeated substitution gives $||\hat{\gamma}_k - \hat{\gamma}_{\text{MLE}}|| = O_p(N^{-(k-1)/2-2b})$. In particular, in the case where the state space is finite and frequency or kernel estimates are used, $b = 1/2$ and $||\hat{\gamma}_k - \hat{\gamma}_{\text{MLE}}|| = O_p(N^{-(k+1)/2})$, where $N^{-(k+1)/2} \to 0$ 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 super-linear. However, we only have access to finite samples in practice, 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. We can therefore 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.[10] We demonstrate this usage of k-EPL in our Monte Carlo experiments in Section 4 and find that it works well.

Aside from k-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_θ such that $G(\theta, Y_\theta) = 0$ in an inner loop. MPEC leverages modern optimization software to search over θ and Y simultaneously, only imposing that $G(\theta, Y) = 0$ at the solution. The algorithm of choice may depend on the structure of the model.

While this section discusses the k-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 often computationally unattractive—or even infeasible—in such games.[11] MPEC, however,

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

[10] Aguirregabiria and Mira (2007) suggest a similar procedure to find ∞-NPL when no initial consistent estimate is available, and multiple starting values are often used when computing the maximum likelihood estimate with other methods (Su and Judd (2012); Egesdal et al. (2015)).

[11] We note that NFXP still performs well in single-agent dynamic models. See Doraszelski and Judd (2012)

remains feasible and performs well, as demonstrated by Egesdal et al. (2015). The key difference here between MPEC and k-EPL is that k-EPL will be able to more heavily exploit the structure of the problem.[12] In Section 3.3, we show that — much like k-NPL in single-agent models — common modeling assumptions lead to EPL iterations composed of two easily-computed parts: solving linear systems to form pseudo-regressors, followed by solving an unconstrained, globally concave maximization problem á la static logit/probit with those pseudo-regressors. Neither of these operations require sophisticated commercial 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.

3.3 Computational Details and Choice of Nuisance Parameter

k-EPL is particularly useful when we are interested in estimating the flow utility parameters, θ_u. In many cases – including our Monte Carlo experiments – the transition parameter, θ_f, is known in advance.[13] So, we focus on estimating only the flow utility parameters and let $\theta \equiv \theta_u$. And we further assume that flow utilities are linear in θ, which is a standard assumption for dynamic discrete choice models.[14]

One of the most useful features of k-NPL in this setting is the computational simplicity of each iteration. In both single-agent models and games, the optimization problem in each k-NPL iteration reduces to two stages: i) solving linear systems to generate pseudo-regressors; and ii) using the generated regressors in solving a globally concave unconstrained optimization problem á la logit/probit. This computational simplicity arises because k-NPL is based on the fixed-point condition that $P = \Psi(\theta, P)$ and follows steps similar to k-EPL but with $\Upsilon(\cdot)$ replaced with $\Psi(\cdot)$. The key is that $\Psi(\theta, P) = \Lambda(\Gamma(\theta, P))$, where $\Gamma(\cdot)$ maps into v space and $\Lambda(\cdot)$ maps into P space. When the distribution of the unobserved shocks, $g(\varepsilon)$, is log-concave – such as with i.i.d. type I extreme value (logit) or standard normal (probit) distributions – the function $\Lambda(\cdot)$ is log-concave. When the flow utilities are linear in θ, Aguirregabiria and Mira (2007) show that $\Gamma(\theta, P)$ is also linear in θ. So, $\Psi(\theta, P) = \Lambda(\Gamma(\theta, P))$

and Arcidiacono, Bayer, Blevins, and Ellickson (2016) for details on the computational burden of computing equilibria in discrete-time dynamic discrete games.

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

[13] Alternatively, θ_f can be estimated in a first stage, with θ_u then estimated via partial maximum likelihood. Similarly to single-agent k-NPL, k-EPL iterations based on this partial MLE problem yield asymptotic equivalence and finite-sample convergence to partial MLE because the zero Jacobian property still applies.

[14] See, e.g., Rust (1987); Aguirregabiria and Mira (2002, 2007); Bajari et al. (2007); Pakes et al. (2007); Pesendorfer and Schmidt-Dengler (2008); Arcidiacono and Miller (2011); Egesdal et al. (2015); Bugni and Bunting (2020).

is log-concave, and the optimization step resembles static unconstrained logit/probit because it solves

$$\hat{\theta}_{k,\text{NPL}} = \arg\max_{\theta\in\Theta} \quad N^{-1}\sum_{i=1}^{N}\sum_{t}\sum_{j} \ln \Psi(x_t, a_t^j; \theta, \hat{P}_{k-1,\text{NPL}}).$$

However, we will lose this computational simplicity if we perform k-EPL estimation with this same equilibrium constraint on P. Setting $Y \equiv P$, we have

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

$$\Upsilon(\theta, \hat{\gamma}_{k-1}) = \hat{P}_{k-1} - \left(I - \nabla_P \Psi(\hat{\theta}_{k-1}, \hat{P}_{k-1})\right)^{-1}\left(P - \Psi(\theta, \hat{P}_{k-1})\right)$$

and the optimization step in the k-EPL algorithm solves

$$\hat{\theta}_{k,\text{EPL}} = \arg\max_{\theta\in\Theta} \quad N^{-1}\sum_{i=1}^{N}\sum_{t}\sum_{j} \ln \Upsilon(x_t, a_t^j; \theta, \hat{\gamma}_{k-1}).$$

Several computational issues arise. First, instead of solving linear systems once before the optimization step, we must repeatedly solve linear systems throughout the optimization because of the need to compute $(I - \nabla_P \Psi(\hat{\theta}_{k-1}, \hat{P}_{k-1}))^{-1}\Psi(\theta, \hat{P}_{k-1})$ for each new value of θ in the search. Second, we will lose the guarantee of concavity of the optimization problem in each step. Even though $\Psi(\theta, \hat{P}_{k-1})$ is log-concave in θ, this does not guarantee log-concavity of $\Upsilon(\theta, \hat{\gamma}_{k-1})$ because affine transformations of log-concave functions are not necessarily log-concave. And third, the Newton-like steps rely on an implicit linearization: even though $\Psi(\theta, \hat{P}_{k-1})$ maps into the probability simplex, $\Upsilon(\theta, \hat{\gamma}_{k-1})$ can arrive at values outside the simplex.[15] Thus, we would need to add constraints to the optimization problem to ensure $\Upsilon(\theta, \hat{\gamma}_{k-1})$ does not leave the unit simplex.

Fortunately, we can restore computational simplicity to the EPL iterations by instead setting $Y \equiv v$ and using the equilibrium representation from Lemma 1: $v = \Phi(\theta, v)$. Suppose flow utilities are linear in θ, so that $u^j(\theta, x, a^j, P^{-j}) = h(x, a^j, P^{-j})'\theta$. Because $P^{-j} = \Lambda^{-j}(v^{-j})$, we can re-write this in terms of v: $u^j(\theta, x, a^j, v^{-j}) = h(x, a^j, v^{-j})'\theta$. Inspecting the form of $\Phi(\theta, v)$, we see that it will be linear in θ and therefore so will $G(\theta, v) = v - \Phi(\theta, v)$:

$$G(\theta, v) = H(v)\theta + z(v),$$

[15] Technically, $\Psi(\cdot)$ maps into a Cartesian product of the interior of the unit simplex due to each player having their own strategies.

where $H(\cdot)$ is a matrix and $z(\cdot)$ is a vector. As a result, $\Upsilon(\theta, \hat{\gamma}_{k-1})$ is also linear in θ:

$$\begin{aligned}
\Upsilon(\theta, \hat{\gamma}_{k-1}) &= \hat{v}_{k-1} - \nabla_v G(\hat{\theta}_{k-1}, \hat{v}_{k-1})^{-1} G(\theta, \hat{v}_{k-1}) \\
&= \hat{v}_{k-1} - \nabla_v G(\hat{\theta}_{k-1}, \hat{v}_{k-1})^{-1} \left(H(\hat{v}_{k-1})\theta + z(\hat{v}_{k-1}) \right) \\
&= -\nabla_v G(\hat{\theta}_{k-1}, \hat{v}_{k-1})^{-1} H(\hat{v}_{k-1})\theta \\
&\quad + \hat{v}_{k-1} - \nabla_v G(\hat{\theta}_{k-1}, \hat{v}_{k-1})^{-1} z(\hat{v}_{k-1}) \\
&= A(\hat{\gamma}_{k-1})\theta + b(\hat{\gamma}_{k-1}).
\end{aligned}$$

The need to solve linear systems arises from computing $\nabla_v G(\hat{\theta}_{k-1}, \hat{v}_{k-1})^{-1} H(\hat{v}_{k-1})$ and $\nabla_v G(\hat{\theta}_{k-1}, \hat{v}_{k-1})^{-1} z(\hat{v}_{k-1})$ in order to obtain $A(\hat{\gamma}_{k-1})$ and $b(\hat{\gamma}_{k-1})$, and this can be done outside the optimization search. The optimization step in Algorithm 1 then becomes

$$\hat{\theta}_k = \arg\max_{\theta \in \Theta} \quad N^{-1} \sum_{i=1}^{N} \sum_t \sum_j \ln \Lambda \left(\Upsilon(x_t, a_t^j; \theta, \hat{\gamma}_{k-1}) \right).$$

This problem is globally concave in θ because $\Lambda(\cdot)$ is log-concave and $\Upsilon(\theta, \hat{\gamma}_{k-1})$ is linear in θ, so that $\ln \Lambda(\Upsilon(\theta, \hat{\gamma}_{k-1}))$ is concave in θ. And there is no need to impose any constraints because v can take any value on a Cartesian product of the real line. Furthermore, $\nabla_v G(\hat{\theta}_{k-1}, \hat{v}_{k-1})$ can be computed analytically when $\Lambda(\cdot)$ and its derivative have an analytic form, such as the logit or probit cases. We therefore recommend this implementation of the k-EPL algorithm in practice, and it is what we use in our Monte Carlo simulations in Section 4.

