

DISCUSSION PAPER SERIES

DP14007

**IMPOSING EQUILIBRIUM
RESTRICTIONS IN THE ESTIMATION OF
DYNAMIC DISCRETE GAMES**

Victor Aguirregabiria and Mathieu Marcoux

INDUSTRIAL ORGANIZATION



IMPOSING EQUILIBRIUM RESTRICTIONS IN THE ESTIMATION OF DYNAMIC DISCRETE GAMES

Victor Aguirregabiria and Mathieu Marcoux

Discussion Paper DP14007
Published 18 September 2019
Submitted 17 September 2019

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Victor Aguirregabiria and Mathieu Marcoux

IMPOSING EQUILIBRIUM RESTRICTIONS IN THE ESTIMATION OF DYNAMIC DISCRETE GAMES

Abstract

Imposing equilibrium restrictions provides substantial gains in the estimation of dynamic discrete games. Estimation algorithms imposing these restrictions -- MPEC, NFXP, NPL, and variations -- have different merits and limitations. MPEC guarantees local convergence, but requires the computation of high-dimensional Jacobians. The NPL algorithm avoids the computation of these matrices, but -- in games -- may fail to converge to the consistent NPL estimator. We study the asymptotic properties of the NPL algorithm treating the iterative procedure as performed in finite samples. We find that there are always samples for which the algorithm fails to converge, and this introduces a selection bias. We also propose a spectral algorithm to compute the NPL estimator. This algorithm satisfies local convergence and avoids the computation of Jacobian matrices. We present simulation evidence illustrating our theoretical results and the good properties of the spectral algorithm.

JEL Classification: C13, C57, C61, C73

Keywords: Dynamic discrete games, convergence, Convergence selection bias, Nested pseudo-likelihood, Spectral algorithms

Victor Aguirregabiria - victor.aguirregabiria@utoronto.ca
University of Toronto and CEPR

Mathieu Marcoux - mathieu.marcoux@umontreal.ca
Université de Montréal

Acknowledgements

We thank Christian Gouriéroux and Yuanyuan Wan for helpful comments. We also thank Martin Burda, Limin Fang, Ashique Habib, Min Seong Kim, Adam Lavecchia, Ismael Mourifié, Thomas Russell and Eduardo Souza-Rodrigues for insightful discussions. The paper benefited from comments from seminar participants at the University of Toronto, the Empirical Microeconomics Workshop (Banff), the CEA Annual Meetings (Vancouver), the CIREQ PhD Students Conference (Montreal), and the Doctoral Workshop in Applied Economics (Toronto). Financial support from Social Sciences and Humanities Research Council (SSHRC) and Ontario Graduate Scholarship (OGS) is gratefully acknowledged.

Imposing equilibrium restrictions in the estimation of dynamic discrete games*

Victor Aguirregabiria[†]
University of Toronto, CEPR

Mathieu Marcoux[‡]
Université de Montréal, CIREQ, CIRANO

September 16, 2019

Abstract

Imposing equilibrium restrictions provides substantial gains in the estimation of dynamic discrete games. Estimation algorithms imposing these restrictions – MPEC, NFXP, NPL, and variations – have different merits and limitations. MPEC guarantees local convergence, but requires the computation of high-dimensional Jacobians. The NPL algorithm avoids the computation of these matrices, but – in games – may fail to converge to the consistent NPL estimator. We study the asymptotic properties of the NPL algorithm treating the iterative procedure as performed in finite samples. We find that there are always samples for which the algorithm fails to converge, and this introduces a selection bias. We also propose a spectral algorithm to compute the NPL estimator. This algorithm satisfies local convergence and avoids the computation of Jacobian matrices. We present simulation evidence illustrating our theoretical results and the good properties of the spectral algorithm.

Keywords: Dynamic discrete game; nested pseudo-likelihood; fixed point algorithms; convergence; convergence selection bias.

JEL codes: C13; C57; C61; C73.

*We thank Christian Gouriéroux and Yuanyuan Wan for helpful comments. We also thank Martin Burda, Limin Fang, Ashique Habib, Min Seong Kim, Adam Lavecchia, Ismael Mourifié, Thomas Russell and Eduardo Souza-Rodrigues for insightful discussions. The paper benefited from comments from seminar participants at the University of Toronto, the Empirical Microeconomics Workshop (Banff), the CEA Annual Meetings (Vancouver), the CIREQ PhD Students Conference (Montreal), and the Doctoral Workshop in Applied Economics (Toronto). Financial support from Social Sciences and Humanities Research Council (SSHRC) and Ontario Graduate Scholarship (OGS) is gratefully acknowledged.

[†]150 St. George Street, Toronto, ON, M5S 3G7, Canada, victor.aguirregabiria@utoronto.ca.

[‡]C.P. 6128, succ. Centre-Ville, Montreal, QC, H3C 3J7, Canada, mathieu.marcoux@umontreal.ca.

1 Introduction

Economic theory often delivers models in which the distribution of agents' decisions is implicitly defined as a fixed point. For instance, this is the case in models where agents' decisions are interpreted via equilibrium conditions based on rational expectations. In this class of models, there are typically two different objects of interest: the conditional choice probabilities (CCPs), which are a probabilistic representation of agents' behavior, and the structural parameters of the model. In order to reduce the computational burden of estimating such models, it is common practice to first estimate the CCPs in a flexible way and then to estimate the structural parameters conditional on these estimates. While consistent in large samples, this two-step procedure is known to suffer from a few drawbacks. In particular, biased and imprecise estimates in the first step may affect the properties of the second-step estimator. Furthermore, consistent first-step estimates may not be available, as it is often the case if choice probabilities depend on unobserved heterogeneity.

These drawbacks motivated the Nested Pseudo Likelihood (NPL) estimator proposed by [Aguirregabiria and Mira \(2002; 2007\)](#): a fixed point of the so-called NPL mapping which embeds both the pseudo maximum likelihood estimation of the structural parameters and the equilibrium restrictions for the CCPs. As a method to obtain the NPL estimator, [Aguirregabiria and Mira \(2002; 2007\)](#) proposed the NPL algorithm, which iteratively updates the estimated CCPs using the NPL mapping.

The NPL algorithm is a fixed point iterative method. Therefore, this algorithm is expected to fail to converge to the consistent estimator when the data come from an unstable fixed point (e.g., [Pesendorfer and Schmidt-Dengler, 2008; 2010](#)). Actually, Monte Carlo studies have shown that this algorithm may fail to converge even when the data generating process satisfies some stability conditions ([Egedal, Lai, and Su, 2015](#)).

In this paper, we take another look at the sequential estimation of dynamic discrete games in order to bridge the literature's incomplete understanding of the NPL algorithm's failure to converge. We depart from existing works in two ways.

First, when studying the asymptotic properties of the NPL algorithm, the updating procedure is assumed to be performed in finite samples. Such an argument highlights how the convergence properties of the NPL algorithm are related to the stability of the stochastic mapping used to update the choice probabilities, rather than the stability of the equilibrium generating the data. In other words, the convergence of the NPL algorithm is driven by the stability of the NPL mapping in the sample, as opposed to the NPL mapping in the population. This feature of our asymptotic approach – though apparently subtle – has important implications. It explains why the NPL algorithm may fail to converge even if the data is generated from a stable equilibrium.

Second, we present a new analysis of the implications of the fixed point algorithm's

lack of convergence on finite sample and asymptotic properties of the NPL estimator. This estimator is defined upon convergence of the fixed point iterations in the sample. This is a stochastic condition which does not have probability one, even when the data have been generated from a stable equilibrium. In contrast, previous studies consider convergence of the algorithm in the population. Using our approach we can show that the possibility of no convergence generates a selection problem in the estimator, which we refer to as *convergence selection*. Random samples for which the algorithm converges have statistical properties that are different from samples where convergence fails. That is, convergence introduces selection issues. We characterize the properties of this selection in terms of the empirical distribution of the spectral radius of the sample fixed point mapping. In our Monte Carlo experiments we show that this convergence selection problem affects especially the structural parameter that represents competition effects: it introduces an attenuation bias in this parameter.

An additional contribution of this paper is that we propose an alternative iterative algorithm to compute the NPL estimator. Instead of fixed point iterations, the proposed algorithm uses a spectral approach to solve for the fixed points of the NPL mapping. Fixed point instability therefore does not prevent this alternative algorithm to converge.

An insightful implication of our findings is that the lack of convergence of the NPL algorithm when the equilibrium generating the data is stable is *entirely* due to the stochastic nature of the sample NPL mapping. In that sense, the current paper fills an important gap in the literature's understanding of the properties of the NPL algorithm. More precisely, even if the data generating process is stable, stability may fail to hold for the sample NPL mapping. In such cases, the NPL algorithm will fail to converge to the NPL estimator. This is exactly the reason why it matters to study the properties of the NPL algorithm iterated in finite samples, i.e. considering the sample size to grow only after the iterative procedure is completed. This important feature of the NPL algorithm is implicitly ignored when treating the iterative procedure as being initiated at the true data generating process.

Another important implication of the randomness of the sample NPL mapping is that the NPL algorithm is likely to fail to converge in finite samples if the data generating process is stable, but close to being unstable. In order to appreciate this property of the NPL algorithm, the Monte Carlo simulations included in this paper are repeated for a grid of data generating processes that satisfy stability conditions to different degrees. We find that [Aguirregabiria and Mira \(2007\)](#)'s data generating process – commonly used as a benchmark for Monte Carlo experiments – is in fact close to being unstable.

The rest of the paper is organized as follows. Section 2 reviews the related literature. The model of interest, a discrete choice dynamic game with incomplete information, is presented in Section 3. The NPL estimator and different sequential estimators are described in Section 4. Section 5 presents our main results. We provide necessary and sufficient conditions for the convergence of the NPL algorithm, and use these conditions to characterize the statistical

properties of the estimator upon convergence. We also describe the selection bias due to convergence. In Section 6, we present evidence that the selection problem generated by the lack of convergence is the most important source of finite sample bias when the data generating process is stable. Section 7 concludes. Codes in R for all our Monte Carlo experiments are available [here](#).

2 Related literature

A motivation behind the NPL estimator for dynamic games introduced by [Aguirregabiria and Mira \(2007\)](#) is the reduction of the finite sample bias typically associated with the two-step approach that is often preferred to a computationally costly full likelihood estimator. One can extend a Nested Fixed Point (NFXP) algorithm similar to [Rust \(1987\)](#) to a game setting to obtain a maximum likelihood estimator (MLE) of the structural parameters of interest. Multiplicity of equilibria in games complicates the implementation of the NFXP algorithm compared to its typical application to single-agent dynamic discrete choice problems. An important reduction in the computational burden can be achieved by using a two-stage approach in which reduced-form CCPs are estimated in the first step and then taken as given when estimating the structural parameters in the second step. This approach corresponds to an application of [Hotz and Miller \(1993\)](#)'s CCP estimator to the dynamic games framework and was initially proposed by [Aguirregabiria and Mira \(2007\)](#).¹ Different two-stage estimators have been proposed in the context of dynamic games: [Bajari, Benkard, and Levin \(2007\)](#)'s simulated minimum distance estimator; [Pakes, Ostrovsky, and Berry \(2007\)](#) advocating for the use of method of moments estimators in the second step; and [Pesendorfer and Schmidt-Dengler \(2008\)](#)'s asymptotic least squares estimator.

Despite being computationally convenient and having nice asymptotic properties, two-stage approaches tend to perform poorly in small samples, especially when there is a large number of states in the model. In that case, the CCPs from the first step are typically biased and imprecise, therefore affecting the structural parameters' estimates. It is also worth noting that – since they are reduced-form estimates – the first-step CCPs do not satisfy the fixed point restrictions *in finite samples*.

[Aguirregabiria and Mira \(2007\)](#) aim at addressing these issues by proposing the NPL estimator and an iterative algorithm to compute this estimator.

The NPL estimator is defined as a vector of structural parameters and a vector of CCPs that satisfy two conditions: the vector of structural parameters maximizes a pseudo likelihood function that depends on the CCPs; and the vector of CCPs satisfies the model equilibrium

¹An assumption commonly maintained when implementing this CCP estimator is that, while the game theoretical model may admit multiple equilibria, the same equilibrium is realized whenever the same game is played.

conditions associated with the vector of structural parameters. These two conditions can be represented as a single fixed point condition: the vector of CCPs is a fixed point of the so-called NPL mapping. If the NPL mapping has multiple fixed points, then the NPL estimator is the solution with the maximum value of the pseudo likelihood function.

To compute the NPL estimator, [Aguirregabiria and Mira \(2007\)](#) propose an iterative procedure, denoted as the NPL algorithm. This procedure consists of fixed point iterations in the NPL mapping. Given a vector of CCPs as input, a single iteration of the NPL algorithm provides a new vector of CCPs that is equal to players' best response to the input CCPs when the vector of structural parameters is the one that maximizes the pseudo likelihood function for the input CCPs. Upon convergence, the NPL algorithm delivers a fixed point of the NPL mapping.

[Aguirregabiria and Mira \(2007\)](#) did not formally show the convergence of the NPL algorithm. In a nutshell, the limitations of the NPL algorithm can be summarized in two important criticisms. First, in some cases, none of the NPL fixed points found by the NPL algorithm initiated at different values of the CCPs correspond to the NPL estimator. Then, the NPL algorithm converges to an inconsistent estimator. Some explanations for such failure of the NPL algorithm are available in the literature. In fact, insights from [Pesendorfer and Schmidt-Dengler \(2010\)](#) and [Kasahara and Shimotsu \(2012\)](#) suggest that the NPL algorithm cannot deliver the NPL estimator if the data is known to be generated from an unstable equilibrium.

Second, even if the data generating process satisfies desirable stability conditions, the NPL algorithm may still fail to converge ([Egesdal, Lai, and Su, 2015](#)). Unfortunately, the literature has so far failed to provide a convincing explanation for this lack of convergence. Numerical issues and tolerance level used to assess convergence have been accused without much justification. [Mogensen \(2015\)](#) provides simulation evidence that non-converging sequences typically coincide with parameters' estimates for which the stability conditions required for local contraction do not hold. This observation is, to the best of our knowledge, the first attempt at understanding the NPL algorithm's possible failure to converge when the data is generated from a stable equilibrium.

It is important to highlight the fundamental difference between the NPL algorithm and the NPL estimator. The NPL estimator, as defined by [Aguirregabiria and Mira \(2007, p.20\)](#), is "the NPL fixed point in the sample with the maximum value of the pseudo likelihood". The nice asymptotic properties of the NPL estimator (stated in [Aguirregabiria and Mira, 2007, Proposition 2](#)) do not depend on stability conditions required for the NPL algorithm to converge. In other words, the NPL estimator exists regardless whether the NPL algorithm converges or not.

A contribution of the current paper is to propose an alternative to the NPL algorithm. The proposed algorithm avoids fixed point iterations by using non-monotone spectral residual

methods to solve the system of nonlinear equations defining the fixed points of the NPL mapping. Importantly, the convergence of the algorithm does not require fixed point stability.

At this point, it is worth mentioning that there are two different ways of investigating the properties of the sequences of estimates generated by the NPL algorithm. The first one – which is the one commonly used in the literature (e.g., [Kasahara and Shimotsu, 2012](#)) – consists in considering an asymptotic approximation to the NPL algorithm, as if the iterations were performed in a large sample. Under this interpretation, the convergence of the NPL algorithm is determined by the stability of the equilibrium generating the data since one is effectively updating the choice probabilities using the NPL mapping in the population. Based on this approach, one expects the NPL algorithm to converge to the consistent NPL estimator if and only if the equilibrium generating the data is stable.

A second interpretation – motivated by [Pesendorfer and Schmidt-Dengler \(2010\)](#)'s criticism – rather iterates the sequential estimator in finite samples and then considers the sample size to grow to infinity. This is the interpretation that is considered in the current paper. In fact, it seems to be the most sensible interpretation when the NPL algorithm is used as a device to reduce the finite sample bias of two-stage estimators. It is also key to understand the convergence issues of the NPL algorithm when the equilibrium generating the data is stable.

Some alternative estimators have been proposed in hope of addressing the convergence issues associated with the NPL algorithm. One of them is the modified NPL algorithm from [Kasahara and Shimotsu \(2012\)](#). This iterative algorithm uses the relaxation method from numerical analysis when updating the CCPs to guarantee that the modified NPL mapping is locally contractive. In order to implement this approach, one must fix the value of the parameter modifying the NPL mapping. [Kasahara and Shimotsu \(2012\)](#) provide simulation evidence that this modified NPL algorithm works well if this parameter is set optimally. However, such an optimal value is not available in practice since it requires the knowledge of the NPL mapping in the population, which is obviously unknown to the econometrician. [Kasahara and Shimotsu \(2012\)](#) therefore recommend using a small arbitrary value for this parameter. Unfortunately, [Egesdal, Lai, and Su \(2015\)](#) have found the properties of the modified NPL algorithm to be much less appealing when using such a small arbitrary value.