While our EPL iterations with $Y \equiv v$ have a similar computational structure to NPL iterations (with $Y \equiv P$) insofar as both require solving linear systems then a globally concave optimization problem, the dimension of the linear systems in k-EPL is larger. Aguirregabiria and Mira (2007) show that k-NPL requires solving $|J|(|\Theta|+1)$ different systems of linear equations, each of dimension $|\mathcal{X}|$, resulting in a worst-case bound of $O((|\Theta|+1)|\mathcal{X}|^3|\mathcal{J}|)$ flops.[16] On the other hand, k-EPL requires solving $|\Theta|+1$ different systems of linear equations, each of dimension $|\mathcal{J}||\mathcal{X}||\mathcal{A}|$, resulting in a larger worst-case bound of $O((|\Theta|+1)|\mathcal{X}|^3|\mathcal{J}|^3|\mathcal{A}|^3)$ flops. Sparsity of the linear systems – a common feature in dynamic discrete choice models (Egesdal et al. (2015)) – can lower these bounds for both k-NPL and k-EPL. Fortunately, in practice, in our largest Monte Carlo experiment the relative difference is much lower than suggested by the worst-case bounds.

We then see a tradeoff between efficiency and computational burden, a common theme in estimating dynamic games. This theme appears in Bugni and Bunting's (2020) analysis

[16] There are $|\Theta|+1$ systems for each of the $|\mathcal{J}|$ players.

comparing their efficient k-MD (minimum distance) estimator to k-NPL. Even in what is essentially the smallest scale game possible – 2 players, 2 actions, 4 states – they report for k-MD a large increase in computational burden over k-NPL, with individual iterations taking about 12 to 26 times longer on average, depending on the number of iterations. It is perhaps reasonable to expect that this difference will grow with the size of the game, which is a concern for practitioners who must balance econometric efficiency with computational feasibility. Our k-EPL estimator, on the other hand, induces a much less severe increase in computational burden while retaining efficiency. In Section 4, we also explore a 2 player, 2 action, 4 state game and find that the difference between computational time for k-EPL and k-NPL iterations is negligible in that setting. Additionally, we explore a much larger-scale game – 5 players, 2 actions, 160 states – that is more representative of empirically relevant models. In this larger-scale game, we find that there is indeed an increase in time per iteration for k-EPL relative to k-NPL, but this increase in the larger-scale game (between 4 to 8 times) is not even as large as the 12 to 26-fold increase for k-MD in the much smaller-scale game.

Even with an increase in computation time per iteration relative to k-NPL, k-EPL can still ultimately be more attractive than k-NPL. First, its asymptotic efficiency, convergence properties, and rapid finite-sample improvements are attractive features that may be worth the increased computational burden of each iteration. Second, even in cases where both k-EPL and k-NPL converge to consistent estimates, k-EPL enjoys a much faster convergence rate than k-NPL, resulting in fewer iterations to convergence. So, iterating to convergence on k-EPL to obtain the finite-sample MLE can still be faster than computing the ∞-NPL estimator (if it converges), even though each individual iteration takes longer.

In many applications, the dominant source of computational burden for either estimator will often be the size of the state space, $|\mathcal{X}|$, since it can be large when $|\mathcal{A}|$ and $|\mathcal{J}|$ are small and also tends to grow with both $|\mathcal{A}|$ and $|\mathcal{J}|$ in dynamic games. One simple yet salient illustration arises when the state is determined by the previous actions of the players, so that $|\mathcal{X}| = |\mathcal{A}|^{|\mathcal{J}|}$. Thus, as $|\mathcal{J}|$ grows, $|\mathcal{X}|$ ultimately becomes the main source of computational burden for the linear systems in both k-NPL and k-EPL. To help deal with large state spaces, Aguirregabiria and Mira (2007) show that the linear systems required for k-NPL can be solved via an iterative process reminiscent of value function Bellman iteration, so that their worst-case computational burden reduces to $O((|\Theta|+1)|\mathcal{X}|^2|\mathcal{J}|)$. Similarly, k-EPL can also use alternative iterative methods to solve the linear systems such as Krylov subspace methods, although it cannot use Bellman-style iteration.

3.4 Single-Agent Dynamic Discrete Choice

We conclude this section by showing that k-NPL in a single-agent dynamic discrete choice model (Aguirregabiria and Mira, 2002) is equivalent to k-EPL with a slightly modified definition of $\Upsilon(\cdot)$. Here, we can work directly in probability space, $Y \equiv P$, and let

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

with $\Psi(\theta, P)$ defined as in Section 2 but with only a single agent.

We now have

$$\nabla_P G(\theta, P) = I - \nabla_P \Psi(\theta, P).$$

Proposition 2 from Aguirregabiria and Mira (2002) shows that $\nabla_P \Psi(\theta, P_\theta) = 0$, where $P_\theta = \Psi(\theta, P_\theta)$. Thus, $\nabla_P G(\theta, P_\theta) = I$ for all θ. So, we can use a modified definition of $\Upsilon(\cdot)$, where $\nabla_P G(\hat{\theta}_{k-1}, \hat{P}_{k-1})$ is simply replaced with the identity matrix, I, and we obtain

$$\begin{aligned}\Upsilon(\theta, \hat{\gamma}_{k-1}) &= \hat{P}_{k-1} - I^{-1}\left(\hat{P}_{k-1} - \Psi(\theta, \hat{P}_{k-1})\right) \\ &= \Psi(\theta, \hat{P}_{k-1}).\end{aligned}$$

This modified implementation of k-EPL is identical to k-NPL.

This equivalence of k-NPL to k-EPL in single agent models is unsurprising for two 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 methods. Second, Aguirregabiria and Mira (2002, Proposition 1(c)) show that their policy iterations are equivalent to Newton-like iterations on the (ex-ante) value function in single-agent models. Since k-EPL is built around Newton iterations, such an equivalence is again suggestive of the relationship shown here.

4 Monte Carlo Simulations

In this section, we present Monte Carlo experiments in two settings. First, we consider a small-scale dynamic entry-exit game with two firms from Pesendorfer and Schmidt-Dengler (2008). Next, we implement a larger-scale dynamic entry-exit game with five heterogeneous firms from Aguirregabiria and Mira (2007). The small-scale game is useful because it shows that k-EPL is reliable even when the game has multiple, possibly unstable equilibria (with data generated from only one of them). The larger-scale game is useful to illustrate how the method performs with a richer state space and to examine the stability of k-EPL and k-NPL

as competition in the market increases, leading to an increase in the spectral radius of the NPL operator. In the Appendix we also apply k-EPL to the static game of Pesendorfer and Schmidt-Dengler (2010), where we can easily compare ∞-EPL to the MLE computed via the nested fixed-point algorithm.

4.1 Pesendorfer and Schmidt-Dengler (2008): Dynamic Game with Multiple Equilibria

Here, we conduct simulations where we estimate the model from Pesendorfer and Schmidt-Dengler (2008), a small-scale dynamic game with multiple possible equilibria. There are two firms indexed by $j \in \{1, 2\}$ who choose an action in each market i, denoted $a_i^j \in \{0, 1\}$, where 1 is entry and 0 is exit. The observed state variable $x_i = (x_i^1, x_i^2) \in \{(0,0), (0,1), (1,0), (1,1)\}$ represents the incumbency status of firms 1 and 2, respectively. Flow utilities are period profits:

$$\tilde{u}^j(x_i, a_i^j = 1) = \theta_{\mathrm{M}} + \theta_{\mathrm{C}} a_i^{-j} + \theta_{\mathrm{EC}}(1 - x_i^j) + \varepsilon_1^j,$$

$$\tilde{u}^j(x_i, a_i^j = 0) = \theta_{\mathrm{SV}} x_i^j + \varepsilon_0^j,$$

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 $(\theta_{\mathrm{M}}, \theta_{\mathrm{C}}, \theta_{\mathrm{EC}}, \theta_{\mathrm{SV}}, \beta) = (1.2, -2.4, -0.2, 0.1, 0.9)$. The private shocks have distribution $\varepsilon_a^j \sim \mathrm{N}(0, 0.5)$. 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 utility values are the same and hence the equilibria are also 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, due to its Newton-like form, 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 theoretical analysis of Kasahara and Shimotsu (2012).[17] In contrast, we expect that k-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

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

fixed at $\theta_{SV} = 0.1$ and $\beta = 0.9$. We present results for estimation using k-NPL and k-EPL. The initial estimates of the conditional choice probabilities are sample frequencies, $\hat{P}^j(x, a^j)$. 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.[18]

In order to implement k-EPL in this context, we use the fixed-point constraint in choice-specific value function space, defined in Section 2. 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 $\hat{v}_0^j(x, a^j) = u^j(x, a^j; \hat{\theta}_0, \hat{P}^{-j}) + \beta f^j(x, a^j; \hat{P}^{-j})' \Gamma^j(\hat{\theta}_0, \hat{P})$, where

$$\Gamma^j(\theta, P) = \left(I - \beta F^j(\theta, P)\right)^{-1} \sum_{a^j} P^j(a^j) * \left(u^j(a^j; \theta, P^{-j}) + e(a^j; P^j)\right)$$

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

4.1.1 Results for NPL-Stable Equilibrium (i)

Table 1 shows results for equilibrium (i), for which k-NPL is stable and consistent. We consider both the one-step k-NPL and k-EPL estimators (for $k = 1$) 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 1 about here.]

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

Comparing 1-NPL to 1-EPL in Table 1, 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 k-NPL is expected to perform well, the efficiency of k-EPL yields improvements. Time per iteration is similar for k-NPL and k-EPL here, so the lower computational times are driven by the significant reduction in the number of iterations to convergence.[19] In this case for the smaller sample size, the results are more mixed. The mean bias of k-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.

4.1.2 Results for NPL-Unstable Equilibrium (ii)

[Table 2 about here.]

Turning to equilibrium (ii), for which the NPL fixed point is unstable, we have a very different picture. The results are presented in Table 2. 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 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 k-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.

4.1.3 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

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

estimates or multiple starting values, or both. First we explore using a single, polluted version of the consistent estimates as the starting value. This is for the purposes of exploring the effect of noise in consistent starting values with respect to the locality of convergence. In practice, of course, we recommend using multiple starting values.

[Table 3 about here.]

In Table 3 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 1, 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 2. 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 k-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 first-stage CCP estimation accuracy is a concern. With that in mind, in the next section we consider moving away from consistent starting values entirely.

4.1.4 k-EPL as a Computational Procedure (Random Starting Values)

Rather than rely solely on a single consistent estimate, we consider here using k-EPL as a computational procedure to compute the MLE using multiple starting values (in practice, ideally 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 k-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.[20] k-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 k-EPL to compute the MLE.

[Table 4 about here.]

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 starting value.[21] Finally, for both ∞-NPL and ∞-EPL we select from among the five estimates the one that maximizes the likelihood function. The results in Table 4 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 k-EPL and k-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 k-EPL is faster, converges more often, and has smaller MSE. In the large sample, k-EPL always has lower bias and MSE in addition to being more stable and computationally lighter. Also as before, k-EPL is robust to the NPL-unstable equilibrium. These results show that ∞-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 ∞-NPL would yield inconsistent estimates.

4.2 Aguirregabiria and Mira (2007): Dynamic Game with Several Heterogeneous Firms

In this section we report the results of a series of Monte Carlo experiments based on the model of Aguirregabiria and Mira (2007). This is an empirically relevant model that forms the basis of many applications but has also been used (sometimes in simplified forms) in simulation studies by Kasahara and Shimotsu (2012), Egesdal et al. (2015), Bugni and Bunting (2020), Blevins and Kim (2019), and Aguirregabiria and Marcoux (2019). In particular, Aguirregabiria and Marcoux (2019) discuss in detail how the spectral radius of the NPL operator increases with the strength of competition in the model. As such, we take as our

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

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

baseline case the parameters as Experiment 2 of Aguirregabiria and Mira (2007). We then follow Aguirregabiria and Marcoux (2019) and increase the competitive effect parameter, θ_{RN} to investigate how k-EPL and k-NPL behave as the spectral radius of the NPL operator increases to the point of instability and beyond.

The model is a dynamic entry-exit game with $|\mathcal{J}| = 5$ firms that operate in N independent markets. The firms have heterogeneous fixed costs. There is a single common market state, the market size, which can take one of five values: $s_{it} \in \{1, 2, 3, 4, 5\}$. Market size follows a 5×5 transition matrix and we use the same transition matrix as Aguirregabiria and Mira (2007). The other observable states are the incumbency statuses of the five firms, denoted $a^j_{i,t-1}$. There are therefore $5 \times 2^5 = 160$ distinct states in the model. Hence the state in market i at time t can be represented in vector form as $x_{it} = (s_{it}, a^1_{i,t-1}, a^2_{i,t-1}, a^3_{i,t-1}, a^4_{i,t-1}, a^5_{i,t-1})$.

Given the state of the model at the beginning of the period, firms simultaneously choose whether to operate in the market, $a^j_{it} = 1$, or not, $a^j_{it} = 0$. They make these decisions in order to maximize expected discounted profits, where the period profit function for an active firm is

$$\bar{u}^j(x_{it}, a^j_{it} = 1, a^{-j}_{it}; \theta) = \theta_{\text{FC},i} + \theta_{\text{RS}} s_{it} - \theta_{\text{RN}} \ln\left(1 + \sum_{l \neq j} a^l_{it}\right) - \theta_{\text{EC}}(1 - a^j_{i,t-1})$$

and $\bar{u}^j(x_{it}, a^j_{it} = 0, a^{-j}_{it}; \theta) = 0$ for inactive firms. The game is dynamic because firms must pay an entry cost θ_{EC} to enter the market and because firms have forward-looking expectations about market size and entry decisions of rival firms. The private information shocks $\varepsilon^j_{it}(a^j_{it})$ are independent and identically distributed across time, markets, players, and actions and follow the standard type I extreme value distribution.

We choose the model parameters following Aguirregabiria and Mira (2007) and Aguirregabiria and Marcoux (2019). In particular, the fixed costs for the five firms are $\theta_{\text{FC},1} = -1.9$, $\theta_{\text{FC},2} = -1.8$, $\theta_{\text{FC},3} = -1.7$, $\theta_{\text{FC},4} = -1.6$, and $\theta_{\text{FC},5} = -1.5$. The coefficient on market size is $\theta_{\text{RS}} = 1$ and the common firm entry cost is $\theta_{\text{EC}} = 1$. The only parameter that differs across our three experiments is the competitive effect θ_{RN}, which we set to be $\theta_{\text{RN}} = 1$ in Experiment 1, $\theta_{\text{RN}} = 2.5$ in Experiment 2, and $\theta_{\text{RN}} = 4$ in Experiment 3. For easy comparison, these parameter values correspond closely to the "very stable," "mildly unstable," and "very unstable" cases of Aguirregabiria and Marcoux (2019).

For each experiment, we carry out 1000 replications using two sample sizes, $N = 1600$ and $N = 6400$, noting that $N = 1600$ is the sample size used by Aguirregabiria and Mira (2007). For each replication and each sample size, we draw a sample of size N. With this

sample we calculate the iterative k-NPL and k-EPL estimates. For k-NPL, we follow the original Aguirregabiria and Mira (2007) implementation. We initialize k-NPL with estimated semiparametric logit choice probabilities. We then initialize 1-EPL using the parameter estimates and value function from the 1-NPL iteration, which is consistent even in cases where further iteration may lead to inconsistency. Aside from the initialization, and using the same sample, the k-EPL iterations proceed independently from the k-NPL iterations. For k-EPL, as before we represent the equilibrium condition in terms of v as $G(\theta, v) = v - \Phi(\theta, v)$ and use analytical derivatives for the Jacobian $\nabla_v G(\theta, v)$.

We report estimates from one, two, and three iterations of each estimator as well as the converged values, which we denote as ∞-NPL and ∞-EPL. We limit the number of iterations to 100 for both estimators, and if the algorithm has not converged we use the estimate from the final iteration. We use the same convergence criteria for both estimators: at each iteration we check the sup norm of the change in the parameter values and choice probabilities. Based on the criteria used by Aguirregabiria and Mira (2007), if both are below $10^{-2}/K$, where K is the number of parameters, we terminate the iterations and return the final converged estimate.

[Figure 1 about here.]

Figure 1 shows the Monte Carlo distributions of two key parameter estimates in the model: $\hat{\theta}_{\text{RN}}$, the competitive effect, and $\hat{\theta}_{\text{EC}}$, the entry cost. Each panel compares two histograms: the histogram above in dark gray corresponds to ∞-EPL and the histogram below in light gray is for ∞-NPL. These histograms are based on data for 1,000 replications of each experiment with 6,400 observations.

The left panels of Figure 1 show the distributions of $\hat{\theta}_{\text{RN}}$, which is related to the strength of competition in the market and is closely related with the spectral radius of the NPL operator (while the spectral radius of the EPL operator is always zero). Indeed, we can see that as the competitive effect becomes large the distribution of ∞-NPL estimates is truncated at around $\hat{\theta}_{\text{RN}} = 2.5$ in Experiment 2. It remains concentrated around $\hat{\theta}_{\text{RN}} = 2.6$ in Experiment 3 even though the value used in the data generating process was $\theta_{\text{RN}} = 4$. Yet for all experiments the distributions of ∞-EPL estimates are centered around the true values.

Although the true entry cost parameter is fixed at $\theta_{\text{EC}} = 1$ in all three experiments, the distribution of estimates can be affected when we vary the competitive effect from $\theta_{\text{RN}} = 1$ to $\theta_{\text{RN}} = 4$. In the right panels of Figure 1, the truncation from above of $\hat{\theta}_{\text{RN}}$ in the ∞-NPL distribution induces truncation from below in $\hat{\theta}_{\text{EC}}$, as lower estimated values of competition result in higher estimated entry costs, leading to distortion of both distributions and to

parameter estimates that would lead a policy-maker to possibly very different economic implications. In Experiment 3, although the distribution of ∞-NPL estimates again appear to be normally distributed for both parameters, they are biased and are not centered around the true values.

[Figure 2 about here.]

Figure 2 shows the distributions of iteration counts and computational times for ∞-EPL and ∞-NPL across the three experiments.[22] The histograms in the left panels show the distribution of iteration counts required to achieve convergence for both estimators. Recall that the maximum number of iterations allowed was 100. ∞-EPL requires fewer iterations for all experiments, especially for Experiments 2 and 3 where ∞-NPL sometimes fails to converge in Experiment 2 and always fails to converge in Experiment 3. ∞-EPL converged for all replications in all experiments.