Another alternative approach is a class of Mathematical Program with Equilibrium Constraints (MPEC) estimators which was proposed by [Su and Judd \(2012\)](#) for single-agent dynamic discrete choice models, and extended to dynamic games by [Egesdal, Lai, and Su \(2015\)](#). This approach consists in computing the constrained MLE implemented using state of the art optimization algorithms. On one hand, it has the important advantage of delivering the MLE, i.e. what one would obtain from the NFXP approach. On the other hand, it now requires optimizing over a large vector of parameters which includes the structural parameters, the equilibrium CCPs, and the Lagrange multipliers of all the equilibrium con-

straints. While the dimensionality of the problem therefore increases with the state space, the optimizer leverages the sparsity of the Jacobian and Hessian matrices involved in the maximization problem. [Egesdal, Lai, and Su \(2015\)](#) provide simulation evidence suggesting better convergence properties (in the sense of finding a solution to the maximization problem) than the estimator obtained using the NPL algorithm in the context of dynamic discrete games.²

[Dearing and Blevins \(2019\)](#) have recently proposed the Efficient Pseudo Likelihood (EPL) algorithm as an alternative to NPL. This algorithm is similar to NPL, but it iterates in the space of conditional choice values, instead of the space of CCPs. The authors show that their algorithm converges in situations where NPL does not. Nevertheless, similarly as other papers in this literature, their asymptotic results only apply when the iterative process is performed in large samples.

While alternatives such as [Kasahara and Shimotsu \(2012\)](#)'s modified NPL algorithm, [Egesdal, Lai, and Su \(2015\)](#)'s MPEC approach and [Dearing and Blevins \(2019\)](#)'s estimator clearly have merit, there is still room for a better understanding of the NPL algorithm's properties. While we are not aware of published work applying these alternative estimators in models of dynamic games, MPEC has been used to estimate single-agent dynamic models ([Aksin, Ata, Emadi, and Su, 2013](#); [Chen, Esteban, and Shum, 2013](#); [Barwick and Pathak, 2015](#); [Frechette, Lizzeri, and Salz, 2019](#)) and in static games ([Vitorino, 2012](#); [Su, 2014](#)). In applications with large state spaces, the NPL algorithm is arguably much easier to implement and has a substantially smaller computational cost. An important goal of the current paper is to improve our understanding of the NPL algorithm's failure to converge in hopes to identify in which settings this method can be safely used in empirical applications. It is worth noting that the NPL algorithm has been used to estimate several empirical applications of single-agent dynamic models ([Copeland and Monnet, 2009](#); [De Pinto and Nelson, 2009](#); [Tomlin, 2014](#); [Aguirregabiria and Alonso-Borrego, 2014](#)), dynamic games ([Sweeting, 2013](#); [Aguirregabiria and Ho, 2012](#); [Collard-Wexler, 2013](#); [Kano, 2013](#); [Huang and Smith, 2014](#); [Lin, 2015](#); [Gayle and Xie, 2018](#)), static games ([Ellickson and Misra, 2008](#); [Han and Hong, 2011](#)) and networks ([Lin and Xu, 2017](#); [Liu and Zhou, 2017](#)).

[Bugni and Bunting \(2018\)](#) recently studied the properties of the variance of iterative K -step estimators. More precisely, they compare the K -step variance of a maximum likelihood (K -ML) and a minimum distance (K -MD) estimators. They find that, while the variance of the K -ML estimator can be quite erratic with K , it is K -invariant for the K -MD. This result is very interesting on its own, but it does not address the convergence and finite sample bias problems that motivate the current paper.

²For single-agent dynamic discrete choice model, [Iskhakov, Lee, Rust, Schjerning, and Seo \(2016\)](#) provide simulation evidence suggesting that the NFXP algorithm may actually have better convergence properties than the MPEC estimator.

3 The model

3.1 General framework

Let \mathcal{Y} be a discrete and finite set and let $P^0(\mathbf{y}|\mathbf{x})$ be the true probability distribution of the realization of a $N \times 1$ vector $\mathbf{y} \in \mathcal{Y}^N$, for $N \geq 1$, conditional on some vector of realized exogenous variables $\mathbf{x} \in \mathcal{X}$. Also, denote $\mathbf{P}^0 = \{P^0(\mathbf{y}|\mathbf{x}) : \mathbf{y} \in \mathcal{Y}^N, \mathbf{x} \in \mathcal{X}\} \in \mathcal{P}$. There is a parametric structural model for this vector of conditional probabilities. Consider a finite-dimensional vector of structural parameters $\boldsymbol{\theta} \in \Theta$ and denote the vector of parametrized probabilities as $\mathbf{P}(\boldsymbol{\theta}) = \{P(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}) : \mathbf{y} \in \mathcal{Y}^N, \mathbf{x} \in \mathcal{X}\}$. The true value of the structural parameters in the population, denoted $\boldsymbol{\theta}^0$ satisfies $\mathbf{P}(\boldsymbol{\theta}^0) = \mathbf{P}^0$.

There is no closed-form analytical expression for $P(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})$. The parametric distribution of \mathbf{y} is implicitly defined as a fixed point of the following mapping in the probability space:

$$\mathbf{P}(\boldsymbol{\theta}) = \boldsymbol{\Psi}(\boldsymbol{\theta}, \mathbf{P}(\boldsymbol{\theta})) \quad (1)$$

where $\boldsymbol{\Psi}(\cdot) : \Theta \times \mathcal{P} \rightarrow \mathcal{P}$ and $\boldsymbol{\Psi}(\boldsymbol{\theta}, \mathbf{P}) = \{\Psi(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}, \mathbf{P}) : \mathbf{y} \in \mathcal{Y}^N, \mathbf{x} \in \mathcal{X}\}$ is a vector with elements organized in the same order as in \mathbf{P} .

Typically, the researcher has a sample of M observations indexed by $m \in \{1, 2, \dots, M\}$. For each observation, the data set includes the realizations $\mathbf{y}_m \in \mathcal{Y}^N$ and $\mathbf{x}_m \in \mathcal{X}$. The index of the data could vary over more than one dimensions, allowing for time and geographical variations for instance. The key element here is that all the asymptotic results depend on one of these dimensions being large. For what follows, this dimension is the one indexed by m , i.e. $M \rightarrow \infty$.

3.2 Dynamic discrete game of market entry and exit

In this section, we introduce the dynamic game of market entry and exit that will be used for the Monte Carlo experiments in Section 6. This model comes from [Aguirregabiria and Mira \(2007\)](#) and it is a special case of the general framework presented above. It has been used in several simulation exercises, namely [Pesendorfer and Schmidt-Dengler \(2008\)](#), [Kasahara and Shimotsu \(2012\)](#), and [Egedal, Lai, and Su \(2015\)](#).³

Each firm – indexed by $i \in \{1, 2, \dots, N\}$ – decides whether or not to operate in a market – $m \in \{1, 2, \dots, M\}$ – at a given period – $t \in \{1, 2, \dots, T\}$. In this context, $\mathcal{Y} = \{0, 1\}$ is each firm's choice set, such that $y_{imt} = 1$ if firm i operates in market m at time t . Let \mathbf{y}_{mt} be the vector $\{y_{imt} : i = 1, \dots, N\}$. A firm maximizes its expected and discounted stream of current and future profits in the market: $E_t[\sum_{s=0}^{\infty} \beta^s U_{im,t+s}]$, where $\beta \in (0, 1)$ is the discount

³Note that the Monte Carlo experiments in [Pesendorfer and Schmidt-Dengler \(2008\)](#) also use a dynamic game of market entry-exit. Their specification can be seen as a particular case of the model in [Aguirregabiria and Mira \(2007\)](#): no market size, and only two players.

factor, and U_{imt} is the period profit.

Suppose that a firm's decision to operate depends on the market size (S_{mt}), the number of competitors that are operating ($\sum_{j \neq i} y_{jmt}$), an entry cost, a firm-specific fixed cost, and the vector of variables $\boldsymbol{\varepsilon}_{imt} = [\varepsilon_{imt}(0), \varepsilon_{imt}(1)]'$, which are player i 's private information. Let $\boldsymbol{\varepsilon}_{mt}$ be the vector $\{\boldsymbol{\varepsilon}_{imt} : i = 1, \dots, N\}$. The dynamic dimension to the firm's decision comes from the firm paying the entry cost only if it has not been operating in the market during the previous period.

An observation is a firm-market-period-tuple. Besides the decision of being active in the market, the econometrician observes the market size (S_{mt}) and the incumbency status of the firms ($\mathbf{y}_{m,t-1}$). To be consistent with the notation introduced above, denote the observable state variables in market m at period t as $\mathbf{x}_{mt} = [S_{mt}, \mathbf{y}'_{m,t-1}]'$. The following assumptions are made by [Aguirregabiria and Mira \(2007\)](#) and are fairly standard in the literature.

Assumption 1 (Additive separability). The contemporaneous utility function is additively separable in its observable and its unobservable components: $U_i(\mathbf{y}_{mt}, \mathbf{x}_{mt}, \varepsilon_{imt}(y_{imt})) = u_i(\mathbf{y}_{mt}, \mathbf{x}_{mt}) + \varepsilon_{imt}(y_{imt})$.

Assumption 2 (Conditional independence). The joint transition probability of $\mathbf{x}_{m,t+1}, \boldsymbol{\varepsilon}_{im,t+1}$ conditional on \mathbf{y}_{mt} is $\Pr(\mathbf{x}_{m,t+1}, \boldsymbol{\varepsilon}_{im,t+1} | \mathbf{x}_{mt}, \boldsymbol{\varepsilon}_{im,t}, \mathbf{y}_{mt}) = \Gamma(\mathbf{x}_{m,t+1} | \mathbf{x}_{mt}, \mathbf{y}_{mt}) G(\boldsymbol{\varepsilon}_{im,t+1})$, where $G(\cdot)$ is the cumulative distribution function of $\boldsymbol{\varepsilon}_{im,t}$, which is iid across players, markets, and time periods.

Assumption 3 (Discrete support for observable state variables). The observable state variables \mathbf{x}_{mt} have a discrete and finite support.

For instance, [Aguirregabiria and Mira \(2007\)](#) consider the following specification of the current utility function $U_i(\mathbf{y}_{mt}, \mathbf{x}_{mt}, \varepsilon_{imt}(y_{imt}))$:

$$\begin{cases} \theta_{RS} S_{mt} - \theta_{RN} \ln \left[1 + \sum_{j \neq i} y_{jmt} \right] - \theta_{EC} (1 - y_{im,t-1}) - \theta_{FC,i} + \varepsilon_{imt}(1) & , y_{imt} = 1 \\ \varepsilon_{imt}(0) & , y_{imt} = 0 \end{cases} \quad (2)$$

where θ_{RS} measures the effect of the market size on the firm's payoff; θ_{RN} is referred to as the strategic interaction parameter; θ_{EC} corresponds to the entry cost; and $\theta_{FC,i}$ is the firm-specific fixed cost.

We assume that players' behavior can be rationalized via a Markov Perfect Equilibrium, which implies that player i 's strategies only depend on this player's payoff-relevant state variables, more precisely \mathbf{x}_{mt} and $\boldsymbol{\varepsilon}_{im,t}$. At a given state – defined by \mathbf{x}_{mt} and $\boldsymbol{\varepsilon}_{im,t}$ – let $\sigma_i(\cdot) : \mathcal{X} \times \mathbb{R}^2 \mapsto \mathcal{Y}$ be player i 's strategy function. Given the other firms' strategies – represented as $\boldsymbol{\sigma}_{-i,mt} \equiv \boldsymbol{\sigma}_{-i}(\mathbf{x}_{mt}, \boldsymbol{\varepsilon}_{-i,mt})$ – firm i 's best response is the solution of a dynamic programming problem with the following value function and Bellman equation:

$$V_i(\mathbf{x}_{mt}, \boldsymbol{\varepsilon}_{imt}) = \max_{y \in \{0,1\}} \left\{ \mathbb{E}_{\boldsymbol{\varepsilon}_{-i,mt}} \left[U_i(y, \boldsymbol{\sigma}_{-i,mt}, \mathbf{x}_{mt}, \boldsymbol{\varepsilon}_{imt}(y)) \right. \right. \\ \left. \left. + \beta \int \sum_{\mathbf{x}_{m,t+1} \in \mathcal{X}} V_i(\mathbf{x}_{m,t+1}, \boldsymbol{\varepsilon}_{im,t+1}) d\Gamma(\mathbf{x}_{m,t+1} | \mathbf{x}_{mt}, y, \boldsymbol{\sigma}_{-i,mt}) dG(\boldsymbol{\varepsilon}_{im,t+1}) \right] \right\} \quad (3)$$

$$\equiv \max_{y \in \{0,1\}} \left\{ \mathbb{E}_{\boldsymbol{\varepsilon}_{-i,mt}} [v_i(y, \boldsymbol{\sigma}_{-i,mt}, \mathbf{x}_{mt})] + \boldsymbol{\varepsilon}_{imt}(y) \right\}. \quad (4)$$

From player i 's point of view, $v_i(y, \boldsymbol{\sigma}_{-i,mt}, \mathbf{x}_{mt})$ is the sum of its current utility (without $\boldsymbol{\varepsilon}_{imt}(y)$) and its discounted future expected utility of choosing y when player i 's competitors behave according to $\boldsymbol{\sigma}_{-i,mt}$ for a given \mathbf{x}_{mt} .

In equilibrium, players' strategies are best responses to each other and – taking the strategies of the other players as given – a player's best response is the solution of a single-agent dynamic programming problem. The set of players' strategies also determine their CCPs such that:

$$P(y_i | \mathbf{x}_{mt}, \boldsymbol{\theta}) = \int \mathbb{1} \{y_i = \sigma_i(\mathbf{x}_{mt}, \boldsymbol{\varepsilon}_{imt})\} dG(\boldsymbol{\varepsilon}_{imt}). \quad (5)$$

One can write the expectation with respect to $\boldsymbol{\varepsilon}_{-i,mt}$ as a function of competitors' CCPs. Let $v_i^{\mathbf{P}}(y, \mathbf{x}_{mt}) \equiv \mathbb{E}_{\boldsymbol{\varepsilon}_{-i,mt}} [v_i(y, \boldsymbol{\sigma}_{-i,mt}, \mathbf{x}_{mt})]$, where the superscript \mathbf{P} is used to make explicit the dependence on the distribution of \mathbf{y} . Then:

$$v_i^{\mathbf{P}}(y, \mathbf{x}_{mt}) = \sum_{\mathbf{y}_{-i} \in \mathcal{Y}^{N-1}} \prod_{j \neq i}^N P(y_j | \mathbf{x}_{mt}, \boldsymbol{\theta}) v_i(y, \mathbf{y}_{-i} = \boldsymbol{\sigma}_{-i,mt}, \mathbf{x}_{mt}). \quad (6)$$

Provided that each firm's strategy is to enter if and only if the expected payoff of entering is greater than the expected payoff of not entering, the CCP of firm i entering market m at time t , denoted $P_i(1 | \mathbf{x}_{mt}, \boldsymbol{\theta}) \equiv P(y_i = 1 | \mathbf{x}_{mt}, \boldsymbol{\theta})$, is:

$$P_i(1 | \mathbf{x}_{mt}, \boldsymbol{\theta}) = \int \mathbb{1} \{v_i^{\mathbf{P}}(1, \mathbf{x}_{mt}) - v_i^{\mathbf{P}}(0, \mathbf{x}_{mt}) \geq \boldsymbol{\varepsilon}_{imt}(0) - \boldsymbol{\varepsilon}_{imt}(1)\} dG(\boldsymbol{\varepsilon}_{imt}) \quad (7)$$

which can be used to construct the fixed point mapping in (1).

Before presenting the estimators of interest, Assumption 4 summarizes the conditions on the data generating process typically supposed to be satisfied in the literature. In essence, these restrictions guarantee that the sample consists of independent draws generated from a single (given observed states) Markov Perfect Equilibrium.

Assumption 4 (Data generating process). (i) $P^0(\mathbf{y}_{mt} | \mathbf{x}_{mt}) = \Psi(\mathbf{y}_{mt} | \mathbf{x}_{mt}, \boldsymbol{\theta}^0, \mathbf{P}^0) \forall m, t$.

(ii) Players expect \mathbf{P}^0 to be played in the future (out of sample periods). (iii) For any \mathbf{P} that solves $\mathbf{P} = \Psi(\boldsymbol{\theta}, \mathbf{P})$, $\mathbf{P} \neq \mathbf{P}^0$ whenever $\boldsymbol{\theta} \neq \boldsymbol{\theta}^0$. (iv) $\{\mathbf{y}_{mt}, \mathbf{x}_{mt} : m = 1, \dots, M; t = 1, \dots, T\}$ are independent across markets and $\Pr(\mathbf{x}_{mt} = \mathbf{x}) > 0 \forall \mathbf{x} \in \mathcal{X}$.

The following regularity conditions will be maintained.

Assumption 5 (Regularity conditions). (i) The mapping $\Psi(\boldsymbol{\theta}, \mathbf{P})$ is twice continuously differentiable. (ii) $\boldsymbol{\theta}^0 \in \text{int}(\Theta)$. (iii) Θ is compact. (iv) There is an open ball around \mathbf{P}^0 that does not include any other fixed point of the mapping $\Psi(\boldsymbol{\theta}^0, \cdot)$.

Throughout the paper, we use a hat superindex to indicate that an object is a statistic or a stochastic function that has sampling variation. Objects without a hat superindex are deterministic functions or parameters.

4 NPL estimator and sequential estimators

In this section, we introduce the NPL mapping and use it to define the (algorithm-free) NPL estimator. We also present a series of sequential estimators including the Hotz-Miller estimator, the K -step estimator, the NPL algorithm, and a new algorithm to compute the NPL estimator.

4.1 NPL mapping and estimator

Let $\hat{Q}(\boldsymbol{\theta}, \mathbf{P}) \equiv M^{-1} \sum_{m=1}^M \ln[\Psi(\mathbf{y}_m | \mathbf{x}_m, \boldsymbol{\theta}, \mathbf{P})]$ be the pseudo log-likelihood function for any $\boldsymbol{\theta} \in \Theta$ and any $\mathbf{P} \in \mathcal{P}$. Its population counterpart is $Q^0(\boldsymbol{\theta}, \mathbf{P}) \equiv E_0[\ln[\Psi(\mathbf{y}_m | \mathbf{x}_m, \boldsymbol{\theta}, \mathbf{P})]]$, where $E_0[\cdot]$ is computed with respect to the true probability distribution \mathbf{P}^0 . By maximizing these functions over Θ for a given $\mathbf{P} \in \mathcal{P}$, one obtains the pseudo maximum likelihood estimator $\hat{\boldsymbol{\theta}}(\mathbf{P}) = \arg \max_{\boldsymbol{\theta} \in \Theta} \hat{Q}(\boldsymbol{\theta}, \mathbf{P})$ and its population counterpart $\boldsymbol{\theta}^0(\mathbf{P}) = \arg \max_{\boldsymbol{\theta} \in \Theta} Q^0(\boldsymbol{\theta}, \mathbf{P})$.

By maximizing the pseudo log-likelihood function evaluated at a consistent reduced-form estimate of the CCPs, say $\hat{\mathbf{P}}_0$, one obtains the two-step pseudo MLE or Hotz-Miller estimator applied to dynamic discrete games. This estimator is formally defined in Definition 1.

Definition 1 (Two-step pseudo MLE). For a given consistent reduced-form estimate of the CCPs $\hat{\mathbf{P}}_0$, the two-step pseudo MLE estimator $\hat{\boldsymbol{\theta}}_1$ satisfies:

$$\hat{\boldsymbol{\theta}}_1 \equiv \hat{\boldsymbol{\theta}}(\hat{\mathbf{P}}_0) = \arg \max_{\boldsymbol{\theta} \in \Theta} \hat{Q}(\boldsymbol{\theta}, \hat{\mathbf{P}}_0).$$

The vectors $\hat{\boldsymbol{\theta}}(\mathbf{P})$ and $\boldsymbol{\theta}^0(\mathbf{P})$ are also used to define the sample and the population NPL mappings. As it is formally stated in Definition 2, for any $\mathbf{P} \in \mathcal{P}$, the NPL mapping is a function that returns the vector of best response CCPs at the corresponding pseudo-maximum

likelihood estimator of $\boldsymbol{\theta}$. The only difference between the population NPL mapping and its sample counterpart is whether the pseudo-maximum likelihood estimator of $\boldsymbol{\theta}$ is computed in the population or in the sample.

Definition 2 (NPL mapping). For any $\mathbf{P} \in \mathcal{P}$, the sample NPL mapping is $\hat{\boldsymbol{\varphi}}(\mathbf{P}) \equiv \boldsymbol{\Psi}(\hat{\boldsymbol{\theta}}(\mathbf{P}), \mathbf{P})$; its population counterpart is $\boldsymbol{\varphi}^0(\mathbf{P}) \equiv \boldsymbol{\Psi}(\boldsymbol{\theta}^0(\mathbf{P}), \mathbf{P})$.

Notice that these NPL mappings are continuous functions of a compact set, i.e. \mathcal{P} , to itself. By Brouwer's fixed point theorem, a fixed point of the NPL mapping always exists, both for the sample and the population mappings. In other words, the vector mappings $\hat{\boldsymbol{\phi}}(\mathbf{P}) \equiv \mathbf{P} - \hat{\boldsymbol{\varphi}}(\mathbf{P})$ and $\boldsymbol{\phi}^0(\mathbf{P}) \equiv \mathbf{P} - \boldsymbol{\varphi}^0(\mathbf{P})$ both have at least one root. However, uniqueness is not guaranteed. Definition 3 formally states the conditions satisfied by a fixed point of the sample NPL mapping.

Definition 3 (Fixed point of the sample NPL mapping). $\hat{\mathbf{P}}_*$ is a fixed point of the sample NPL mapping if $\hat{\mathbf{P}}_* = \hat{\boldsymbol{\varphi}}(\hat{\mathbf{P}}_*)$. This implies that for $\hat{\boldsymbol{\theta}}_* \equiv \hat{\boldsymbol{\theta}}(\hat{\mathbf{P}}_*)$, the following conditions hold: (i) $\hat{\boldsymbol{\theta}}_* = \arg \max_{\boldsymbol{\theta} \in \Theta} \hat{Q}(\boldsymbol{\theta}, \hat{\mathbf{P}}_*)$; and (ii) $\hat{\mathbf{P}}_* = \boldsymbol{\Psi}(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)$.

The NPL estimator is formally stated in Definition 4.

Definition 4 (NPL estimator). The NPL estimator $(\hat{\boldsymbol{\theta}}_{\text{NPL}}, \hat{\mathbf{P}}_{\text{NPL}})$ is defined as the fixed point of the sample NPL mapping achieving the largest value of the pseudo-likelihood function. That is, it satisfies the following conditions: (i) $\hat{\mathbf{P}}_{\text{NPL}} = \hat{\boldsymbol{\varphi}}(\hat{\mathbf{P}}_{\text{NPL}})$; (ii) $\hat{\boldsymbol{\theta}}_{\text{NPL}} = \hat{\boldsymbol{\theta}}(\hat{\mathbf{P}}_{\text{NPL}})$; and (iii) $(\hat{\boldsymbol{\theta}}_{\text{NPL}}, \hat{\mathbf{P}}_{\text{NPL}})$ maximizes $\hat{Q}(\boldsymbol{\theta}, \mathbf{P})$ within the set of fixed points of the sample NPL mapping.

The NPL estimator is consistent and asymptotically normal. Furthermore, if the spectral radius⁴ of the Jacobian matrix $\partial \boldsymbol{\varphi}^0(\mathbf{P}) / \partial \mathbf{P}$ evaluated at \mathbf{P}^0 is smaller than one, it is asymptotically more efficient than the two-step estimator (Proposition 2 in Aguirregabiria and Mira, 2007). Proposition 1 is an alternative statement of the consistency and asymptotic normality of the NPL estimator. Our proof is different from the one in Aguirregabiria and Mira (2007). Some notation must be introduced before stating the result. Let the empirical measure be denoted by \mathbb{P} where:

$$\mathbb{P} = \left\{ \mathbb{P}(\mathbf{y} | \mathbf{x}) = \frac{\sum_{m=1}^M \mathbb{1}\{\mathbf{y}_m = \mathbf{y}\} \mathbb{1}\{\mathbf{x}_m = \mathbf{x}\}}{\sum_{m=1}^M \mathbb{1}\{\mathbf{x}_m = \mathbf{x}\}} : \mathbf{y} \in \mathcal{Y}^N, \mathbf{x} \in \mathcal{X} \right\} \quad (8)$$

with the elements of \mathbb{P} ordered in the same way as in \mathbf{P} . Let \mathbf{H} be a random vector such that $\mathbf{H} \sim \text{Normal}(\mathbf{0}, \text{Var}[\mathbf{H}])$ where $\text{Var}[\mathbf{H}]$ is a block diagonal matrix with \mathbf{x} -specific block corresponding to:

⁴For a given square matrix \mathbf{A} , let $\rho(\mathbf{A})$ denote its spectral radius, i.e. its largest eigenvalue in absolute value. The condition $\rho(\mathbf{A}) < 1$ is sufficient, but not necessary, for \mathbf{A}^K to converge to $\mathbf{0}$ for large enough K (where \mathbf{A}^K denotes the matrix product of K matrices \mathbf{A}).

$$\text{Var} [\mathbb{P}(\mathbf{y}|\mathbf{x})] = \frac{M}{\sum_{m=1}^M \mathbb{1}\{\mathbf{x}_m = \mathbf{x}\}} [\text{diag}\{\mathbf{P}^0(\mathbf{y}|\mathbf{x})\} - \mathbf{P}^0(\mathbf{y}|\mathbf{x})\mathbf{P}^0(\mathbf{y}|\mathbf{x})'] \quad (9)$$

and covariances between $\mathbb{P}(\mathbf{y}|\mathbf{x})$ for different values of \mathbf{x} are 0. It follows that, as $M \rightarrow \infty$:

$$\sqrt{M}(\mathbb{P} - \mathbf{P}^0) \xrightarrow{d} \mathbf{H}. \quad (10)$$

Let $\nabla_{\mathbf{P}}\Psi_{(\boldsymbol{\theta},\mathbf{P})}$ and $\nabla_{\boldsymbol{\theta}}\Psi_{(\boldsymbol{\theta},\mathbf{P})}$ respectively denote the Jacobians of $\Psi(\boldsymbol{\theta},\mathbf{P})$ with respect to \mathbf{P} and $\boldsymbol{\theta}$, when evaluated at $(\boldsymbol{\theta},\mathbf{P})$. Let $\nabla\hat{\boldsymbol{\varphi}}_{(\mathbf{P})}$ be the Jacobian of the sample NPL mapping evaluated at \mathbf{P} . Similarly, let $\nabla\boldsymbol{\varphi}_{(\mathbf{P})}^0$ be its population counterpart. Let $\nabla\hat{\boldsymbol{\phi}}_{(\mathbf{P})}$ denote the Jacobian of $\hat{\boldsymbol{\phi}}(\mathbf{P})$ evaluated at \mathbf{P} . Define $\nabla\boldsymbol{\phi}_{(\mathbf{P})}^0$ similarly for $\boldsymbol{\phi}^0(\mathbf{P})$. Finally, let $\nabla_{\boldsymbol{\theta}}\hat{Q}_{(\boldsymbol{\theta},\mathbf{P})}$, $\nabla_{\mathbf{P}}\hat{Q}_{(\boldsymbol{\theta},\mathbf{P})}$, $\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2\hat{Q}_{(\boldsymbol{\theta},\mathbf{P})}$ and $\nabla_{\boldsymbol{\theta}\mathbf{P}}^2\hat{Q}_{(\boldsymbol{\theta},\mathbf{P})}$ be the matrices of first and second order derivatives of $\hat{Q}(\cdot)$ evaluated at $(\boldsymbol{\theta},\mathbf{P})$. The matrices $\nabla_{\boldsymbol{\theta}}Q_{(\boldsymbol{\theta},\mathbf{P})}^0$, $\nabla_{\mathbf{P}}Q_{(\boldsymbol{\theta},\mathbf{P})}^0$, $\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2Q_{(\boldsymbol{\theta},\mathbf{P})}^0$ and $\nabla_{\boldsymbol{\theta}\mathbf{P}}^2Q_{(\boldsymbol{\theta},\mathbf{P})}^0$ are equivalently defined for $Q^0(\cdot)$.

The asymptotic properties of the NPL estimator stated in Proposition 1 rely on the asymptotic distribution of the sample NPL mapping which is given in Lemma 1.⁵

Lemma 1 (Asymptotic distribution of the sample NPL mapping). *Let Assumptions 1-5 be satisfied. Then, for any $\mathbf{P} \in \mathcal{P}$, as $M \rightarrow \infty$, the maximum likelihood estimator $\hat{\boldsymbol{\theta}}(\mathbf{P})$ has limiting distribution:*

$$\sqrt{M}(\hat{\boldsymbol{\theta}}(\mathbf{P}) - \boldsymbol{\theta}^0(\mathbf{P})) \xrightarrow{d} \mathbf{L}(\mathbf{P}) \sim \text{Normal}\left(\mathbf{0}, \left[\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2Q_{(\boldsymbol{\theta}^0(\mathbf{P}),\mathbf{P})}^0\right]^{-1}\right). \quad (11)$$

The sample mappings $\hat{\boldsymbol{\phi}}(\mathbf{P})$ and $\hat{\boldsymbol{\varphi}}(\mathbf{P})$ have a limiting distribution such that:

$$\sqrt{M}(\hat{\boldsymbol{\phi}}(\mathbf{P}) - \boldsymbol{\phi}^0(\mathbf{P})) = -\sqrt{M}(\hat{\boldsymbol{\varphi}}(\mathbf{P}) - \boldsymbol{\varphi}^0(\mathbf{P})) \xrightarrow{d} \mathbf{Z}(\mathbf{P}) \sim \text{Normal}\left(\mathbf{0}, \boldsymbol{\Sigma}_{\hat{\boldsymbol{\phi}}(\mathbf{P})}\right)$$

$$\text{with } \boldsymbol{\Sigma}_{\hat{\boldsymbol{\phi}}(\mathbf{P})} \equiv \nabla_{\boldsymbol{\theta}}\Psi_{(\boldsymbol{\theta}^0(\mathbf{P}),\mathbf{P})} \left[\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2Q_{(\boldsymbol{\theta}^0(\mathbf{P}),\mathbf{P})}^0 \right]^{-1} \nabla_{\boldsymbol{\theta}}\Psi'_{(\boldsymbol{\theta}^0(\mathbf{P}),\mathbf{P})}.$$

Proof. See Appendix A.1. □

Proposition 1 (Asymptotic properties of the NPL estimator). *Let Assumptions 1-5 be satisfied. As $M \rightarrow \infty$, $(\hat{\boldsymbol{\theta}}_{NPL}, \hat{\mathbf{P}}_{NPL})$ satisfy:*

⁵It could be interesting to extend Proposition 1 to the weak convergence of $\sqrt{M}(\hat{\mathbf{P}}_{NPL} - \mathbf{P}^0)$ to a tight process. Two sufficient conditions for this would be that: (i) $\nabla\boldsymbol{\phi}_{(\mathbf{P})}^0$ is continuously invertible over $\mathbf{P} \in \mathcal{P}$; and (ii) $\mathbf{Z}(\mathbf{P})$ converges uniformly in $\mathbf{P} \in \mathcal{P}$ to a tight Gaussian process. One could then apply Kosorok (2008, Corollary 13.6) to get the desired result. This result would still be valid even if $\boldsymbol{\phi}^0(\mathbf{P}^0) = \mathbf{0}$ may have multiple solutions.

$$\begin{aligned}\sqrt{M} \left(\hat{\mathbf{P}}_{NPL} - \mathbf{P}^0 \right) &= - [\nabla \phi_{(\mathbf{P}^0)}^0]^{-1} \sqrt{M} \left(\hat{\phi}(\mathbf{P}^0) - \phi^0(\mathbf{P}^0) \right) + O_p \left(\sqrt{M} \left\| \hat{\mathbf{P}}_{NPL} - \mathbf{P}^0 \right\|^2 \right) \\ &\xrightarrow{d} Normal \left(\mathbf{0}, [\nabla \phi_{(\mathbf{P}^0)}^0]^{-1} \Sigma_{\hat{\phi}(\mathbf{P}^0)} [\nabla_{\mathbf{P}} \phi_{(\mathbf{P}^0)}^{0'}]^{-1} \right)\end{aligned}$$

and

$$\begin{aligned}\sqrt{M} \left(\hat{\boldsymbol{\theta}}_{NPL} - \boldsymbol{\theta}^0 \right) &= - \left[\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 Q_{(\boldsymbol{\theta}^0, \mathbf{P}^0)}^0 \right]^{-1} \left\{ \nabla_{\boldsymbol{\theta}\mathbf{P}}^2 Q_{(\boldsymbol{\theta}^0, \mathbf{P}^0)}^0 \sqrt{M} \left(\hat{\mathbf{P}}_{NPL} - \mathbf{P}^0 \right) \right. \\ &\quad \left. + \nabla_{\boldsymbol{\theta}} \Psi'_{(\boldsymbol{\theta}^0, \mathbf{P}^0)} \mathbf{D}^0 \sqrt{M} (\mathbf{P} - \mathbf{P}^0) \right\} + O_p \left(\sqrt{M} \left\| \hat{\mathbf{P}}_{NPL} - \mathbf{P}^0 \right\|^2 \right) \\ &\xrightarrow{d} \left[\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 Q_{(\boldsymbol{\theta}^0, \mathbf{P}^0)}^0 \right]^{-1} \left\{ \nabla_{\boldsymbol{\theta}\mathbf{P}}^2 Q_{(\boldsymbol{\theta}^0, \mathbf{P}^0)}^0 [\nabla \phi_{(\mathbf{P}^0)}^0]^{-1} \mathbf{Z}(\mathbf{P}^0) - \nabla_{\boldsymbol{\theta}} \Psi'_{(\boldsymbol{\theta}^0, \mathbf{P}^0)} \mathbf{D}^0 \mathbf{H} \right\}\end{aligned}$$

where $\mathbf{D}^0 \equiv \text{diag} \{ \mathbf{P}^0 \}^{-1}$.