The histograms in the right panels of Figure 2 show the distribution of total computational time, in seconds, across the 1000 replications. In Experiment 1, where k-NPL is stable, ∞-NPL is faster even though it requires more iterations on average. In other words, each k-EPL iteration is more expensive on average, but fewer are required. This is in line with the analysis in Section 3.3. However, in Experiments 2 and 3 the computational times for ∞-NPL increase as it requires more iterations, eventually overtaking the time required for ∞-EPL yet still frequently (Experiment 2) or always (Experiment 3) failing to converge.

Summary statistics for all parameter estimates and all three experiments can be found in Tables 5–10. Each table reports summary statistics for 1-, 2-, 3-, and ∞-NPL and 1-, 2-, 3-, and ∞-EPL over the 1000 Monte Carlo replications for $N = 1600$ or $N = 6400$. The upper two panels report the mean bias and mean square error (MSE) for each parameter and each sequential estimator. The next panel reports summary statistics for the number of iterations completed until either convergence or failure at 100 iterations. This includes the median, maximum, and inter-quartile range (IQR) of iteration counts across the replications as well as the number of replications for which ∞-NPL failed to converge. Finally, the bottom panel reports the computational times for each estimator including the mean and median total time (for all completed iterations) across the 1000 replications and the median time per iteration.

[Table 5 about here.]

[Table 6 about here.]

[22] Computational times are reported using Matlab R2020b on a 2019 Mac Pro with a 3.5 GHz 8-Core Intel Xeon W processor.

For Experiment 1, both the k-NPL and k-EPL estimators perform equally well with low bias, as can be seen in Tables 5 and 6. The parameter with the most finite-sample variation is also the main parameter of interest in our study: θ_{RN}. Note that in this model, in order to obtain estimates with performance similar to the converged estimates it would suffice to stop at 3 iterations with either estimator. Typically, both ∞-NPL and ∞-EPL converge in 4 to 5 iterations. However, even in this specification where k-NPL performs well, in one case out of 1000, ∞-NPL fails to converge in 100 iterations or less while ∞-EPL always converges in at most 7 iterations. In terms of computational time, in this model due to the computational complexity, in the median experiment one iteration of k-EPL is more expensive (0.3 seconds) than one iteration of k-NPL (0.045 seconds). Because roughly the same number of iterations are required in this model, the overall times for ∞-EPL are also longer than for ∞-NPL. However, even with the increased complexity each replication of ∞-EPL takes only around 1.3 seconds to estimate, so it remains quite feasible.

[Table 7 about here.]

[Table 8 about here.]

In Experiment 2, we begin to see a divergence between the two methods. As reported by Aguirregabiria and Marcoux (2019), the spectral radius of the population NPL operator is slightly larger than one for this specification. In finite random samples, sometimes the sample counterpart is stable and sometimes it is unstable. This leads to the situation illustrated by Table 7, where ∞-NPL fails to converge in 612 of 1000 replications. ∞-EPL, on the other hand, is stable and converges for all 1000 replications. Importantly, the 1-NPL estimator obtained without further iterations is always consistent. However, there is substantial bias in the ∞-NPL estimates. In an apparent contradiction, the MSE for θ_{RN} is actually lower for ∞-NPL than for ∞-EPL. This pattern of larger bias and lower MSE is seen again with the large sample size in Table 8, however it can be understood simply by recalling the histogram of the θ_{RN} estimates (panel (b) of Figure 1). The ∞-NPL sampling distribution appears to be truncated near 2.4, which perhaps not coincidentally is near the value where the spectral radius exceeds one Aguirregabiria and Marcoux (2019). Since this happens to be close to the true parameter value, the MSE is artificially low. However, the sampling distribution is neither normally distributed nor centered at the true value. A K-S test for normality of the ∞-NPL estimates has a p-value equal to zero up to three decimal places, while for ∞-EPL the p-value is 0.622.

In Experiment 2 we also see a reversal of the order of computational times: the non-convergent ∞-NPL cases require more iterations and more time per iteration, while ∞-EPL

always converges in 8 or fewer iterations. Thus, in thinking about the trade-off between robustness and computational time we should also consider convergence. A non-convergent estimator may take longer and yield worse results in the end. k-EPL does require more time per iteration in this model, but it is more robust to the strength of competition in the model.

[Table 9 about here.]

[Table 10 about here.]

Turning to Experiment 3, we increase the effect of competition even further to where $\theta_{\text{RN}} = 4$, well beyond the point where k-NPL becomes unstable. In this case, both with small samples and large samples, ∞-NPL fails to converge in all 1000 replications but ∞-EPL converges in all 1000 replications. In these cases, the bias in ∞-NPL estimates is larger and in this case and so is the MSE, since the value of θ_{RN} is farther from the point of truncation than in Experiment 2. The ∞-NPL estimator systematically underestimates the competitive effect θ_{RN} and overestimates the entry cost θ_{EC}.

Overall, across the three experiments the performance of k-EPL is stable and of similar quality despite the increasing competitive effect. This result agrees with our theoretical analysis showing that k-EPL is stable, convergent, and efficient.

Finally, we note that with the k-EPL estimator there is very little performance improvement after the first three iterations. The performance of ∞-EPL is achieved already, up to two decimal places, by 3-EPL. Thus, for this model one can reduce the computational time required while retaining efficiency and robustness by carrying out only a few iterations of k-EPL.

5 Conclusion

We proposed an iterative k-step Efficient Pseudo-Likelihood (k-EPL) estimation sequence that extends the attractive econometric and computational properties of the single-agent k-NPL sequence to games. The nice econometric properties arise because k-EPL uses Newton-like steps on the fixed point constraint at each iteration. As a result, k-EPL is stable for all regular Markov perfect equilibria, each EPL iteration has the same limiting distribution as the MLE, and further iterations achieve higher-order equivalence and quickly converge to the finite-sample MLE almost surely. Computational advantages follow from defining the equilibrium conditions with choice-specific value functions, with standard modeling assumptions reducing each EPL iteration to two steps: (i) solving linear systems to generate pseudo-regressors, followed by (ii) solving a globally concave static logit/probit maximum likelihood

problem using the pseudo-regressors. Our Monte Carlo simulations show that k-EPL performs favorably in finite samples, is robust to data-generating processes where standard k-NPL encounters serious problems, and scales better than other iterative alternatives to k-NPL.

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 is likely that one could modify the k-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.

Blevins, J. R. and M. Kim (2019). Nested pseudo likelihood estimation of continuous-time dynamic discrete games. Working paper, The Ohio State University.

Bugni, F. and J. Bunting (2020). On the iterated estimation of dynamic discrete choice games. Working Paper. arXiv:1802.06665v4.

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.

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

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. Using the definition of $\Phi(\cdot)$ from Section 2 and stacking all states for player j, we have

$$\Phi^j(a^j; v^j, v^{-j}, \theta) = u^j(a^j; \Lambda^{-j}(v^{-j}), \theta_u) + \beta F^j(a^j; \Lambda^{-j}(v^{-j}); \theta_f) S(v^j).$$

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 $\phi^j(\theta, v^j, P^{-j}) = \Phi^j_a(\theta, v^j, v^{-j})$. 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.

A.2 Proof of Lemma 2

Let $\gamma_\theta = (\theta, Y_\theta)$. Then we have

$$\Upsilon(\theta, \gamma_\theta) = Y_\theta - \left(\nabla_Y G(\theta, Y_\theta)\right)^{-1} G(\theta, Y_\theta)$$

Result 1 follows because $G(\theta, Y_\theta) = 0$. For Results 2 and 3, first consider partial derivatives of $\Upsilon(\theta, \gamma)$, evaluated at (θ, γ_θ):

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

Implicit differentiation on $G(\theta, Y_\theta) = 0$ yields

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

proving Result 2. For Result 3, we have

$$\begin{aligned}\nabla_\gamma \Upsilon(\theta,\gamma_\theta) &= I - (\nabla_Y G(\theta,Y_\theta))^{-1}\nabla_Y G(\theta,Y_\theta) - \frac{\partial \nabla_\gamma G(\gamma_\theta)^{-1}}{\partial \gamma} G(\theta,Y_\theta) \\ &= -\frac{\partial \nabla_Y G(\theta,Y_\theta)^{-1}}{\partial Y} G(\theta,Y_\theta) \\ &= 0,\end{aligned}$$

where the last equality arises because $G(\theta, Y_\theta) = 0$.

A.3 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.[23] 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 2.

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

A.3.1 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 $\Upsilon(\theta,\gamma)$ and the Mann-Wald theorem then give almost sure convergence of $\hat{Y}_k$ to Y^*.

A.3.2 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 $\Upsilon(\cdot)$ defined in Algorithm

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

1, 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

$$\nabla_\gamma\Upsilon(\theta,\gamma) = I - \nabla_Y G(\breve{\theta},Y)^{-1}\nabla_Y G(\theta,Y) - \begin{bmatrix} \frac{\partial \nabla_Y G(\breve{\theta},Y)^{-1}}{\partial \breve{\theta}} \\ \frac{\partial \nabla_Y G(\breve{\theta},Y)^{-1}}{\partial Y}\end{bmatrix} G(\theta,Y). \tag{1}$$

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 $\nabla_{\theta\theta'}\tilde{Q}_N(\bar{\theta},\bar{\gamma}) \xrightarrow{a.s.} -\Omega^*_{\theta\theta}$ 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 $\hat{Y}_{\mathrm{MLE}}$ follow from asymptotic equivalence of $\hat{\theta}_k$ and $\hat{\theta}_{\mathrm{MLE}}$ and the properties of Υ. Strong $\sqrt{N}$-consistency of $\hat{\gamma}_0$ completes the proof by induction.