Proof. See Appendix A.2. □

The main challenge in obtaining the NPL estimator is how to compute the fixed point(s) of the sample NPL mapping. [Aguirregabiria and Mira \(2007\)](#) propose to use the NPL algorithm that is defined in the next subsection.

4.2 NPL algorithm estimator

Let $\hat{\mathbf{P}}_0$ be a vector of nonparametric or reduced-form estimates of the CCPs. Let $\hat{\boldsymbol{\theta}}_1$ be the two-step pseudo maximum likelihood estimator – or Hotz-Miller estimator – as in Definition 1. These estimates can be used to update the CCPs using $\hat{\mathbf{P}}_1 = \Psi \left(\hat{\boldsymbol{\theta}}_1, \hat{\mathbf{P}}_0 \right)$, construct a new pseudo log-likelihood function $\hat{Q} \left(\boldsymbol{\theta}, \hat{\mathbf{P}}_1 \right)$ to be maximized, and so on. Such an iterative process is referred to as the NPL algorithm. When iterated for a fixed K number of times, this approach delivers the K -step estimator.

Definition 5 (K -step estimator). The k -th iteration of the NPL algorithm for $k \in \mathbb{N}$ is defined as $\hat{\mathbf{P}}_k = \hat{\varphi} \left(\hat{\mathbf{P}}_{k-1} \right)$, or equivalently:

$$\hat{\mathbf{P}}_k = \Psi \left(\hat{\boldsymbol{\theta}}_k, \hat{\mathbf{P}}_{k-1} \right)$$

with

$$\hat{\boldsymbol{\theta}}_k \equiv \hat{\boldsymbol{\theta}} \left(\hat{\mathbf{P}}_{k-1} \right) = \arg \max_{\boldsymbol{\theta} \in \Theta} \hat{Q} \left(\boldsymbol{\theta}, \hat{\mathbf{P}}_{k-1} \right)$$

and $\hat{\mathbf{P}}_0$ is a reduced-form estimate of the CCPs. The K -step estimator iterates until $k = K$.

It is easy to see that, upon convergence, the NPL algorithm delivers a fixed point of the sample NPL mapping. It is worth repeating that the estimator obtained upon convergence of the NPL algorithm does not necessarily correspond to the NPL estimator in Definition 4 as the NPL estimator is one of the potentially multiple fixed points of the sample NPL mapping.

It is important to emphasize that iterating the NPL algorithm is one possible method to obtain some NPL fixed point(s). The NPL algorithm failing to converge should therefore not be interpreted as evidence against the existence of a NPL fixed point since – as already mentioned above – the NPL mapping has at least one fixed point. Furthermore, even if the NPL algorithm may converge to different NPL fixed points when initiated at distinct starting values, it may still fail to obtain all NPL fixed points. The reasons for such lack of convergence and failure to deliver the full set of NPL fixed points will be presented in Section 5, after having formally investigated the properties of the NPL algorithm iterated in small samples.

Nonetheless, there are two cases where the NPL algorithm is guaranteed to converge to the consistent NPL estimator in every sample. First, the NPL algorithm always converges when applied to single-agent dynamic discrete choice models. For these models, convergence follows from the zero-Jacobian property shown by [Aguirregabiria and Mira \(2002, Proposition 2\)](#). Second, [Aguirregabiria and Mira \(2007, Example 5\)](#) highlight several appealing properties of the NPL algorithm when the number of structural parameters in θ is exactly equal to the number of free probabilities in \mathbf{P} . In fact, the NPL mapping then has a unique fixed point, this fixed point is stable, and the NPL algorithm reaches this consistent fixed point in only two iterations.

4.3 Spectral algorithm estimator

In this section, we propose an alternative iterative algorithm to compute the NPL estimator. More precisely, computing the NPL estimator requires solving a system of nonlinear equations. We now describe how a derivative-free non-monotone spectral residual method – proposed by [La Cruz, Martinez, and Raydan \(2006\)](#) – can be implemented to find these fixed points. In contrast to the NPL algorithm, the proposed method is not based on fixed point iterations and can therefore be used to find unstable fixed points.

Notice that $\hat{\phi}(\mathbf{P})$ is a stochastic function since $\hat{\varphi}(\mathbf{P})$ is stochastic through $\hat{\theta}(\mathbf{P})$. Finding the solution(s) to $\hat{\phi}(\mathbf{P}) = \mathbf{0}$, where $\hat{\phi}(\cdot) : \mathcal{P} \mapsto (-1, 1)^{\dim\{\mathcal{P}\}}$ is a nonlinear function with continuous partial derivatives, boils down to solving a large dimensional system of nonlinear equations. A derivative-free spectral approach is especially appealing in the current setting for two reasons. First, it does not require the computation, nor the approximation of the Jacobian of $\hat{\phi}(\cdot)$. Second, its small memory needs make it applicable to large-dimensional \mathbf{P} 's.

Spectral methods initially proposed by [Barzilai and Borwein \(1988\)](#) are most easily understood by comparing them with other existing methods commonly used to solve nonlinear systems of equations. For instance, Newton's method would update $\hat{\mathbf{P}}_k$ to $\hat{\mathbf{P}}_{k+1}$ as:

$$\hat{\mathbf{P}}_{k+1} = \hat{\mathbf{P}}_k - \left[\nabla \hat{\phi}(\hat{\mathbf{P}}_k) \right]^{-1} \hat{\phi}(\hat{\mathbf{P}}_k) \quad (12)$$

which unfortunately requires computing and inverting the Jacobian matrix $\nabla \hat{\phi}(\hat{\mathbf{P}}_k)$ at each iteration. Alternatively, quasi-Newton methods such as Broyden's would use:

$$\hat{\mathbf{P}}_{k+1} = \hat{\mathbf{P}}_k - \mathbf{B}_k^{-1} \hat{\phi}(\hat{\mathbf{P}}_k) \quad (13)$$

where the matrices \mathbf{B}_k are defined sequentially as,

$$\mathbf{B}_{k+1} = \mathbf{B}_k + \frac{\hat{\phi}(\hat{\mathbf{P}}_{k+1}) \Delta \hat{\mathbf{P}}'_{k+1}}{\Delta \hat{\mathbf{P}}'_{k+1} \Delta \hat{\mathbf{P}}_{k+1}} \quad (14)$$

with $\Delta \hat{\mathbf{P}}_{k+1} \equiv \hat{\mathbf{P}}_{k+1} - \hat{\mathbf{P}}_k$ and \mathbf{B}_0 is typically the identity matrix. While the latter avoids having to compute $\nabla \hat{\phi}(\hat{\mathbf{P}}_k)$, it still requires the inversion of a potentially large dimensional matrix. Spectral methods would rather propose to update the CCPs using:

$$\hat{\mathbf{P}}_{k+1} = \hat{\mathbf{P}}_k - \alpha_k \hat{\phi}(\hat{\mathbf{P}}_k) \quad (15)$$

where α_k is the spectral steplength. This spectral steplength can be interpreted as a Rayleigh quotient with respect to a secant approximation of the Jacobian of $\hat{\phi}(\cdot)$. More importantly, this object is a scalar. Different α_k 's have been proposed in the literature. Let $\Delta \hat{\phi}(\hat{\mathbf{P}}_k) \equiv \hat{\phi}(\hat{\mathbf{P}}_k) - \hat{\phi}(\hat{\mathbf{P}}_{k-1})$. [Barzilai and Borwein \(1988\)](#) proposed using:

$$\alpha_k = \frac{\Delta \hat{\mathbf{P}}'_k \Delta \hat{\phi}(\hat{\mathbf{P}}_k)}{\Delta \hat{\phi}(\hat{\mathbf{P}}_k)' \Delta \hat{\phi}(\hat{\mathbf{P}}_k)}. \quad (16)$$

Alternatively, [La Cruz, Martinez, and Raydan \(2006\)](#) proposed:

$$\alpha_k = \frac{\Delta \hat{\mathbf{P}}'_k \Delta \hat{\mathbf{P}}_k}{\Delta \hat{\mathbf{P}}'_k \Delta \hat{\phi}(\hat{\mathbf{P}}_k)}. \quad (17)$$

Finally, [Varadhan and Roland \(2008\)](#) proposed using:

$$\alpha_k = \text{sgn} \left(\Delta \hat{\mathbf{P}}'_k \Delta \hat{\phi}(\hat{\mathbf{P}}_k) \right) \frac{\left\| \Delta \hat{\mathbf{P}}_k \right\|}{\left\| \Delta \hat{\phi}(\hat{\mathbf{P}}_k) \right\|} \quad (18)$$

where $\text{sgn}(a) = a/|a|$ if $a \neq 0$; $\text{sgn}(a) = 0$ otherwise. Commonly used values of α_0 are

$$\alpha_0 = 1 \text{ or } \alpha_0 = \min \left\{ 1, 1/ \left| \hat{\phi}(\mathbf{P}_0) \right| \right\}.$$

A consequence of using spectral steplengths as the ones presented above is that they generate non-monotone updating processes. An important contribution of [La Cruz, Martinez, and Raydan \(2006\)](#) is that they propose a non-monotone line search as a globalization strategy that allows for such non-monotone behavior. More precisely, for some merit function $f(\cdot)$ such that $f(\mathbf{P}) = 0 \Leftrightarrow \hat{\phi}(\mathbf{P}) = \mathbf{0}$, each step must satisfy:

$$f(\hat{\mathbf{P}}_{k+1}) \leq \max_{0 \leq j \leq J-1} f(\hat{\mathbf{P}}_{k-j}) + \eta_k - \gamma \alpha_k^2 f(\hat{\mathbf{P}}_k) \quad (19)$$

where J is a non-negative integer, $0 < \gamma < 1$ and $\sum_k \eta_k < \infty$. This globalization strategy is a safeguard that can improve the performance of the algorithm. [La Cruz, Martinez, and Raydan \(2006, Theorem 4\)](#) show that if the point used to initialize the algorithm is close enough to some strongly isolated solution,⁶ the whole sequence converges to this solution.

There are three important remarks that are worth mentioning. First, since there is no guarantee that the NPL sample mapping has a unique fixed point, this iterative procedure should be initiated at different starting values. Second, it is worth emphasizing that the procedure described above is also iterative, just like the NPL algorithm. However, it does not iterate over the NPL mapping. For this reason, the spectral approach is able to find unstable fixed points that would not be attainable via the NPL algorithm. Third, it should be noted that the spectral approach is not much harder to code than NPL: one basically uses the code of a $K = 1$ step iteration to construct the system of nonlinear equations to be solved.

4.4 Comparing domains of attraction

The spectral algorithm described above is therefore an alternative approach that can be used to compute the NPL estimator. In order to appreciate how its convergence properties compare to the NPL algorithm and to other iterative approaches available in the literature, it is useful to highlight how a given solution's domain of attraction may vary across different methods.

There is an important consequence of using the NPL algorithm to find fixed points of the sample NPL mapping: if the stability condition required for convergence – which will be provided below – fails to hold, the relevant domain of attraction is a singleton. In other words, the only way to find the NPL fixed point corresponding to the NPL estimator would be to initiate the algorithm exactly at the solution. As a result, the domain of attraction of the NPL estimator has zero Lebesgue measure when stability fails. This feature is a consequence

⁶They consider a solution $\hat{\mathbf{P}}_*$ to be strongly isolated if $\hat{\phi}(\hat{\mathbf{P}}_*) = \mathbf{0}$ and there exists $\delta > 0$ such that $\|\mathbf{P} - \hat{\mathbf{P}}_*\| \leq \delta$ implies that $\hat{\phi}(\mathbf{P})' \left[\nabla \hat{\phi}(\mathbf{P}) \right]' \hat{\phi}(\mathbf{P}) \neq 0$.

of using fixed point iterations in an NPL mapping that – in the case of games – is not a contraction.

As discussed above, fixed point iterations are not the only method available to solve for the fixed points of a nonlinear system of equations. Under very general conditions – which hold in our problem – Newton or quasi-Newton algorithms have desirable local convergence properties. In fact, for these algorithms, the solution’s domain of attraction is a singleton if and only if the Jacobian matrix of the system of equations to solve is singular when evaluated at the solution. Under very mild conditions, the set of such unregular solutions has zero Lebesgue measure, such that the domain of attraction of any solution is generically larger than a singleton. This observation is true even if stability fails. It follows that, when stability fails, estimating algorithms based on Newton methods are associated with domains of attraction leading to a consistent estimator that are larger than the NPL algorithm’s.

This distinction between fixed point iterations and Newton methods is key to understand the properties of different algorithms used to estimate dynamic discrete games. For instance, the properties of the NFXP applied to dynamic games crucially depend on the method used to solve for the fixed point in the inner algorithm. If one uses fixed point iterations, instability of the best response mapping may prevent the inner algorithm to converge unless this algorithm is initiated at the fixed point of interest. In such cases, the inner algorithm’s domain of attraction being a singleton implies that the domain of attraction of the NFXP algorithm must also have zero Lebesgue measure. For this reason, it may be preferable to implement a NFXP algorithm that uses Newton methods to solve for the fixed point to ensure that the algorithm locally converges to the consistent estimator.

MPEC methods can also be interpreted as a set of variations of Newton methods used to find the solution of a constrained optimization problem. As a result, the MPEC algorithm applied to dynamic discrete games should locally converge to a solution. Of course, this algorithm may still fail to converge for a given initial value. However, the set of starting values leading to a consistent estimator has a strictly positive Lebesgue measure. In fact, simulation evidence from [Egesdal, Lai, and Su \(2015\)](#) suggests that the MPEC could always be initiated at a starting value leading to a converging sequence, despite the algorithm failing to converge for some sequences.

To sum up, one way to guarantee that the NPL algorithm locally converges to the NPL estimator would be to propose a Newton NPL algorithm that does not rely on fixed point iterations to find sample NPL fixed points. This is how the spectral algorithm described above should be interpreted. It is worth emphasizing a considerable advantage of the spectral approach over Newton and quasi-Newton algorithms that is especially relevant in the context of dynamic discrete games with large state spaces: replacing the approximation of the Jacobian to be inverted by a scalar. For instance, MPEC methods typically leverage sparsity properties to reduce the computational burden in large state spaces, but they still

require inverting matrices. Furthermore, the dimensionality of the problem remains smaller than implementing the MPEC estimator of [Egedal, Lai, and Su \(2015\)](#). While solving the system of nonlinear equations requires searching over the CCPs' space, MPEC is searching over the spaces of the CCPs, the parameters, and the Lagrange multipliers. For that reason, we expect the spectral approach to be feasible whenever MPEC is.

Should one conclude that the NPL algorithm is necessarily dominated by algorithms based on Newton or spectral methods? A single iteration of the NPL algorithm is computationally much cheaper than such alternative methods in very large state spaces. There may be relevant cases where even state of the art algorithms are not applicable. Better understanding under which conditions the NPL algorithm can be applied is therefore needed.

5 Convergence of the NPL algorithm in small samples

We now turn to assessing the convergence properties of the NPL algorithm when it is iterated in small samples. We first define stability of the data generating process as it is understood in the literature. Then, we investigate the convergence of the NPL algorithm iterated in small samples to some NPL fixed point ([Lemma 2](#)). This result is key to explain why the NPL algorithm may fail to converge even if the equilibrium generating the data is stable. Whether the NPL algorithm converges to the NPL estimator is closely related to the properties of the sample NPL mapping. We can therefore characterize the probability that the NPL algorithm delivers the NPL estimator ([Lemma 3](#)). Furthermore, we investigate how the NPL algorithm could lead to an estimator that is asymptotically different from the NPL estimator. We consider two cases. First, the algorithm could converge to an inconsistent fixed point of the sample NPL mapping ([Proposition 2](#)). Second, we show that the positive probability of no convergence of the algorithm introduces a convergence selection in the NPL algorithm estimator ([Proposition 3](#)). To the best of our knowledge, the implications of this selection have never been explored in the literature. The relative relevance of these two cases depends on the data generating process, as it will be highlighted in [Section 6](#).

[Kasahara and Shimotsu \(2012, Propositions 1 and 2\)](#) show that the spectral radius of $\nabla \varphi_{\mathbf{P}^0}^0$, denoted $\rho\left(\nabla \varphi_{\mathbf{P}^0}^0\right)$, is an important determinant of the convergence of the sequence of estimators generated by the NPL algorithm. The condition defining a stable population NPL mapping at the true CCPs is stated in [Definition 6](#).

Definition 6 (Stable population NPL mapping at true CCPs). Under [Assumptions 1-5](#), the data generating process is stable if $\rho\left(\nabla \varphi_{\mathbf{P}^0}^0\right) < 1$.

The condition established in [Definition 6](#) implies that there is a neighborhood of \mathbf{P}^0 – say $\mathcal{N}(\mathbf{P}^0)$ – such that a fixed point algorithm that starts at a vector $\mathbf{P}_0 \in \mathcal{N}(\mathbf{P}^0)$, and iterates in the population NPL mapping, converges to \mathbf{P}^0 . However, this stability condition in the

population NPL mapping is not sufficient to guarantee the convergence of the NPL algorithm in finite samples.

We now formally state Lemma 2 which establishes a sufficient condition for the NPL algorithm to converge to some fixed point of the sample NPL mapping in finite samples.

Lemma 2 (Convergence of the NPL algorithm in finite samples). *Consider a finite sample size M . Let $\hat{\mathbf{P}}_*$ be a fixed point of the sample NPL mapping. Suppose that $\rho\left(\nabla\hat{\varphi}_{(\hat{\mathbf{P}}_*)}\right) < 1$. By continuity of $\rho(\cdot)$, there is a neighborhood of $\hat{\mathbf{P}}_*$, say $\mathcal{N}\left(\hat{\mathbf{P}}_*\right)$, such that $\rho\left(\nabla\hat{\varphi}_{(\mathbf{P})}\right) < 1$ for any $\mathbf{P} \in \mathcal{N}\left(\hat{\mathbf{P}}_*\right)$. Then, the NPL algorithm converges to $\hat{\mathbf{P}}_*$ if $\hat{\mathbf{P}}_0 \in \mathcal{N}\left(\hat{\mathbf{P}}_*\right)$.*

Proof. See Appendix A.3. □

Convergence of the NPL algorithm is studied by Kasahara and Shimotsu (2012, Proposition 1) under the condition that the sample size is large enough such that: (i) the sample NPL mapping can be approximated – arbitrarily well – using the population NPL mapping; and (ii) a consistent estimator $\hat{\mathbf{P}}_0$ is arbitrarily close to \mathbf{P}^0 . This is the reason why their convergence conditions are stated in terms of $\nabla\varphi_{(\mathbf{P}^0)}^0$ as in Definition 6. Our Lemma 2 relaxes such conditions and allows the algorithm to be iterated in small samples. Our approach echoes criticisms – formulated by Pesendorfer and Schmidt-Dengler (2010) – to the asymptotic approximation. For these reasons, the convergence properties of the NPL algorithm in Lemma 2 are more in line with the idea of using the NPL algorithm as a device to reduce the finite sample bias.

An important consequence of considering a fixed M is that the updating procedure is now iterated over the sample NPL mapping. The convergence of the NPL algorithm to a given NPL fixed point does not (directly) depend on the stability of the data generating process – as stated in Definition 6 – but rather on the contraction properties of the sample NPL mapping evaluated at this fixed point. More precisely, it depends on $\rho\left(\nabla\hat{\varphi}_{(\hat{\mathbf{P}}_*)}\right)$. Furthermore, if this fixed point corresponds to the NPL estimator – to the couple $(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)$ that maximizes the pseudo likelihood function among the set of NPL fixed points – it follows that one can deduce conditions which must be satisfied for the NPL algorithm iterated in small samples to deliver the NPL estimator.

To alleviate the notation, let $\hat{\rho}_{\text{NPL}} \equiv \rho\left(\nabla\hat{\varphi}_{(\hat{\mathbf{P}}_{\text{NPL}})}\right)$ and $\rho^0 \equiv \rho\left(\nabla\varphi_{(\mathbf{P}^0)}^0\right)$. Since the sample NPL mapping is a random object, one cannot guarantee that the stability of the equilibrium generating the data implies that $\hat{\rho}_{\text{NPL}} < 1$. The randomness of the stability condition for the sample NPL mapping implies that there is always a strictly positive probability that $\hat{\rho}_{\text{NPL}} > 1$, even when the population is such that $\rho^0 < 1$. This observation can explain the lack of convergence of the NPL algorithm when the data generating process is stable.

Lemma 3 establishes the asymptotic distribution of the statistic $\hat{\rho}_{\text{NPL}}$. The result is stated as follows. It is worth mentioning that the required differentiability of the spectral radius

operator holds under mild conditions.⁷

Lemma 3 (Asymptotic distribution of the sample spectral radius at NPL estimates). *Suppose that the spectral radius function $\rho(\cdot)$ is differentiable at the Jacobian matrices $\nabla \hat{\boldsymbol{\varphi}}_{(\mathbf{P}^0)} = \nabla_{\mathbf{P}} \boldsymbol{\Psi}(\hat{\boldsymbol{\theta}}_{(\mathbf{P}^0), \mathbf{P}^0}) + \nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}(\hat{\boldsymbol{\theta}}_{(\mathbf{P}^0), \mathbf{P}^0}) \nabla \hat{\boldsymbol{\theta}}_{(\mathbf{P}^0)}$ and $\nabla \boldsymbol{\varphi}_{(\mathbf{P}^0)}^0 = \nabla_{\mathbf{P}} \boldsymbol{\Psi}(\boldsymbol{\theta}^0_{(\mathbf{P}^0), \mathbf{P}^0}) + \nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}(\boldsymbol{\theta}^0_{(\mathbf{P}^0), \mathbf{P}^0}) \nabla \boldsymbol{\theta}^0_{(\mathbf{P}^0)}$, where $\nabla \hat{\boldsymbol{\theta}}_{(\mathbf{P}^0)}$ and $\nabla \boldsymbol{\theta}^0_{(\mathbf{P}^0)}$ are the Jacobians of $\hat{\boldsymbol{\theta}}(\mathbf{P})$ and $\boldsymbol{\theta}^0(\mathbf{P})$ at \mathbf{P}^0 . Then, as $M \rightarrow \infty$:*

$$\sqrt{M} (\hat{\rho}_{NPL} - \rho^0) \xrightarrow{d} \mathbf{G}^0 \mathbf{Z}(\mathbf{P}^0) + \mathbf{J}^0 \mathbf{L}(\mathbf{P}^0)$$

where \mathbf{G}^0 and \mathbf{J}^0 are deterministic matrices defined in Appendix A.4; $\mathbf{Z}(\mathbf{P}^0)$ and $\mathbf{L}(\mathbf{P}^0)$ are the vectors of normal random variables defined in Lemma 1. Therefore, as $M \rightarrow \infty$, the probability that the sample NPL mapping is stable at $\hat{\mathbf{P}}_{NPL}$ is:

$$\Pr(\hat{\rho}_{NPL} < 1) = \Phi\left(\frac{\sqrt{M}[1 - \rho^0]}{\sqrt{\text{Var}[\hat{\rho}_{NPL}]}}\right)$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal.

Proof. See Appendix A.4. □

In Lemma 3, the expression for the probability of a stable NPL mapping is such that, as $M \rightarrow \infty$: it goes to 1 if the equilibrium in the population is stable – i.e., if $\rho^0 < 1$; and it goes to 0 if the equilibrium in the population is unstable – i.e., if $\rho^0 > 1$.

It is worth briefly discussing how Lemma 3 relates to single-agent dynamic discrete choice models. Single-agent models can be interpreted as special cases of games with strategic interactions being 0 in the data generating process. By the zero-Jacobian property (Aguirregabiria and Mira, 2002, Proposition 2), both ρ^0 and $\hat{\rho}_{NPL}$ are equal to 0 in such models. Therefore, the distribution of $\hat{\rho}_{NPL}$ should be degenerate at 0 as a consequence of strategic interactions being constrained to 0 in the estimation of single-agent models. However, when estimating an empirical game, strategic interactions are not constrained to 0 even if their true value in the population may be 0 and $\rho^0 = 0$. In other words, even if $\rho^0 = 0$, not restricting the value of strategic interactions implies that the distribution of $\hat{\rho}_{NPL}$ is not degenerate, as suggested by Lemma 3.

Let $(\hat{\boldsymbol{\theta}}_K, \hat{\mathbf{P}}_K)$ be the K -step NPL algorithm estimator. As a corollary of Lemma 2, we have that if the sample is such that $\hat{\rho}_{NPL} < 1$, then there exists a large enough $\bar{K} \in \mathbb{N}$ and a starting value $\hat{\mathbf{P}}_0$ such that for any $K \geq \bar{K}$, $\hat{\boldsymbol{\theta}}_K = \hat{\boldsymbol{\theta}}_{NPL}$. That is, under this condition the NPL algorithm returns the NPL estimator. Given this result, one could be tempted to claim that – under the condition $\hat{\rho}_{NPL} < 1$ – the NPL algorithm estimator has the same (good) asymptotic properties as the NPL estimator. However, one must take into account the

⁷In fact, eigenvalues are continuous and differentiable functions of the elements of a matrix. The spectral radius is the maximum of the absolute values of these eigenvalues and the set of matrices for which there is a tie between two or more eigenvalues for the maximum has zero Lebesgue measure.

randomness introduced by this stochastic stability condition when defining the properties of the K -step NPL algorithm estimator. For the rest of this section, we study how the randomness of the stability condition of the sample NPL mapping affects the asymptotic properties of the NPL algorithm estimator. Propositions 2 and 3 characterize the two sources of asymptotic bias of the NPL algorithm estimator.

For a given sample, we can distinguish three different scenarios for the convergence of the NPL algorithm: (i) the sample stability conditions holds such that the algorithm converges to the NPL estimator; (ii) the sample stability condition does not hold, but the algorithm converges to a fixed point of the sample NPL mapping – say, $\hat{\theta}_*$ – which is not the NPL estimator; and (iii) the sample stability condition does not hold, the algorithm generates cycles and it never converges. These possible results have the following implications for the NPL algorithm estimator, as $K \rightarrow \infty$:

$$\hat{\theta}_K = \begin{cases} \hat{\theta}_{\text{NPL}} & \text{if the NPL algorithm converges and } \hat{\rho}_{\text{NPL}} < 1 \\ \hat{\theta}_* \neq \hat{\theta}_{\text{NPL}} & \text{if the NPL algorithm converges and } \hat{\rho}_{\text{NPL}} \geq 1 \\ \text{is not well defined} & \text{if the NPL algorithm fails to converge} \end{cases} \quad (20)$$

where we treat the K -step NPL estimator that fails to converge when $K \rightarrow \infty$ as not well defined since its properties would vary with K , even for very large values of K .

Expression (20) illustrates two sources of discrepancy between $\hat{\theta}_K$ and $\hat{\theta}_{\text{NPL}}$, which correspond to two sources of bias in the NPL algorithm estimator that may arise when the sample NPL fixed point does not satisfy the stability condition. First, if $\hat{\rho}_{\text{NPL}} \geq 1$ the algorithm may still converge to a fixed point of the sample NPL mapping which satisfies the stability condition. In such cases, the NPL algorithm converges to an inconsistent estimator. Second, the randomness in $\hat{\rho}_{\text{NPL}}$ explicated in Lemma 3 implies that two samples drawn from the same data generating process – and therefore associated with the same value of ρ^0 – may not equally satisfy the convergence condition $\hat{\rho}_{\text{NPL}} < 1$. In other words, the randomness in the sample NPL mapping introduces a convergence selection bias.

Another way to highlight this selection problem is to consider the following unfeasible estimator:

$$\tilde{\theta}_{\text{NPL}} = \begin{cases} \lim_{K \rightarrow \infty} \hat{\theta}_K = \hat{\theta}_{\text{NPL}} & \text{if } \hat{\rho}_{\text{NPL}} < 1 \\ \text{is not well defined} & \text{if } \hat{\rho}_{\text{NPL}} \geq 1 \end{cases} \quad (21)$$

which is unfeasible precisely because the stability condition $\hat{\rho}_{\text{NPL}} < 1$ is not observable to the researcher. Since $\hat{\theta}_{\text{NPL}}$ and $\hat{\rho}_{\text{NPL}}$ are not independent random variables, the sampling distribution of $\tilde{\theta}_{\text{NPL}}$ must be different from the one of the NPL estimator.

Proposition 2 focuses on the bias introduced by convergence to an inconsistent estimator: for large enough K , $\hat{\theta}_K = \hat{\theta}_* \neq \hat{\theta}_{\text{NPL}}$. For the sake of illustration and to isolate this source

of bias we assume that the NPL algorithm always converges to $\hat{\boldsymbol{\theta}}_*$ when $\hat{\rho}_{\text{NPL}} \geq 1$.

Proposition 2 (Convergence to an inconsistent estimator). *Let Assumptions 1-5 be satisfied. Suppose that – with probability 1 – the NPL algorithm converges to some $\hat{\boldsymbol{\theta}}_*$ whenever $\hat{\rho}_{\text{NPL}} \geq 1$. Define \mathbf{W}_* as $\sqrt{M} (\hat{\boldsymbol{\theta}}_* - \hat{\boldsymbol{\theta}}_{\text{NPL}}) \xrightarrow{d} \mathbf{W}_*$. Then, there exists a large enough $\bar{K} \in \mathbb{N}$ such that $\forall K \geq \bar{K}$, $\hat{\boldsymbol{\theta}}_K = \hat{\boldsymbol{\theta}}_{\text{NPL}} + \mathbb{1} \{ \hat{\rho}_{\text{NPL}} \geq 1 \} [\hat{\boldsymbol{\theta}}_* - \hat{\boldsymbol{\theta}}_{\text{NPL}}]$ and:*

$$\sqrt{M} (\hat{\boldsymbol{\theta}}_K - \boldsymbol{\theta}^0) \xrightarrow{d} \sqrt{M} (\hat{\boldsymbol{\theta}}_{\text{NPL}} - \boldsymbol{\theta}^0) + \mathbf{W}_* \left[1 - \Phi \left(\frac{\sqrt{M} [1 - \rho^0]}{\sqrt{\text{Var}[\hat{\rho}_{\text{NPL}}]}} \right) \right].$$

Proof. See Appendix A.5. □

The second term in the asymptotic distribution of $\sqrt{M} (\hat{\boldsymbol{\theta}}_K - \boldsymbol{\theta}^0)$ stated in Proposition 2 represents the effect of the possibility of converging to an inconsistent estimator when the sample NPL mapping is unstable. It depends on the probability of that event. Even if $\rho^0 < 1$, this term contributes to the asymptotic distribution of the NPL algorithm estimator.

Proposition 3 focuses on the bias introduced by the convergence selection problem: the estimator $\hat{\boldsymbol{\theta}}_K$ is not well defined when the algorithm does not converge. To isolate this source of bias, we assume that the algorithm never converges when $\hat{\rho}_{\text{NPL}} \geq 1$.

Proposition 3 (Convergence selection). *Let Assumptions 1-5 be satisfied. Suppose that – with probability 1 – the NPL algorithm never converges whenever $\hat{\rho}_{\text{NPL}} \geq 1$. Then, there exists a large enough $\bar{K} \in \mathbb{N}$ such that $\forall K \geq \bar{K}$, any well-defined NPL algorithm estimator satisfies $\hat{\boldsymbol{\theta}}_K = \tilde{\boldsymbol{\theta}}_{\text{NPL}}$ and:*

$$\sqrt{M} (\hat{\boldsymbol{\theta}}_K - \boldsymbol{\theta}^0) \xrightarrow{d} \sqrt{M} (\hat{\boldsymbol{\theta}}_{\text{NPL}} - \boldsymbol{\theta}^0) \mathbb{1} \left\{ \sqrt{M} (\hat{\rho}_{\text{NPL}} - \rho^0) < \sqrt{M} (1 - \rho^0) \right\}.$$

Proof. See Appendix A.6. □

The asymptotic distribution of $\sqrt{M} (\hat{\boldsymbol{\theta}}_K - \boldsymbol{\theta}^0)$ stated in Proposition 3 corresponds to a truncated normal distribution with the variable defining the truncation distributed as in Lemma 3. As $M \rightarrow \infty$, $\Pr(\hat{\rho}_{\text{NPL}} < 1) \rightarrow 1$ if the population NPL fixed point corresponding to the data generating process is stable, such that the distribution of the well-defined estimator is asymptotically equivalent to the NPL estimator's distribution. However, if the equilibrium generating the data is unstable, $\Pr(\hat{\rho}_{\text{NPL}} < 1) \rightarrow 0$ as $M \rightarrow \infty$ such that the asymptotic distribution does not exist, which is consistent with the estimator not being well-defined in this case.

The results stated in Propositions 2 and 3 lead to interesting conclusions that shed light on convergence issues associated with the NPL algorithm. These implications, which are supported by the simulation evidence reported in Section 6, are summarized as follows.

First, $\rho^0 < 1$ is not sufficient for convergence when the iterative procedure is implemented if finite samples. In fact, it fails to account for the randomness in the mapping used to update the CCPs.

Second, the possibility of no convergence introduces a discrepancy between the NPL algorithm estimator and the NPL estimator – even in the unfeasible case where the researcher would know $\hat{\rho}_{\text{NPL}}$ and could select the NPL algorithm estimator $\hat{\theta}_K$ only when the convergence condition $\hat{\rho}_{\text{NPL}} < 1$ holds.