A.3.3 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 $\nabla_{\theta\gamma}\tilde{Q}^*(\theta^*,\gamma^*) = 0$, so the spectral radius is also zero.

A.4 Proof of Theorem 2

By examining the first-order conditions, we can see that $\hat{\theta}_{\mathrm{MLE}} = \tilde{\theta}_N(\hat{\gamma}_{\mathrm{MLE}})$, so that $\hat{Y}_{\mathrm{MLE}} = \Upsilon(\tilde{\theta}_N(\hat{\gamma}_{\mathrm{MLE}}), \hat{\gamma}_{\mathrm{MLE}})$. 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 $\hat{\gamma}_{\mathrm{MLE}}$. First we have

$$\nabla_\gamma H_{1,N}(\hat{\gamma}_{\mathrm{MLE}}) = \nabla_\gamma \tilde{\theta}_N(\hat{\gamma}_{\mathrm{MLE}}) \xrightarrow{a.s.} \nabla_\gamma \tilde{\theta}^*(\gamma^*) = 0 \tag{2}$$

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

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

recalling that $\tilde{\theta}_N(\hat{\gamma}_{\mathrm{MLE}}) = \hat{\theta}_{\mathrm{MLE}}$. The first term converges to zero in probability as in (2). The second term is zero due to the zero Jacobian property of Υ, as can be seen by evaluating (1).

The above analysis implies that $\hat{\gamma}_{\mathrm{MLE}} = H_N(\hat{\gamma}_{\mathrm{MLE}})$ and $\nabla_\gamma H_N(\hat{\gamma}_{\mathrm{MLE}}) \xrightarrow{a.s.} 0$, which are

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

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

It follows that

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

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

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 11 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 2.

[Table 11 about here.]

B.1 Pesendorfer and Schmidt-Dengler (2010): Static Game, Unstable Equilibrium

As a simple illustration of the performance of k-EPL, we consider estimating the static game ($\beta = 0$) of incomplete information from Pesendorfer and Schmidt-Dengler (2010). This example is particularly interesting because it was constructed as an example where inconsistency of ∞-NPL can be shown analytically. 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 $F_\alpha(v) = v$ for $v \in [\alpha, 1-\alpha)$ and a more complicated form for $v \in \mathbb{R} \setminus [\alpha, 1-\alpha)$ 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 in P space:

$$\begin{aligned}\begin{bmatrix} P^1 \\ P^2 \end{bmatrix} &= \begin{bmatrix} 1+\theta P^2 \\ 1+\theta P^1 \end{bmatrix} \\ &= \begin{bmatrix} 1 \\ 1 \end{bmatrix} + \begin{bmatrix} 0 & \theta \\ \theta & 0 \end{bmatrix} \begin{bmatrix} P^1 \\ P^2 \end{bmatrix},\end{aligned}$$

or more compactly, $P = \Psi(\theta, P)$, so this can be used in k-NPL. This linear system has a unique solution if and only if $\theta \neq -1$, and the solution is $P^1 = P^2 = \frac{1}{1-\theta}$. Pesendorfer and Schmidt-Dengler (2010) consider $\theta^* = -2$, so that $P_1^* = P_2^* = \frac{1}{3}$. 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, this type of adjustment may not induce convergence in more general settings.

Consider, instead, estimating θ^* with k-EPL, with a change of variable to v space. Noting that $P^{-j} = F_\alpha(v^{-j})$, in the equilibrium we have $v^j = \theta F_\alpha(v^{-j})$ where $F_\alpha(v^{-j}) = v^{-j}$ in the region of interest. The fixed-point equation then reduces to:

$$\begin{bmatrix} v^1 \\ v^2 \end{bmatrix} = \begin{bmatrix} (1+v^2)\theta \\ (1+v^1)\theta \end{bmatrix}.$$

So, we can define $Y \equiv v = (v^1, v^2)$ and therefore

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

Because θ is a scalar, $G(\theta, v)$ is linear in θ and v separately (holding the other fixed) but not jointly. Linearity in v is important because $G(\theta, v) = 0$ can be solved analytically via a linear system. So, we can easily compute the finite-sample MLE via nested fixed point as well as via EPL iterations. Additionally, we see that $\nabla_v G(\theta, v) = I - A\theta$ and we can easily verify that this is invertible if and only if $\theta \neq -1$. And since $v^j = \theta P^{-j}$, we also have $q_i(\theta, v) = q_i(v)$, so that θ only influences q_i through $v(\theta)$. This modification is made without loss of generality in full MLE subject to the equilibrium constraint, so it is also valid here. For k-EPL, we have $\Upsilon(\theta, \hat{\gamma}_{k-1}) = \hat{v}_{k-1} - \nabla_v G(\hat{\theta}_{k-1}, \hat{v}_{k-1})^{-1} G(\theta, \hat{v}_{k-1})$, as in Algorithm (1).

All that remains now is to obtain $\hat{v}_0$ 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 $\hat{P}_0^1$ and $\hat{P}_0^2$. We then use these to construct

$$\begin{aligned} \hat{\theta}_0 &= \frac{\frac{\hat{P}_0^1 - 1}{\hat{P}_0^2} + \frac{\hat{P}_0^2 - 1}{\hat{P}_0^1}}{2}, \\ \hat{v}_0^j &= \hat{\theta}_0 \hat{P}_0^{-j}. \end{aligned}$$

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.[24] The results are summarized in Table 12. 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 12 about here.]

On the other hand, ∞-NPL performs poorly as expected since this model was constructed as an example where ∞-NPL is inconsistent. The estimate has a mean of -1.0342 and MSE

[24] By ∞-EPL and ∞-NPL, we mean that we iterate until $||\hat{\theta}_k - \hat{\theta}_{k-1}||_\infty < 10^{-6}$ or k reaches 20. Estimation was performed with Matlab R2017a using `fmincon`. The default tolerance of 10^{-6} is used for the solver.

of 0.9651. Almost all of the MSE is due to the asymptotic bias, so the estimate is reliably converging to the wrong number.[25]

C Alternative Choices of $\Upsilon(\cdot)$

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

Proof. All of the listed $\Upsilon(\cdot)$ functions satisfy the zero Jacobian property, so the results from the proofs of Theorems 1 and 2 carry through. □

The first definition of Υ in the Theorem 3 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.[26] 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, 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 introduce additional nonlinearities in the objective function for optimization.[27] However, we include it for completeness. For all of the definitions of Υ in the theorem, the results of Lemma 2 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.

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

[26] 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.

[27] 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.