Third, one should expect convergence to the NPL estimator to be a much relevant issue in cases where the data generating process is stable, but close to being unstable. For fixed M , the randomness in the sample NPL mapping may make the iterative procedure unstable even if it would be stable if one could update the CCPs using the mapping in the population.

Fourth, one should expect the randomness in the sample NPL mapping to be less concerning when $M \rightarrow \infty$, even in stable data generating processes that are close to being unstable. In other words, the NPL algorithm should deliver the NPL estimator with large samples even if the data generating process is almost unstable.

Finally, if one could implement the NPL algorithm by updating the CCPs using the population NPL mapping, then $\rho^0 < 1$ would be sufficient for the NPL algorithm to converge, even in samples of modest size.

6 Monte Carlo experiments

6.1 Data-generating process

The data generating process used for our simulations is based on Experiment 3 in [Aguirregabiria and Mira \(2007\)](#). The same data generating process, or a simplified version of it, was also used by [Kasahara and Shimotsu \(2012\)](#) and by [Egedal, Lai, and Su \(2015\)](#).

The model corresponds to the dynamic game of market entry and exit presented in Section 3 above. A number $N = 5$ of firms must decide whether or not to operate in M independent markets. There are $N + 1 = 6$ observable state variables: the market size (S_{mt}) and the incumbency status of each player ($\mathbf{y}_{m,t-1}$). Let the support of S_{mt} be $\mathcal{S} = \{1, 2, 3, 4, 5\}$. The observable state space is denoted by $\mathcal{X} \equiv \mathcal{S} \times \mathcal{Y}^N$ and the total number of possible different states is given by $\dim(\mathcal{X}) = 5 \times 2^5 = 160$.

A transition probability distribution has to be specified for the vector of observable state variables. The transition probability of incumbency statuses trivially corresponds to the firms' decisions CCPs. Let the transition matrix for market size S_{mt} be given by:




Figure 1: Spectral radii for different θ_{RN} – Population NPL mapping

$$\begin{bmatrix} 0.8 & 0.2 & 0 & 0 & 0 \\ 0.2 & 0.6 & 0.2 & 0 & 0 \\ 0 & 0.2 & 0.6 & 0.2 & 0 \\ 0 & 0 & 0.2 & 0.6 & 0.2 \\ 0 & 0 & 0 & 0.2 & 0.8 \end{bmatrix}. \quad (22)$$

Let $\theta_{RS} = 1$, $\theta_{EC} = 1$, $\theta_{FC,1} = 1.9$, $\theta_{FC,2} = 1.8$, $\theta_{FC,3} = 1.7$, $\theta_{FC,4} = 1.6$ and $\theta_{FC,5} = 1.5$. The discount rate is assumed to be known and equal to $\beta = 0.95$. With the values of these parameters fixed, the stability of the data generating process is determined by the strategic interaction parameter θ_{RN} . Notice that $\theta_{RN} = 0$ would correspond to a single-agent dynamic discrete choice model and stability in the sense of Definition 6 is guaranteed through the zero-Jacobian property (Aguirregabiria and Mira, 2002, Proposition 2). The data generating process will therefore be stable or unstable for positive values of θ_{RN} . Spectral radii ρ^0 for different values of θ_{RN} are graphed in Figure 1.⁸

Simulation exercises available in the literature – that have used this data generating process – have typically focused on two specific values of θ_{RN} : $\theta_{RN} = 2$ – which will be referred to as the *mildly stable* case since the corresponding spectral radius at the true CCPs is 0.9237 – and $\theta_{RN} = 4$ – which will be the *very unstable* case with spectral radius of 1.6748. We present below Monte Carlo experiments and convergence results for these two cases.

⁸The Jacobians used to compute these spectral radii have been approximated using numerical derivatives.

Furthermore, to provide a finer picture of the properties of the sequential estimators, we also consider a *very stable* case by setting $\theta_{RN} = 1$ – corresponding spectral radius 0.4623 – and a *mildly unstable* case where $\theta_{RN} = 2.4$ – spectral radius equal to 1.1168. These four cases allow us to study how different degrees of (in)stability – at the population level – affect the statistical properties of the sample NPL mapping and the estimators.

Besides potentially affecting the convergence properties of the NPL algorithm, these different values of the strategic interaction parameter have important economic implications. In order to appreciate how the strategic interaction parameter affects the level of competition, we simulated firms’ decisions in 50000 markets drawn from the ergodic distribution of the state variables. Summary statistics computed over these simulated markets are reported in Table 1. Increasing the value of the strategic interaction parameter considerably reduces the average number of active firms, generating larger reductions in the probabilities of being active for the firms with larger fixed costs. Interestingly, the effect on the number of entry and exits is non-monotonic: these numbers are larger for the “mildly stable” and the “mildly unstable” cases than the other two data generating processes.

Table 1: Competition statistics for the different data generating processes

	Very stable	Mildly stable	Mildly unstable	Very unstable
Number of active firms				
Average	2.7652	1.9939	1.8451	1.2225
Std. dev.	1.6622	1.4320	1.3635	1.0024
AR(1) parameter for number of active firms	0.7070	0.5691	0.5300	0.3519
Average number of entries	0.6917	0.7473	0.7380	0.5492
Average number of exits	0.6933	0.7569	0.7448	0.5558
Average excess turnover	0.4600	0.5110	0.4972	0.2879
Correlation between entries and exits	-0.1743	-0.2225	-0.2201	-0.1854
Probabilities of being active				
Firm 1	0.4993	0.3222	0.2866	0.1239
Firm 2	0.5222	0.3552	0.3219	0.1457
Firm 3	0.5536	0.3975	0.3662	0.1871
Firm 4	0.5797	0.4363	0.4098	0.2689
Firm 5	0.6103	0.4827	0.4606	0.4968

Notes: Statistics computed using $M = 50000$ markets drawn from the ergodic distribution of the state variables. Excess turnover defined as $(\# \text{ entries} + \# \text{ exits}) - \text{abs}(\# \text{ entries} - \# \text{ exits})$.

The equilibrium CCPs needed to generate the data solve the fixed point mapping in (1) when the unobservable state variables follow the extreme value type 1 distribution. The equilibrium CCPs can be found by using a built-in solver for a system of nonlinear equations.⁹ For all our data generating processes, we select the equilibrium that this solver returns when we initialize it with a vector of CCPs equal to 0.5 for all players and states.

For each data generating process, we take a random sample of M markets from the ergodic distribution of the state variables, and use the equilibrium CCPs to generate firms’ entry/exit

⁹We used R’s BBSolve function with its default settings. Similar solvers can be found in other softwares.

decisions in each market. We implement experiments with $M = 400$ and with $M = 5000$ markets. For each experiment, we draw 500 Monte Carlo samples.

In order to illustrate the theoretical results derived in previous sections, we compute and compare estimates obtained from the NPL and the spectral algorithms. In particular, we iterate both algorithms $K = 100$ times, regardless whether convergence is reached, and study the properties of the sequences of estimates for various subsets of samples to separate the different sources of bias identified in Propositions 2 and 3. Both the NPL and the spectral algorithms were initialized with 5 different starting values, including the frequency estimator¹⁰ of the CCPs as one of them. For a given algorithm, if different starting points generate different estimates, then the solution with the highest value of the pseudo likelihood is selected as the estimator using that algorithm.

The spectral approach estimator was implemented in two different ways. First, we considered a simple version using the spectral steplength proposed by [La Cruz, Martinez, and Raydan \(2006\)](#), without their non-monotone line search. This version was used to reduce the computational burden when comparing the sequence of estimates generated by the spectral approach with the sequence generated by the NPL algorithm. Second, spectral estimates are obtained using R's BBSolve function ([Varadhan and Gilbert, 2009](#)) which incorporates the non-monotone line search along with other safeguards. In each model specification (very stable, mildly stable, mildly unstable, and very unstable) and each sample size considered, the BBSolve algorithm only failed to converge in at most 0.4% Monte Carlo samples. Interestingly, the estimates obtained from the simple version of the spectral approach coincided with the more involved implementation never less than in 98.6% Monte Carlo samples.

6.2 Results

6.2.1 Comparing convergence

We first assess whether, upon convergence, the NPL and the spectral algorithms deliver the NPL estimator. More precisely, for each Monte Carlo sample, we compute the fixed points that maximize the log-likelihood for the converged NPL algorithm and the spectral solver. We call the NPL estimator the estimates that correspond to the fixed point that maximizes the log-likelihood across both estimators (i.e., the best log-likelihood). Then, we compare the value of the log-likelihood function achieved by each estimator to this best log-likelihood. We consider that an estimator has reached the NPL estimator if it delivers a value of the log-likelihood function that is within 10^{-5} of the best log-likelihood. The results are reported in Table 2. We find that whenever the spectral solver converges (which happens in at least 99.6% of the samples in each experiment) it delivers the NPL estimator. Furthermore, whenever the NPL algorithm converges (which happens almost always in the very stable case, but almost

¹⁰Estimated zeros and ones are respectively replaced with 10^{-10} and $1 - 10^{-10}$.

never in the very unstable case), it delivers the NPL estimator. Of course, one caveat of defining the NPL estimator as we do is that we cannot rule out the case where none of the estimators we consider actually finds the NPL fixed point that truly maximizes the log-likelihood function. However, the nice properties of the spectral solver estimator that will be presented below suggest that the spectral algorithm succeeds in finding the NPL estimator in our simulation experiments.

Table 2: Convergence to the NPL estimator

	Very stable		Mildly stable		Mildly unstable		Very unstable	
	400	5K	400	5K	400	5K	400	5K
% NPL algorithm = NPL estimator	92.4	99.6	57.6	74.0	43.6	29.2	3.2	0.0
% Spectral solver = NPL estimator	100.0	99.8	100.0	99.6	99.8	100.0	99.6	99.6
% Both = NPL estimator	92.4	99.6	57.6	74.0	43.6	29.2	3.2	0.0

Notes: Percentages computed over 500 Monte Carlo samples. For each sample, the NPL estimates correspond to the fixed point of the sample NPL mapping that maximizes the log-likelihood function among all fixed points found by the NPL algorithm and the spectral solver. An estimator is deemed converging to the NPL estimator if it generates a value of the log-likelihood function that is within 10^{-5} of the NPL estimator’s log-likelihood.

Table 3: Sample stability and convergence

	Very stable		Mildly stable		Mildly unstable		Very unstable	
	400	5K	400	5K	400	5K	400	5K
Spectral radius < 1	95.8	99.6	63.8	88.2	50.0	47.0	4.6	0.0
% NPL algo converged	92.4	99.4	57.2	73.6	43.4	29.2	3.2	0.0
% NPL algo did not converged	3.4	0.2	6.6	14.6	6.6	17.8	1.4	0.0
% Spectral algo converged	95.2	99.6	63.6	87.6	49.6	46.8	4.6	0.0
% Spectral algo did not converged	0.6	0.0	0.2	0.6	0.4	0.2	0.0	0.0
Spectral radius ≥ 1	4.2	0.2	36.2	11.4	49.8	53.0	95.0	99.6
% NPL algo converged	0.0	0.2	0.4	0.4	0.2	0.0	0.0	0.0
% NPL algo did not converged	4.2	0.0	35.8	11.0	49.6	53.0	95.0	99.6
% Spectral algo converged	4.2	0.2	35.2	10.8	49.8	52.8	94.8	99.2
% Spectral algo did not converged	0.0	0.0	1.0	0.6	0.0	0.2	0.2	0.4
Spectral solver did not converge	0.0	0.2	0.0	0.4	0.2	0.0	0.4	0.4

Notes: Percentages computed over 500 Monte Carlo samples. “Spectral radius” refers to $\hat{\rho}_{\text{NPL}}$. The NPL and spectral algorithms are deemed having converged if $\max\left\{\left|\hat{\theta}_{100} - \hat{\theta}_{99}\right|\right\} < 10^{-5}$.

In Table 3, we compare the convergence frequencies of the NPL and the spectral algorithms, broken down with respect to the value of the spectral radius of the Jacobian of the sample NPL mapping evaluated at $\hat{\mathbf{P}}_{\text{NPL}}$.¹¹ For both algorithms, we generated a sequence of $K = 100$ estimates and we consider the sequence to have converged if $\max\left\{\left|\hat{\theta}_{100} - \hat{\theta}_{99}\right|\right\} < 10^{-5}$. The simulation results provide evidence supporting that the convergence properties of the NPL algorithm iterated in fixed samples are driven by the spectral radius of the sample

¹¹Given our results in Table 2, the spectral solver estimates are used as $\hat{\mathbf{P}}_{\text{NPL}}$.

NPL mapping evaluated at the NPL estimator. In fact, the fraction of samples for which the NPL algorithm fails to converge and the share of $\hat{\rho}_{\text{NPL}}$ being greater than 1 both increase when the data generating process is further from being stable. When the data generating process is “mildly stable”, the NPL algorithm is more likely to converge and $\hat{\rho}_{\text{NPL}}$ is more likely to be smaller than 1 in larger samples. One may notice that relatively large fractions of Monte Carlo samples are such that the NPL algorithm fails to converge at $K = 100$ despite $\hat{\rho}_{\text{NPL}} < 1$. However, further investigation of these cases showed that the $K = 99$ and the $K = 100$ estimates are actually very close to each other (exactly the same up to the second decimal for all parameters) despite not satisfying the convergence criterion. Interestingly, these estimates are also associated with values of sample spectral radius that are relatively close to 1 (greater than 0.9 on average).

In Propositions 2 and 3, we highlighted two potential sources of bias for the estimates obtained from the converged NPL algorithm. On one hand, Proposition 2 implies that the NPL algorithm could converge to an inconsistent estimator when $\hat{\rho}_{\text{NPL}} \geq 1$. This would be the case if there exist some stable fixed points of the sample NPL mapping in addition to the unstable NPL fixed points defining the NPL estimator. According to Table 3, the NPL algorithm almost never converges when $\hat{\rho}_{\text{NPL}} \geq 1$. It follows that convergence to an inconsistent estimator is unlikely to be a convincing explanation of the properties of the NPL algorithm estimates in our experiments. On the other hand, Proposition 3 suggests that the converged NPL algorithm sequences represent a selected subset of sequences that may have properties that differ from those of the NPL estimator. Table 3 shows that the NPL algorithm does not converge in a non-negligible share of samples even when the data generating process is stable – 25.6% in the mildly stable case with $M = 5000$ – and it does converge in a non-negligible proportion of samples when the population is unstable – 29.2% in the mildly unstable case with $M = 5000$. It seems that convergence selection is a suitable candidate to explain how the NPL algorithm estimator may differ from the consistent NPL estimator.

In order to appreciate how the distribution of $\hat{\rho}_{\text{NPL}}$ varies with ρ^0 we drew boxplots of $\hat{\rho}_{\text{NPL}}$ from 50 Monte Carlo samples of $M = 400$ markets for the same grid of data generating processes as in Figure 1. The results, which provide evidence supporting Lemma 3, are summarized in Figure 2.¹² As expected, $\hat{\rho}_{\text{NPL}} < 1$ is a more likely event in more stable data generating processes. Once again, it is worth emphasizing that a large fraction of the sample spectral radii can be on both sides of 1 when the data generating process is close to the stability cut-off.

An important feature of the spectral algorithm that is observable from Table 3 is that it almost always converges, regardless of the data generating process and the value of the sample

¹²For $\theta_{\text{RN}} = 1.4$, there was a value of $\hat{\rho}_{\text{NPL}}$ which was close to 6. We dropped this outlier when drawing Figure 2 to make the figure easier to read.




Figure 2: Boxplots of sample spectral radii at NPL estimates – $M = 400$

spectral radius. This observation confirms the importance of separating the properties of the NPL estimator itself from the properties of the NPL algorithm. The definition of the NPL estimator does not depend on the stability of the sample NPL fixed point: it exists even if the NPL algorithm may fail to deliver it.

6.2.2 Average estimates and standard errors

Tables 4 and 5 report the averages and standard errors for five estimators of interest, respectively for $M = 400$ and $M = 5000$, which are used to illustrate Propositions 2 and 3. First of all, both for the NPL and the spectral algorithms, we report the properties of the estimates obtained at $K = 100$ regardless whether the sequence of estimates has converged or not. Such estimators always exist, but their properties may be negatively affected by the fact that they contain estimates obtained from non-converging sequences. Then, we report the properties of the estimates obtained from the converged NPL algorithm – where a sequence is deemed having converged at $K = 100$ if $\max \left\{ \left| \hat{\theta}_{100} - \hat{\theta}_{99} \right| \right\} < 10^{-5}$ – and the spectral solver. Finally, we consider a conditional estimator which corresponds to the $K = 100$ NPL algorithm estimator when the spectral radius of the Jacobian of the sample NPL mapping at the estimates obtained from the spectral solver is smaller than 1.