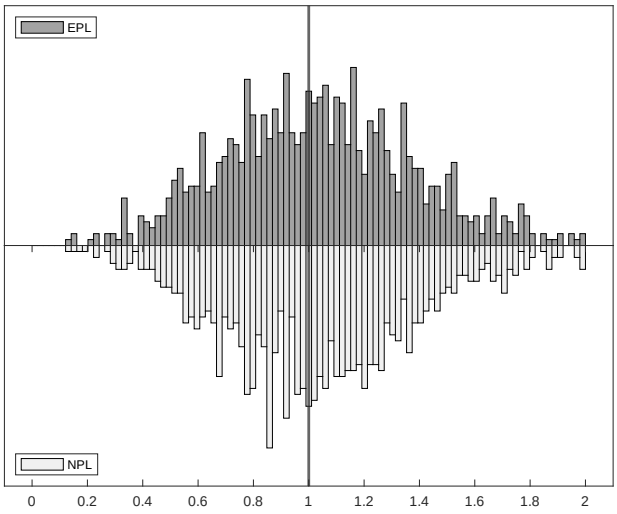

(a) Experiment 1: $\hat{\theta}_{\mathrm{RN}}$ when $\theta_{\mathrm{RN}} = 1$

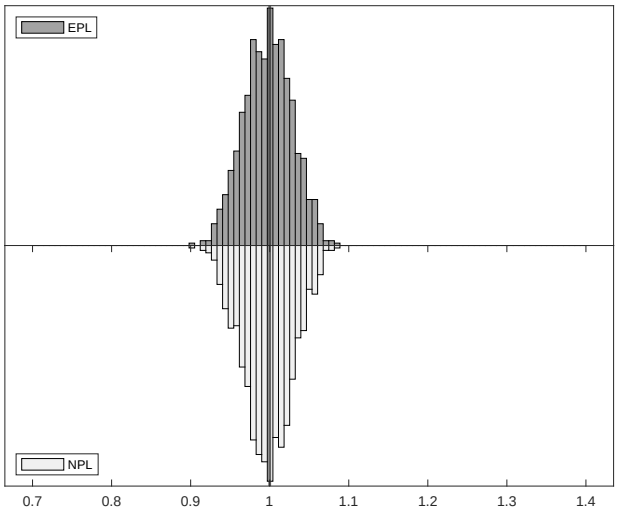

(d) Experiment 1: $\hat{\theta}_{\mathrm{EC}}$ when $\theta_{\mathrm{RN}} = 1$

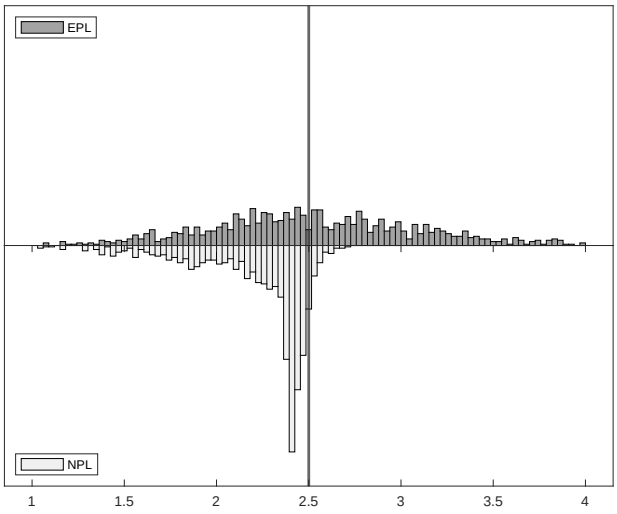

(b) Experiment 2: $\hat{\theta}_{\mathrm{RN}}$ when $\theta_{\mathrm{RN}} = 2.5$

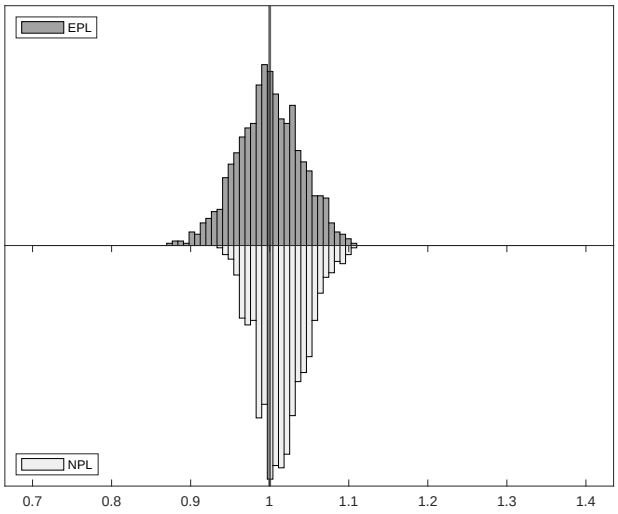

(e) Experiment 2: $\hat{\theta}_{\mathrm{EC}}$ when $\theta_{\mathrm{RN}} = 2.5$

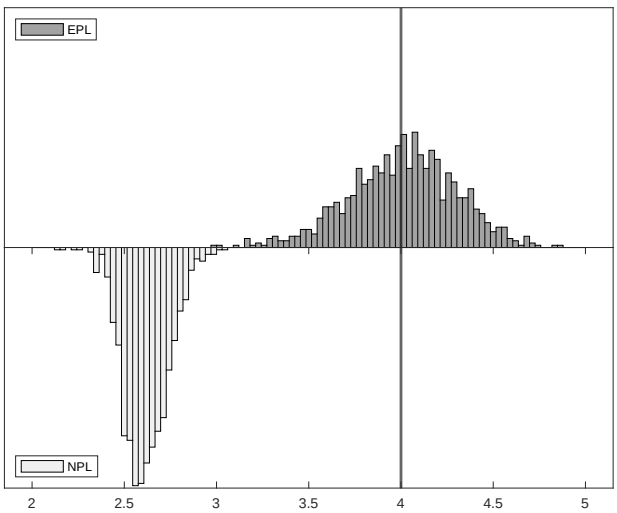

(c) Experiment 3: $\hat{\theta}_{\mathrm{RN}}$ when $\theta_{\mathrm{RN}} = 4$

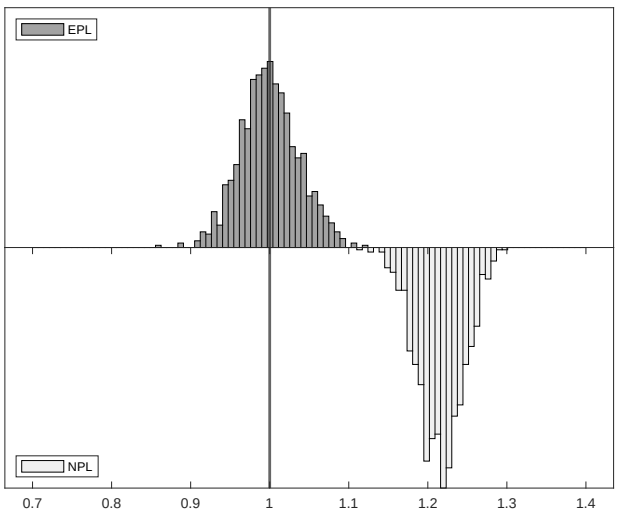

(f) Experiment 3: $\hat{\theta}_{\mathrm{EC}}$ when $\theta_{\mathrm{RN}} = 4$

Figure 1: Distributions of $\hat{\theta}_{\mathrm{RN}}$ and $\hat{\theta}_{\mathrm{EC}}$ for ∞-EPL and ∞-NPL Over 1000 Replications ($N = 6400$)

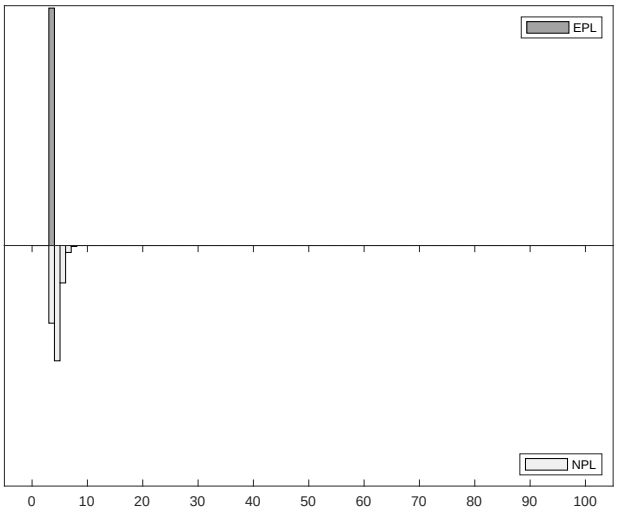

(a) Experiment 1: Iterations

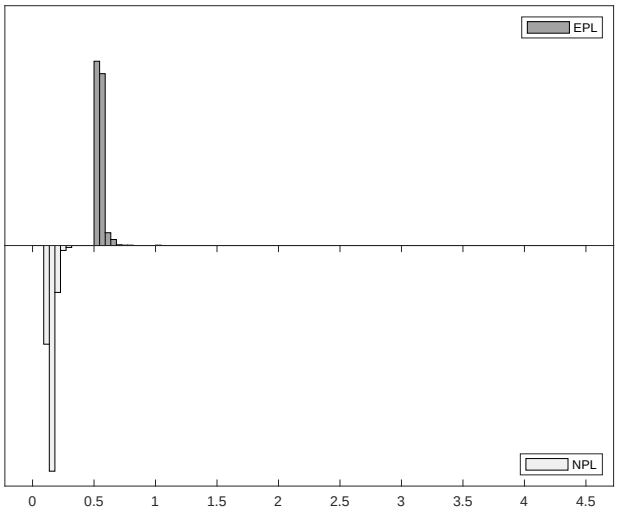

(d) Experiment 1: Computational Times (sec.)

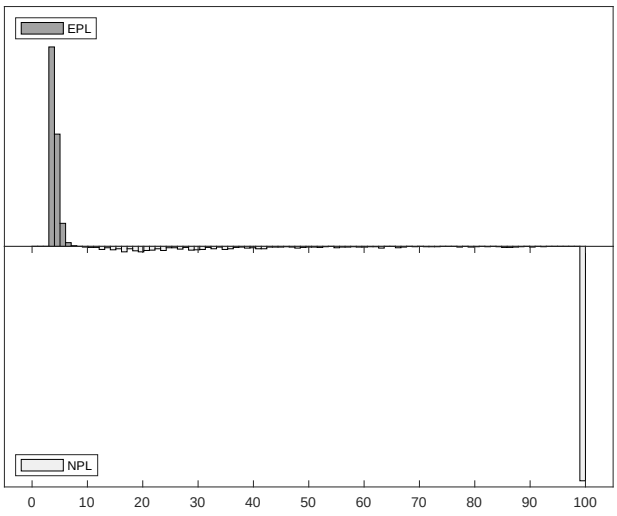

(b) Experiment 2: Iterations

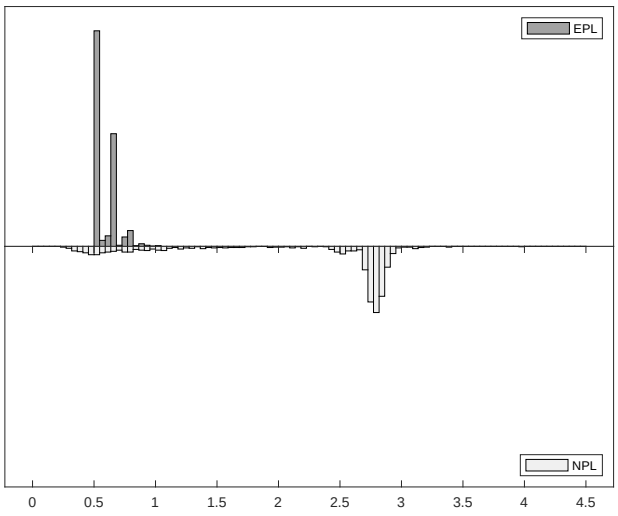

(e) Experiment 2: Computational Times (sec.)

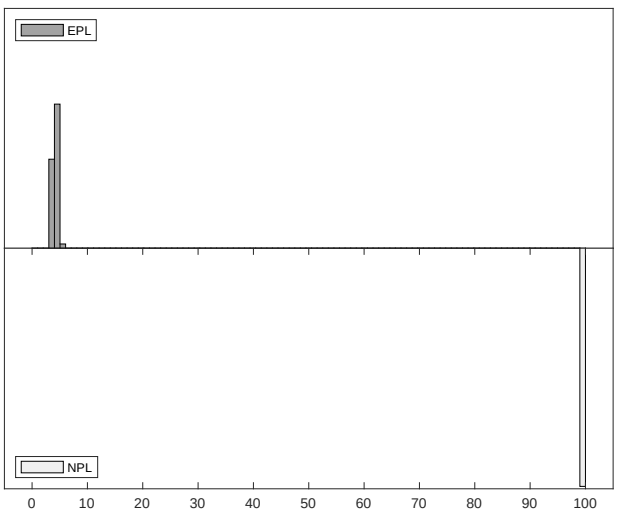

(c) Experiment 3: Iterations

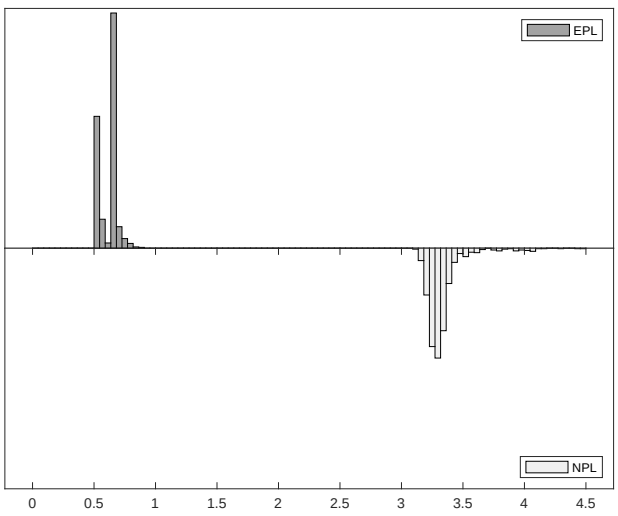

(f) Experiment 3: Computational Times (sec.)

Figure 2: Distributions of Iterations and Computational Time for ∞-EPL and ∞-NPL Over 1000 Replications ($N = 6400$)

Table 1: 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

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

Table 3: 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

Table 4: 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

Bias	True	1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
$\theta_{\text{FC},1}$	-1.9	-0.018	0.007	0.008	0.004	0.002	0.004	0.004	0.005
$\theta_{\text{FC},2}$	-1.8	-0.016	0.007	0.007	0.003	0.002	0.004	0.003	0.004
$\theta_{\text{FC},3}$	-1.7	-0.018	0.003	0.003	-0.001	-0.002	0.000	-0.001	0.000
$\theta_{\text{FC},4}$	-1.6	-0.015	0.003	0.004	0.000	-0.001	0.001	0.000	0.001
$\theta_{\text{FC},5}$	-1.5	-0.013	0.003	0.004	0.000	-0.001	0.000	0.000	0.000
θ_{RS}	1.0	-0.018	0.010	0.010	0.013	0.014	0.015	0.014	0.015
θ_{RN}	1.0	-0.068	0.034	0.034	0.038	0.040	0.044	0.041	0.044
θ_{EC}	1.0	-0.000	0.002	0.003	-0.001	-0.002	-0.001	-0.001	-0.001
MSE	True	1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
$\theta_{\text{FC},1}$	-1.9	0.012	0.013	0.013	0.013	0.013	0.013	0.013	0.013
$\theta_{\text{FC},2}$	-1.8	0.012	0.012	0.012	0.012	0.012	0.012	0.012	0.012
$\theta_{\text{FC},3}$	-1.7	0.011	0.012	0.012	0.012	0.012	0.012	0.012	0.012
$\theta_{\text{FC},4}$	-1.6	0.010	0.011	0.011	0.011	0.011	0.011	0.011	0.011
$\theta_{\text{FC},5}$	-1.5	0.009	0.010	0.010	0.009	0.009	0.009	0.009	0.009
θ_{RS}	1.0	0.010	0.013	0.013	0.013	0.014	0.014	0.014	0.014
θ_{RN}	1.0	0.091	0.123	0.123	0.128	0.128	0.130	0.129	0.131
θ_{EC}	1.0	0.004	0.004	0.004	0.004	0.004	0.004	0.004	0.004
Iterations		1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
Median		1	1	2	2	3	3	5	4
Max		1	1	2	2	3	3	100	7
IQR								2	0
Non-Conv.								0.1%	0%
Time (sec.)		1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
Total		17.88	118.42	31.33	234.82	44.35	350.07	76.55	481.70
Mean		0.02	0.12	0.03	0.23	0.04	0.35	0.08	0.48
Median		0.02	0.12	0.03	0.23	0.04	0.35	0.07	0.46
Med./Iter.		0.018	0.118	0.016	0.117	0.015	0.116	0.014	0.116

Table 5: Monte Carlo Results for Aguirregabiria and Mira (2007): Experiment 1 ($\theta_{\text{RN}} = 1$), Small Sample ($N = 1600$)

Bias	True	1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
$\theta_{\text{FC},1}$	-1.9	-0.018	0.003	0.003	0.000	-0.001	0.000	0.000	0.000
$\theta_{\text{FC},2}$	-1.8	-0.017	0.002	0.002	-0.001	-0.002	-0.001	-0.001	-0.001
$\theta_{\text{FC},3}$	-1.7	-0.016	0.001	0.001	-0.002	-0.002	-0.001	-0.001	-0.001
$\theta_{\text{FC},4}$	-1.6	-0.013	0.003	0.003	-0.000	-0.001	-0.000	-0.000	-0.000
$\theta_{\text{FC},5}$	-1.5	-0.012	0.001	0.001	-0.002	-0.003	-0.002	-0.002	-0.002
θ_{RS}	1.0	-0.022	0.003	0.003	0.004	0.005	0.005	0.005	0.005
θ_{RN}	1.0	-0.078	0.010	0.010	0.011	0.012	0.013	0.013	0.013
θ_{EC}	1.0	-0.001	0.000	0.001	-0.002	-0.003	-0.002	-0.002	-0.002
MSE	True	1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
$\theta_{\text{FC},1}$	-1.9	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003
$\theta_{\text{FC},2}$	-1.8	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003
$\theta_{\text{FC},3}$	-1.7	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003
$\theta_{\text{FC},4}$	-1.6	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003
$\theta_{\text{FC},5}$	-1.5	0.003	0.003	0.003	0.002	0.003	0.002	0.003	0.002
θ_{RS}	1.0	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003
θ_{RN}	1.0	0.027	0.030	0.030	0.030	0.030	0.030	0.031	0.030
θ_{EC}	1.0	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
Iterations		1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
Median		1	1	2	2	3	3	5	4
Max		1	1	2	2	3	3	8	4
IQR								1	0
Non-Conv.								0%	0%
Time (sec.)		1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
Total		43.91	148.28	75.66	284.75	105.45	416.53	154.35	548.53
Mean		0.04	0.15	0.08	0.28	0.11	0.42	0.15	0.55
Median		0.04	0.15	0.07	0.28	0.10	0.41	0.15	0.55
Med./Iter.		0.043	0.147	0.037	0.141	0.034	0.138	0.031	0.136

Table 6: Monte Carlo Results for Aguirregabiria and Mira (2007): Experiment 1 ($\theta_{\text{RN}} = 1$), Large Sample ($N = 6400$)

Bias	True	1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
$\theta_{\text{FC},1}$	-1.9	-0.005	-0.002	-0.002	-0.001	-0.008	-0.003	0.006	-0.003
$\theta_{\text{FC},2}$	-1.8	-0.002	-0.004	-0.005	-0.002	-0.005	-0.004	0.010	-0.004
$\theta_{\text{FC},3}$	-1.7	0.003	-0.006	-0.007	-0.001	-0.000	-0.004	0.014	-0.004
$\theta_{\text{FC},4}$	-1.6	0.007	-0.008	-0.009	-0.002	0.002	-0.006	0.016	-0.006
$\theta_{\text{FC},5}$	-1.5	0.011	-0.009	-0.010	-0.002	0.005	-0.007	0.017	-0.008
θ_{RS}	1.0	0.012	0.007	0.006	-0.006	-0.006	0.007	-0.063	0.008
θ_{RN}	1.0	0.056	0.024	0.017	-0.026	-0.027	0.025	-0.261	0.030
θ_{EC}	1.0	-0.002	-0.001	-0.001	0.004	0.001	-0.002	0.027	-0.003
MSE	True	1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
$\theta_{\text{FC},1}$	-1.9	0.016	0.015	0.015	0.015	0.016	0.015	0.014	0.015
$\theta_{\text{FC},2}$	-1.8	0.016	0.015	0.015	0.014	0.015	0.015	0.013	0.015
$\theta_{\text{FC},3}$	-1.7	0.016	0.015	0.015	0.014	0.015	0.015	0.013	0.015
$\theta_{\text{FC},4}$	-1.6	0.016	0.015	0.015	0.014	0.015	0.015	0.012	0.015
$\theta_{\text{FC},5}$	-1.5	0.018	0.017	0.017	0.016	0.016	0.016	0.013	0.017
θ_{RS}	1.0	0.029	0.023	0.023	0.019	0.018	0.021	0.010	0.021
θ_{RN}	1.0	0.468	0.365	0.359	0.310	0.282	0.327	0.162	0.335
θ_{EC}	1.0	0.008	0.007	0.007	0.006	0.006	0.006	0.004	0.006
Iterations		1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
Median		1	1	2	2	3	3	100	5
Max		1	1	2	2	3	3	100	30
IQR								75.5	2
Non-Conv.								61.2%	0%
Time (sec.)		1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
Total		18.26	120.53	32.92	239.61	47.60	358.72	1030.36	672.18
Mean		0.02	0.12	0.03	0.24	0.05	0.36	1.03	0.67
Median		0.02	0.12	0.03	0.24	0.05	0.36	1.44	0.60
Med./Iter.		0.018	0.120	0.016	0.119	0.016	0.119	0.014	0.118