We find that restricting the NPL algorithm’s estimates to converged sequences can introduce a substantial bias. What’s more, notice that none of the sequences converged in the “very unstable” case with $M = 5000$. However, the results are much more encouraging for

Table 4: Simulation results – $M = 400$

	$\theta_{RS} = 1$	θ_{RN}	$\theta_{EC} = 1$	$\theta_{FC,1} = 1.9$	$\theta_{FC,2} = 1.8$	$\theta_{FC,3} = 1.7$	$\theta_{FC,4} = 1.6$	$\theta_{FC,5} = 1.5$
<i>All $K = 100$ NPL algorithm estimates</i>								
Very stable ($\theta_{RN} = 1$)	1.0355 (0.2217)	1.0882 (0.6761)	0.9990 (0.1177)	1.9058 (0.2304)	1.8206 (0.2313)	1.7124 (0.2181)	1.6076 (0.2055)	1.5156 (0.1998)
Mildly stable ($\theta_{RN} = 2$)	0.9603 (0.1953)	1.8406 (0.6930)	1.0181 (0.1098)	1.9082 (0.2218)	1.8046 (0.2130)	1.7043 (0.2062)	1.6004 (0.2069)	1.4957 (0.2059)
Mildly unstable ($\theta_{RN} = 2.4$)	0.9154 (0.1640)	2.0475 (0.6134)	1.0476 (0.1133)	1.8998 (0.2291)	1.7899 (0.2165)	1.6834 (0.2100)	1.5802 (0.2077)	1.4721 (0.2083)
Very unstable ($\theta_{RN} = 4$)	0.7892 (0.0869)	2.7388 (0.2833)	1.2489 (0.1105)	1.9232 (0.2769)	1.8002 (0.2620)	1.6744 (0.2491)	1.5110 (0.2426)	1.2762 (0.2311)
<i>All $K = 100$ spectral algorithm estimates</i>								
Very stable ($\theta_{RN} = 1$)	1.0416 (0.2377)	1.1071 (0.7264)	0.9973 (0.1204)	1.9043 (0.2328)	1.8192 (0.2337)	1.7113 (0.2205)	1.6064 (0.2075)	1.5148 (0.2014)
Mildly stable ($\theta_{RN} = 2$)	1.0308 (0.2922)	2.0941 (1.0461)	0.9869 (0.1383)	1.9209 (0.2300)	1.8212 (0.2220)	1.7240 (0.2189)	1.6244 (0.2232)	1.5258 (0.2269)
Mildly unstable ($\theta_{RN} = 2.4$)	1.0135 (0.2722)	2.4333 (1.0528)	0.9951 (0.1509)	1.9230 (0.2416)	1.8200 (0.2318)	1.7192 (0.2289)	1.6218 (0.2335)	1.5230 (0.2429)
Very unstable ($\theta_{RN} = 4$)	0.9818 (0.1591)	3.8546 (0.7949)	1.0316 (0.1793)	1.9337 (0.3046)	1.8250 (0.2882)	1.7229 (0.2769)	1.6106 (0.2762)	1.4913 (0.2989)
<i>Converged $K = 100$ NPL algorithm estimates</i>								
Very stable ($\theta_{RN} = 1$)	1.0038 (0.1949)	0.9891 (0.5923)	1.0047 (0.1156)	1.9186 (0.2250)	1.8331 (0.2224)	1.7238 (0.2101)	1.6151 (0.1999)	1.5231 (0.1935)
Mildly stable ($\theta_{RN} = 2$)	0.8339 (0.1503)	1.3737 (0.5272)	1.0547 (0.1094)	1.9126 (0.2081)	1.8124 (0.1928)	1.6974 (0.1882)	1.5849 (0.1915)	1.4729 (0.1876)
Mildly unstable ($\theta_{RN} = 2.4$)	0.7677 (0.1246)	1.4587 (0.4646)	1.0974 (0.1132)	1.9062 (0.2113)	1.7887 (0.1932)	1.6723 (0.1841)	1.5512 (0.1855)	1.4282 (0.1883)
Very unstable ($\theta_{RN} = 4$)	0.6444 (0.0930)	1.8959 (0.2967)	1.3416 (0.1131)	2.0073 (0.2946)	1.8728 (0.3246)	1.6933 (0.2873)	1.4787 (0.2853)	1.1595 (0.2722)
<i>Spectral solver estimates</i>								
Very stable ($\theta_{RN} = 1$)	1.0416 (0.2377)	1.1071 (0.7264)	0.9973 (0.1204)	1.9043 (0.2328)	1.8192 (0.2337)	1.7113 (0.2205)	1.6064 (0.2075)	1.5148 (0.2014)
Mildly Stable ($\theta_{RN} = 2$)	1.0308 (0.2922)	2.0941 (1.0461)	0.9869 (0.1383)	1.9209 (0.2300)	1.8212 (0.2220)	1.7240 (0.2189)	1.6244 (0.2232)	1.5258 (0.2269)
Mildly Unstable ($\theta_{RN} = 2.4$)	1.0135 (0.2725)	2.4331 (1.0539)	0.9953 (0.1510)	1.9228 (0.2418)	1.8199 (0.2320)	1.7190 (0.2291)	1.6217 (0.2338)	1.5231 (0.2432)
Very unstable ($\theta_{RN} = 4$)	0.9813 (0.1592)	3.8541 (0.7964)	1.0320 (0.1796)	1.9315 (0.3031)	1.8235 (0.2878)	1.7213 (0.2763)	1.6089 (0.2754)	1.4898 (0.2987)
<i>Conditional estimator</i>								
Very stable ($\theta_{RN} = 1$)	1.0161 (0.2050)	1.0269 (0.6224)	1.0047 (0.1150)	1.9148 (0.2252)	1.8295 (0.2231)	1.7214 (0.2104)	1.6139 (0.1997)	1.5212 (0.1935)
Mildly stable ($\theta_{RN} = 2$)	0.8575 (0.1614)	1.4615 (0.5667)	1.0470 (0.1105)	1.9133 (0.2165)	1.8112 (0.1995)	1.6974 (0.1950)	1.5867 (0.1985)	1.4753 (0.1944)
Mildly unstable ($\theta_{RN} = 2.4$)	0.7935 (0.1364)	1.5634 (0.5106)	1.0880 (0.1148)	1.9051 (0.2132)	1.7896 (0.1946)	1.6743 (0.1856)	1.5550 (0.1879)	1.4332 (0.1892)
Very unstable ($\theta_{RN} = 4$)	0.6577 (0.0831)	2.0094 (0.3030)	1.3414 (0.1052)	2.0072 (0.2638)	1.8501 (0.2829)	1.6777 (0.2564)	1.4751 (0.2502)	1.1581 (0.2308)

Notes: Averages and standard errors (in brackets) computed over 500 Monte Carlo samples. The $K = 100$ NPL algorithm is deemed having converged if $\max\{|\hat{\theta}_{100} - \hat{\theta}_{99}|\} < 10^{-5}$. The conditional estimator is defined as the $K = 100$ NPL algorithm estimates provided that $\hat{\rho}_{NPL} < 1$ with $\hat{\mathbf{P}}_{NPL}$ computed using the spectral solver.

Table 5: Simulation results – $M = 5000$

	$\theta_{RS} = 1$	θ_{RN}	$\theta_{EC} = 1$	$\theta_{FC,1} = 1.9$	$\theta_{FC,2} = 1.8$	$\theta_{FC,3} = 1.7$	$\theta_{FC,4} = 1.6$	$\theta_{FC,5} = 1.5$
<i>All $K = 100$ NPL algorithm estimates</i>								
Very stable ($\theta_{RN} = 1$)	1.0032 (0.0659)	1.0086 (0.2052)	1.0007 (0.0355)	1.9010 (0.0661)	1.8019 (0.0664)	1.7016 (0.0642)	1.6012 (0.0600)	1.5004 (0.0582)
Mildly stable ($\theta_{RN} = 2$)	0.9963 (0.0806)	1.9875 (0.2944)	1.0038 (0.0387)	1.9005 (0.0661)	1.7990 (0.0664)	1.7010 (0.0622)	1.5988 (0.0625)	1.4994 (0.0647)
Mildly unstable ($\theta_{RN} = 2.4$)	0.9696 (0.0508)	2.2761 (0.1927)	1.0123 (0.0353)	1.8991 (0.0699)	1.7949 (0.0685)	1.6953 (0.0661)	1.5947 (0.0666)	1.4954 (0.0681)
Very unstable ($\theta_{RN} = 4$)	0.7667 (0.0240)	2.6684 (0.0603)	1.2213 (0.0338)	1.8942 (0.0820)	1.7736 (0.0791)	1.6647 (0.0760)	1.4770 (0.0744)	1.2901 (0.0678)
<i>All $K = 100$ spectral algorithm estimates</i>								
Very stable ($\theta_{RN} = 1$)	1.0032 (0.0659)	1.0086 (0.2052)	1.0007 (0.0355)	1.9010 (0.0661)	1.8019 (0.0664)	1.7016 (0.0642)	1.6012 (0.0600)	1.5004 (0.0582)
Mildly stable ($\theta_{RN} = 2$)	0.9990 (0.0853)	1.9970 (0.3117)	1.0029 (0.0399)	1.9010 (0.0661)	1.7996 (0.0665)	1.7016 (0.0624)	1.5993 (0.0628)	1.4997 (0.0650)
Mildly unstable ($\theta_{RN} = 2.4$)	1.0035 (0.0886)	2.4121 (0.3486)	0.9991 (0.0465)	1.9049 (0.0708)	1.8027 (0.0707)	1.7042 (0.0689)	1.6037 (0.0706)	1.5034 (0.0736)
Very unstable ($\theta_{RN} = 4$)	0.9993 (0.0432)	3.9922 (0.2132)	1.0025 (0.0497)	1.9046 (0.0913)	1.8019 (0.0882)	1.7036 (0.0852)	1.6020 (0.0838)	1.5008 (0.0880)
<i>Converged $K = 100$ NPL algorithm estimates</i>								
Very stable ($\theta_{RN} = 1$)	1.0030 (0.0659)	1.0079 (0.2053)	1.0008 (0.0355)	1.9011 (0.0662)	1.8020 (0.0665)	1.7017 (0.0643)	1.6013 (0.0601)	1.5004 (0.0583)
Mildly stable ($\theta_{RN} = 2$)	0.9665 (0.0658)	1.8753 (0.2374)	1.0143 (0.0350)	1.9037 (0.0634)	1.7999 (0.0647)	1.7004 (0.0590)	1.5958 (0.0602)	1.4931 (0.0609)
Mildly unstable ($\theta_{RN} = 2.4$)	0.9069 (0.0408)	2.0172 (0.1508)	1.0402 (0.0333)	1.9066 (0.0650)	1.7977 (0.0629)	1.6918 (0.0641)	1.5840 (0.0621)	1.4715 (0.0636)
Very unstable ($\theta_{RN} = 4$)	– (–)	– (–)	– (–)	– (–)	– (–)	– (–)	– (–)	– (–)
<i>Spectral solver estimates</i>								
Very stable ($\theta_{RN} = 1$)	1.0031 (0.0659)	1.0082 (0.2052)	1.0008 (0.0355)	1.9011 (0.0661)	1.8019 (0.0664)	1.7017 (0.0643)	1.6013 (0.0600)	1.5004 (0.0583)
Mildly stable ($\theta_{RN} = 2$)	0.9990 (0.0854)	1.9972 (0.3122)	1.0029 (0.0400)	1.9011 (0.0661)	1.7998 (0.0665)	1.7017 (0.0625)	1.5994 (0.0629)	1.4998 (0.0651)
Mildly unstable ($\theta_{RN} = 2.4$)	1.0035 (0.0886)	2.4121 (0.3486)	0.9991 (0.0465)	1.9049 (0.0708)	1.8027 (0.0701)	1.7042 (0.0689)	1.6037 (0.0706)	1.5034 (0.0736)
Very unstable ($\theta_{RN} = 4$)	0.9993 (0.0431)	3.9918 (0.2131)	1.0025 (0.0498)	1.9046 (0.0913)	1.8018 (0.0882)	1.7035 (0.0851)	1.6018 (0.0836)	1.5006 (0.0877)
<i>Conditional estimator</i>								
Very stable ($\theta_{RN} = 1$)	1.0030 (0.0659)	1.0081 (0.2054)	1.0008 (0.0355)	1.9009 (0.0662)	1.8019 (0.0665)	1.7015 (0.0642)	1.6012 (0.0600)	1.5003 (0.0582)
Mildly stable ($\theta_{RN} = 2$)	0.9995 (0.0850)	1.9992 (0.3109)	1.0026 (0.0397)	1.9009 (0.0662)	1.7995 (0.0665)	1.7016 (0.0625)	1.5994 (0.0627)	1.4996 (0.0650)
Mildly unstable ($\theta_{RN} = 2.4$)	0.9308 (0.0467)	2.1205 (0.1810)	1.0292 (0.0341)	1.9027 (0.0670)	1.7950 (0.0649)	1.6911 (0.0637)	1.5851 (0.0632)	1.4751 (0.0635)
Very unstable ($\theta_{RN} = 4$)	– (–)	– (–)	– (–)	– (–)	– (–)	– (–)	– (–)	– (–)

Notes: Averages and standard errors (in brackets) computed over 500 Monte Carlo samples. The $K = 100$ NPL algorithm is deemed having converged if $\max\{|\hat{\theta}_{100} - \hat{\theta}_{99}|\} < 10^{-5}$. The conditional estimator is defined as the $K = 100$ NPL algorithm estimates provided that $\hat{\rho}_{NPL} < 1$ with $\hat{\mathbf{P}}_{NPL}$ computed using the spectral solver.

the converged NPL algorithm estimates in the “very stable” data generating process. To the best of our knowledge, this case had not been previously considered in the literature. The inclusion of this data generating process in our simulation study serves two purposes. First, it provides a setting where estimates obtained from the converged NPL algorithm do have nice properties. Second, by comparing the results from the “very stable” and the “mildly stable” cases, it is obvious that the properties of the converged NPL algorithm estimator vary within the category of stable data generating processes, with more stable equilibria being more suitable for this estimator. While this insight is quite intuitive, the current simulation exercise seems to be the first one to make this point.

Interestingly, except in the “very stable” case, one typically improves upon the properties of the NPL algorithm estimator by considering any sequence of estimates, regardless whether it converged, instead of focusing on converged ones. We are not aware of existing work in the literature that has made this comparison. This feature of the NPL algorithm is related to the convergence selection exposed in Proposition 3.

For most cases, the spectral approach estimator has nice properties as one expects from the frequency at which this approach delivers the NPL estimator. In particular, it has much better properties than the NPL algorithm estimator when the equilibrium generating the data is not stable. This observation is coherent with the fact that the NPL estimator should be consistent even if stability of the sample NPL mapping fails at the NPL estimates.