Table 7: Monte Carlo Results for Aguirregabiria and Mira (2007): Experiment 2 ($\theta_{\text{RN}} = 2.5$), Small Sample ($N = 1600$)

Bias	True	1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
$\theta_{\text{FC},1}$	-1.9	-0.002	-0.000	-0.001	-0.002	-0.005	-0.002	0.006	-0.002
$\theta_{\text{FC},2}$	-1.8	0.002	-0.002	-0.003	-0.002	-0.002	-0.002	0.008	-0.002
$\theta_{\text{FC},3}$	-1.7	0.006	-0.004	-0.005	-0.002	0.000	-0.003	0.009	-0.003
$\theta_{\text{FC},4}$	-1.6	0.012	-0.004	-0.005	-0.001	0.004	-0.002	0.010	-0.002
$\theta_{\text{FC},5}$	-1.5	0.019	-0.004	-0.005	-0.001	0.007	-0.003	0.008	-0.003
θ_{RS}	1.0	-0.001	0.001	0.001	-0.003	-0.005	0.001	-0.040	0.001
θ_{RN}	1.0	0.003	0.002	-0.003	-0.013	-0.018	0.001	-0.163	0.002
θ_{EC}	1.0	0.003	-0.000	0.000	0.000	0.000	-0.001	0.015	-0.001
MSE	True	1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
$\theta_{\text{FC},1}$	-1.9	0.004	0.004	0.004	0.004	0.004	0.004	0.004	0.004
$\theta_{\text{FC},2}$	-1.8	0.004	0.004	0.004	0.004	0.004	0.004	0.003	0.004
$\theta_{\text{FC},3}$	-1.7	0.004	0.004	0.004	0.004	0.004	0.004	0.003	0.004
$\theta_{\text{FC},4}$	-1.6	0.004	0.004	0.004	0.004	0.004	0.004	0.003	0.004
$\theta_{\text{FC},5}$	-1.5	0.004	0.004	0.004	0.004	0.004	0.004	0.004	0.004
θ_{RS}	1.0	0.007	0.006	0.006	0.005	0.005	0.006	0.003	0.006
θ_{RN}	1.0	0.109	0.096	0.095	0.085	0.084	0.087	0.048	0.088
θ_{EC}	1.0	0.002	0.002	0.002	0.002	0.002	0.002	0.001	0.002
Iterations		1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
Median		1	1	2	2	3	3	100	4
Max		1	1	2	2	3	3	100	8
IQR								52	1
Non-Conv.								68.9%	0%
Time (sec.)		1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
Total		35.34	135.62	64.32	266.26	92.87	395.24	2186.69	586.21
Mean		0.04	0.14	0.06	0.27	0.09	0.40	2.19	0.59
Median		0.03	0.13	0.06	0.26	0.09	0.39	2.75	0.53
Med./Iter.		0.034	0.135	0.032	0.132	0.031	0.131	0.028	0.130

Table 8: Monte Carlo Results for Aguirregabiria and Mira (2007): Experiment 2 ($\theta_{\text{RN}} = 2.5$), Large Sample ($N = 6400$)

Bias	True	1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
$\theta_{\text{FC},1}$	-1.9	-0.100	0.010	0.008	-0.002	-0.095	0.002	0.012	0.002
$\theta_{\text{FC},2}$	-1.8	-0.088	0.000	-0.001	-0.000	-0.074	0.002	0.031	0.002
$\theta_{\text{FC},3}$	-1.7	-0.066	-0.010	-0.011	0.000	-0.040	0.001	0.062	0.002
$\theta_{\text{FC},4}$	-1.6	-0.016	-0.016	-0.017	0.004	0.026	0.002	0.131	0.002
$\theta_{\text{FC},5}$	-1.5	0.050	-0.007	-0.005	0.009	0.156	0.000	0.211	-0.000
θ_{RS}	1.0	0.042	-0.005	-0.019	-0.009	-0.108	0.001	-0.244	0.001
θ_{RN}	1.0	0.198	-0.067	-0.147	-0.055	-0.640	0.004	-1.382	0.008
θ_{EC}	1.0	-0.009	0.025	0.031	0.010	0.073	0.001	0.218	-0.000
MSE	True	1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
$\theta_{\text{FC},1}$	-1.9	0.038	0.022	0.022	0.023	0.032	0.023	0.019	0.023
$\theta_{\text{FC},2}$	-1.8	0.034	0.020	0.020	0.021	0.026	0.021	0.018	0.021
$\theta_{\text{FC},3}$	-1.7	0.030	0.020	0.019	0.019	0.021	0.020	0.020	0.020
$\theta_{\text{FC},4}$	-1.6	0.028	0.019	0.019	0.019	0.019	0.019	0.032	0.019
$\theta_{\text{FC},5}$	-1.5	0.039	0.022	0.021	0.019	0.041	0.019	0.057	0.019
θ_{RS}	1.0	0.022	0.006	0.006	0.004	0.015	0.004	0.061	0.004
θ_{RN}	1.0	0.602	0.126	0.131	0.090	0.456	0.086	1.924	0.086
θ_{EC}	1.0	0.022	0.007	0.007	0.005	0.009	0.005	0.051	0.005
Iterations		1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
Median		1	1	2	2	3	3	100	5
Max		1	1	2	2	3	3	100	9
IQR								0	1
Non-Conv.								100%	0%
Time (sec.)		1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
Total		18.64	125.88	33.88	246.85	48.96	367.25	1512.68	654.05
Mean		0.02	0.13	0.03	0.25	0.05	0.37	1.51	0.65
Median		0.02	0.12	0.03	0.24	0.05	0.36	1.48	0.62
Med./Iter.		0.018	0.124	0.017	0.122	0.016	0.120	0.015	0.119

Table 9: Monte Carlo Results for Aguirregabiria and Mira (2007): Experiment 3 ($\theta_{\text{RN}} = 4$), Small Sample ($N = 1600$)

Bias	True	1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
$\theta_{\text{FC},1}$	-1.9	-0.098	0.008	0.007	0.000	-0.098	0.000	0.013	0.000
$\theta_{\text{FC},2}$	-1.8	-0.086	-0.003	-0.003	0.000	-0.078	0.000	0.031	0.000
$\theta_{\text{FC},3}$	-1.7	-0.063	-0.015	-0.016	0.001	-0.045	-0.001	0.062	-0.001
$\theta_{\text{FC},4}$	-1.6	-0.013	-0.024	-0.025	0.004	0.017	0.000	0.130	0.000
$\theta_{\text{FC},5}$	-1.5	0.057	-0.025	-0.023	0.009	0.139	0.000	0.207	0.000
θ_{RS}	1.0	0.034	0.011	-0.003	-0.010	-0.096	0.000	-0.243	0.000
θ_{RN}	1.0	0.155	0.024	-0.054	-0.057	-0.570	-0.001	-1.377	0.000
θ_{EC}	1.0	0.001	0.011	0.016	0.008	0.060	0.000	0.216	0.000
MSE	True	1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
$\theta_{\text{FC},1}$	-1.9	0.016	0.005	0.005	0.005	0.015	0.005	0.005	0.005
$\theta_{\text{FC},2}$	-1.8	0.014	0.005	0.005	0.005	0.011	0.005	0.005	0.005
$\theta_{\text{FC},3}$	-1.7	0.011	0.005	0.005	0.005	0.007	0.005	0.008	0.005
$\theta_{\text{FC},4}$	-1.6	0.007	0.006	0.006	0.005	0.005	0.005	0.020	0.005
$\theta_{\text{FC},5}$	-1.5	0.012	0.006	0.006	0.005	0.024	0.005	0.046	0.005
θ_{RS}	1.0	0.006	0.002	0.001	0.001	0.010	0.001	0.060	0.001
θ_{RN}	1.0	0.167	0.037	0.034	0.023	0.335	0.023	1.900	0.023
θ_{EC}	1.0	0.005	0.002	0.002	0.001	0.004	0.001	0.048	0.001
Iterations		1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
Median		1	1	2	2	3	3	100	5
Max		1	1	2	2	3	3	100	6
IQR								0	1
Non-Conv.								100%	0%
Time (sec.)		1-NPL	1-EPL	2-NPL	2-EPL	3-NPL	3-EPL	∞-NPL	∞-EPL
Total		40.48	145.93	75.22	278.93	108.36	410.64	3332.77	620.36
Mean		0.04	0.15	0.08	0.28	0.11	0.41	3.33	0.62
Median		0.04	0.14	0.07	0.28	0.11	0.41	3.30	0.65
Med./Iter.		0.039	0.145	0.037	0.138	0.035	0.136	0.033	0.133

Table 10: Monte Carlo Results for Aguirregabiria and Mira (2007): Experiment 3 ($\theta_{\text{RN}} = 4$), Large Sample ($N = 6400$)

Table 11: Monte Carlo Results for Pesendorfer and Schmidt-Dengler (2008) Equilibrium (iii) with $N = 1000$

Parameter	Statistic	1-NPL	1-EPL	∞-NPL	∞-EPL
$\theta_{\text{M}} = 1.2$	Mean Bias	-0.0419	-0.0383	-0.2099	-0.0003
	MSE	0.0204	0.0194	0.0480	0.0174
$\theta_{\text{C}} = -2.4$	Mean Bias	0.0948	0.0625	0.6806	0.0043
	MSE	0.1127	0.1000	0.4683	0.0987
$\theta_{\text{EC}} = -0.2$	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

Table 12: 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