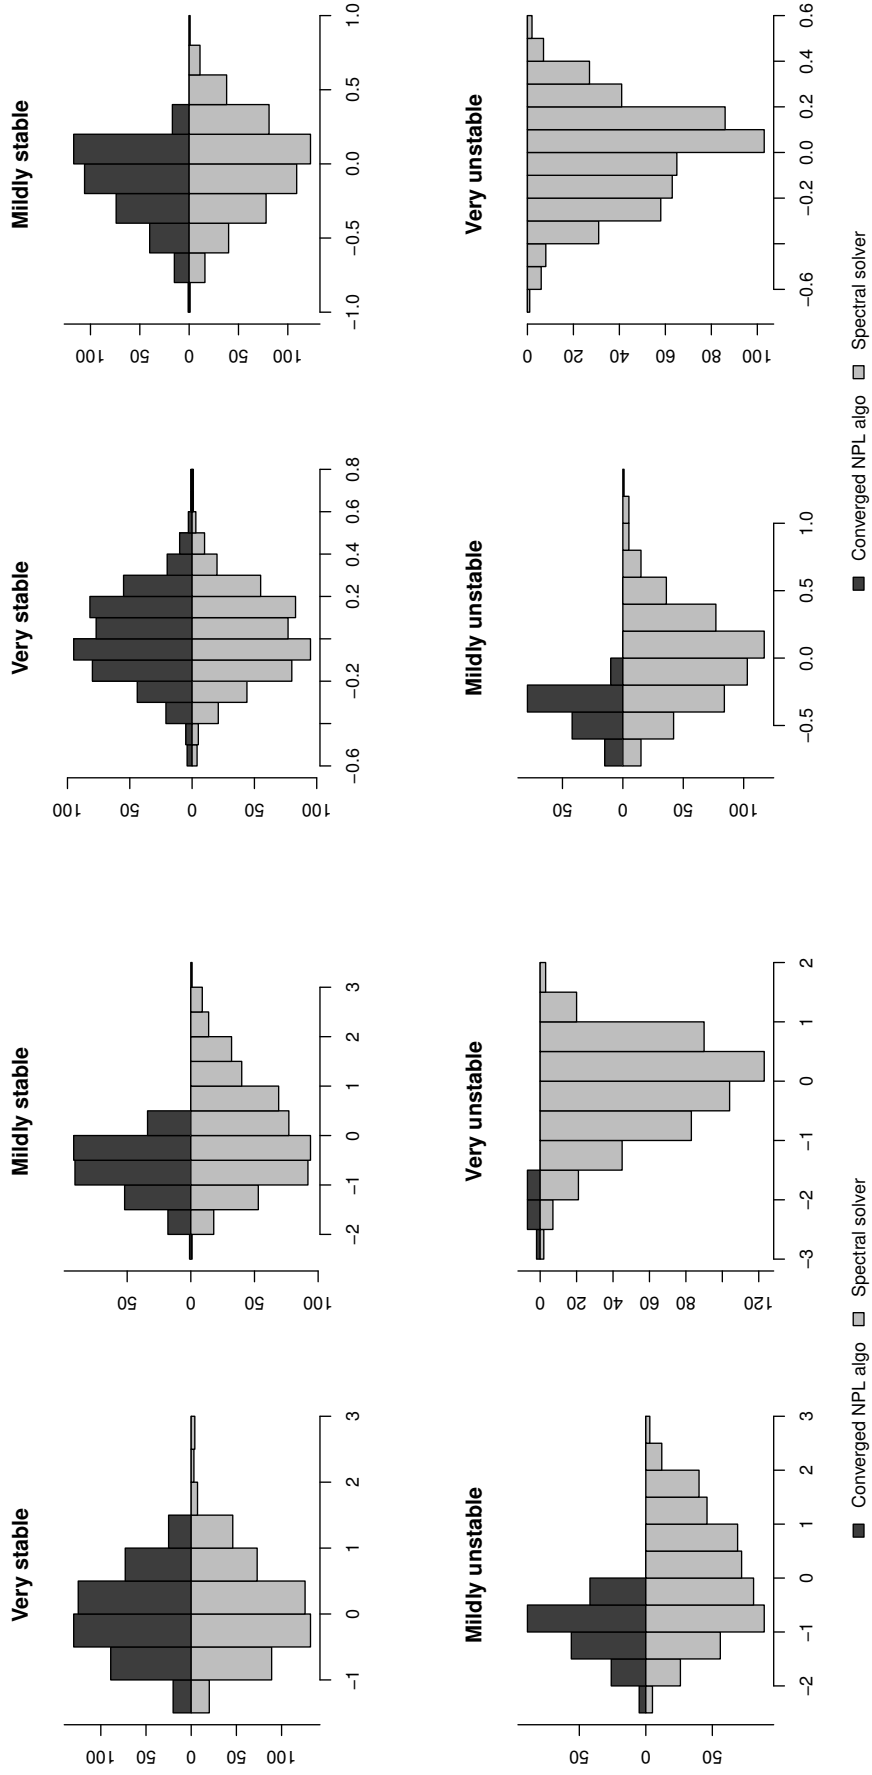
The convergence selection bias from Proposition 3 is also clearly apparent in the results from the conditional estimator. Table 3 has already showed that fewer sequences of the NPL algorithm actually converge as data generating processes become less stable. Tables 4 and 5 further suggest that the few selected samples used to construct the conditional estimator in unstable cases lead to biased estimates. In other words, samples generating converging sequences for the NPL algorithm in unstable data generating processes can be quite different from the samples in which the NPL algorithm fails to converge.

6.2.3 Distribution of estimates

In order to appreciate the properties of the $K = 100$ NPL algorithm estimator and the NPL estimator obtained via the spectral solver, we now compare the full distributions of estimates obtained from the 500 Monte Carlo samples. Since the strategic interaction parameter θ_{RN} is the most problematic one according to the results presented in Tables 4 and 5, we focus our attention on this parameter.¹³

We first consider Figure 3, which compares the histograms of $\hat{\theta}_{RN} - \theta_{RN}^0$ obtained from the converged $K = 100$ NPL algorithm estimates with the spectral solver’s estimates. Sizeable differences appear between the distributions of the two estimators as the data generating process becomes less stable. In fact, while the distributions are practically identical for the

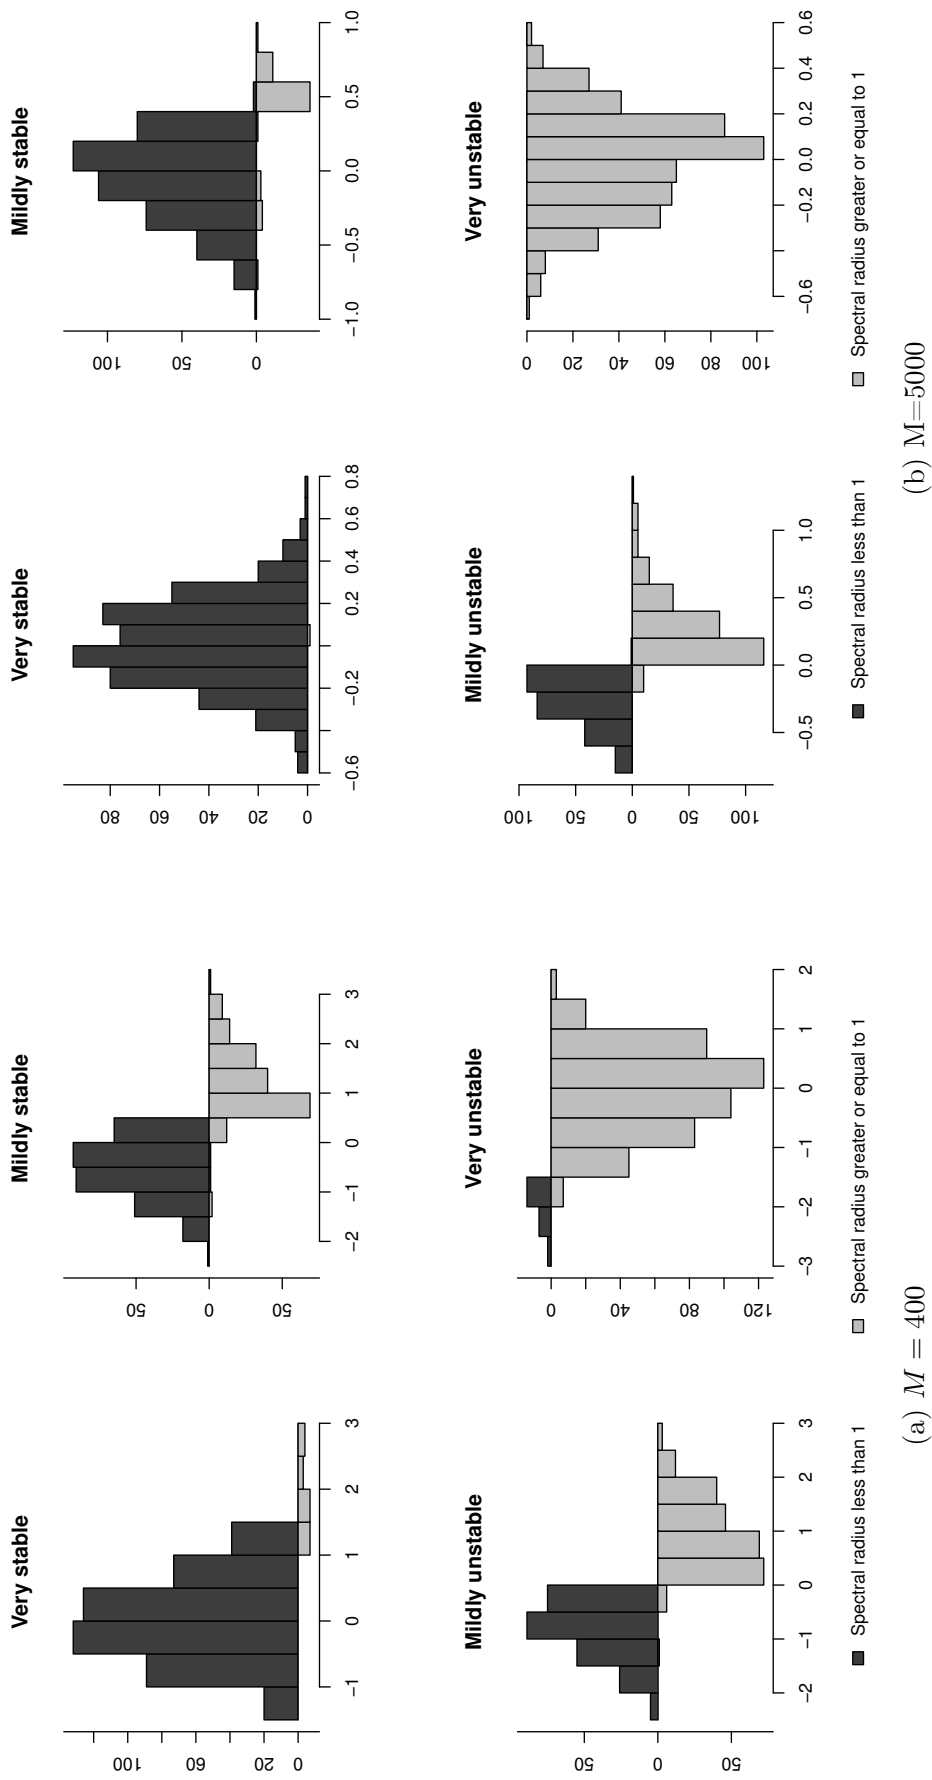
¹³Similar distributions for all other parameters are available from the authors upon request.



(a) $M = 400$

(b) $M = 5000$

Figure 3: Histograms of $\hat{\theta}_{RN} - \theta_{RN}^0$ (Converged NPL algorithm vs spectral solver)



Notes:

Figure 4: Histograms of $\hat{\theta}_{RN} - \theta_{RN}^0$ from spectral solver (by value of $\hat{\rho}_{NPL}$)

“very stable” case, the distribution of the converged NPL algorithm estimates does not even exist in the “very unstable” case with $M = 5000$ since none of the sequences converged. Furthermore, sample size does affect the level of similitude between distributions. While the distributions are still quite similar in the “mildly stable” data generating process with $M = 5000$, they are already quite different with $M = 400$.

The distribution of the spectral approach estimator is roughly normal (perhaps a bit skewed in smaller samples) and centered close to the true value of the parameter. This observation is coherent with the properties of the NPL estimator repeated in Proposition 1. However, the distributions of the converged NPL algorithm estimates tend to be shifted away from the true value in more unstable cases and look like truncated normals. Such a truncation is consistent with Proposition 3. The direction of the truncation is also as expected. Converging sequences of the NPL algorithm tend to be associated with smaller values of the estimated strategic interaction parameter. A similar finding was noted by Mogensen (2015). This pattern is not surprising given the spectral radii reported in Figure 1. It also offers an obvious explanation for the direction of the bias associated with the NPL algorithm estimator restricted to converging sequences and the conditional estimator: since converging sequences typically lead to smaller values of estimated strategic interaction parameters, the convergence selection implies an attenuation bias.

Histograms in Figure 4 split the distributions of spectral solver estimates according to the values of $\hat{\rho}_{\text{NPL}}$. It emphasizes that the NPL estimator exists regardless the stability of the sample NPL fixed point defining the NPL estimator. As expected from Figure 3, $\hat{\rho}_{\text{NPL}} \geq 1$ becomes more frequent as the data generating process is less stable. Once again, there are noticeable truncations in the distributions conditional on $\hat{\rho}_{\text{NPL}} < 1$, as predicted by Proposition 3.

7 Concluding remarks

There are multiple reasons why structural econometricians want to impose the equilibrium restrictions of their models when estimating structural parameters: efficiency, reduction of finite sample bias, and using estimates of choice probabilities – and of other endogeneous objects – that are perfectly comparable with those obtained in counterfactual experiments. Imposing these equilibrium restrictions requires researchers to go beyond simple two-step estimators.

In the estimation of dynamic discrete games, the existing algorithms that impose the model equilibrium restrictions – NFXP, NPL, MPEC, and their variations – have different merits and limitations. As a quasi-Newton algorithm, MPEC guarantees local convergence to the consistent estimator, but it requires the repeated computation of high-dimensional Jacobian matrices at each iteration of the algorithm. In contrast, the NPL algorithm does

not involve the computation of these Jacobian matrices. However, since the NPL algorithm uses fixed point iterations, it requires some stability conditions to converge. Simulation evidence has shown that convergence may fail even if the stability conditions are satisfied at the population level. One solution to by-pass these convergence issues would be to propose a Newton-Kantorovich version of the NPL algorithm. Such an algorithm would satisfy local convergence, but it would involve – as the MPEC algorithm – the repeated computation of high-dimensional Jacobian matrices.

This paper has two main contributions. First, we investigate the reasons and the implications of the lack of convergence of the NPL algorithm. We show that the randomness in the sample NPL mapping – used to update the CCPs – is the culprit. While typically ignored in existing asymptotic studies of the properties of the NPL algorithm, the role of this sample randomness is made obvious when considering that the NPL algorithm is, in practice, iterated in fixed samples. Randomness of the NPL mapping is more likely to hinder convergence if the data generating process is stable, but close to being unstable, especially in samples of modest sizes.

We show that any data generating process – any population – always has a strictly positive probability to generate samples where the algorithm does not converge, as well as samples where the algorithm does converge. We show that the samples for which the algorithm converges are different in a particular and substantial way. This difference generates a convergence selection bias in this type of estimation algorithms. We characterize the nature of this bias and show in Monte Carlo experiments that it can be substantial. This selection bias is more serious when the data generating process is either “mildly stable” or “mildly unstable”. Furthermore, it has a particularly important impact on the structural parameters that capture strategic interactions or competition effects: the estimates of these parameters are biased towards zero. Intuitively, the algorithm converges only for samples consistent with weaker strategic interactions, and this introduces an attenuation bias in the estimation of these strategic parameters.

Second, we propose and implement a spectral algorithm to obtain the NPL estimator. This algorithm has advantages relative to the NPL and the MPEC algorithms. Compared to NPL, the spectral algorithm satisfies local convergence to the consistent estimator. And in contrast to the MPEC algorithm, it does not require the calculation of high-dimensional Jacobian matrices. Furthermore, the dimensionality of the spectral approach’s optimization problem is smaller than the MPEC method – which searches over the joint spaces of CCPs, Lagrange multipliers and structural parameters.

Appendix A Proofs

A.1 Proof of Lemma 1

By definition of the NPL mapping:

$$\sqrt{M} \left(\hat{\phi}(\mathbf{P}) - \phi^0(\mathbf{P}) \right) = -\sqrt{M} \left(\Psi \left(\hat{\theta}(\mathbf{P}), \mathbf{P} \right) - \Psi \left(\theta^0(\mathbf{P}), \mathbf{P} \right) \right) \quad (23)$$

and is therefore, given \mathbf{P} , a continuous nonlinear transformation of a maximum likelihood estimator. A standard delta method argument gives the multivariate normal distribution of $\mathbf{Z}(\mathbf{P})$ stated in Lemma 1.

A.2 Proof of Proposition 1

Consider a first-order expansion of $\hat{\phi}(\hat{\mathbf{P}}_{\text{NPL}})$ around \mathbf{P}^0 :

$$\hat{\phi}(\hat{\mathbf{P}}_{\text{NPL}}) = \hat{\phi}(\mathbf{P}^0) + \nabla \hat{\phi}_{(\mathbf{P}^0)} \left(\hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0 \right) + O_p \left(\left\| \hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0 \right\|^2 \right). \quad (24)$$

Suppose that $\nabla \phi_{(\mathbf{P}^0)}^0$ is invertible. Rearranging (24) and noting that $\hat{\phi}(\hat{\mathbf{P}}_{\text{NPL}}) = \phi^0(\mathbf{P}^0) = \mathbf{0}$ as well as $\left\{ \nabla \hat{\phi}_{(\mathbf{P}^0)} - \nabla \phi_{(\mathbf{P}^0)}^0 \right\} \left(\hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0 \right) \leq O_p \left(\left\| \hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0 \right\|^2 \right)$ give:

$$\hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0 = - \left[\nabla \phi_{(\mathbf{P}^0)}^0 \right]^{-1} \left\{ \hat{\phi}(\mathbf{P}^0) - \phi^0(\mathbf{P}^0) \right\} + O_p \left(\left\| \hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0 \right\|^2 \right) \quad (25)$$

as stated in Proposition 1 (after multiplying by \sqrt{M}).

Consider the following Taylor expansion of $\nabla_{\theta} \hat{Q}_{(\hat{\theta}_{\text{NPL}}, \hat{\mathbf{P}}_{\text{NPL}})} = \mathbf{0}$ around θ^0 and \mathbf{P}^0 :

$$\begin{aligned} \mathbf{0} = & \nabla_{\theta} \hat{Q}_{(\theta^0, \mathbf{P}^0)} + \nabla_{\theta\theta}^2 \hat{Q}_{(\theta^0, \mathbf{P}^0)} \left(\hat{\theta}_{\text{NPL}} - \theta^0 \right) \\ & + \nabla_{\theta\mathbf{P}}^2 \hat{Q}_{(\theta^0, \mathbf{P}^0)} \left(\hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0 \right) + O_p \left(\left\| \hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0 \right\|^2 \right). \end{aligned} \quad (26)$$

Notice that $\nabla_{\theta\theta}^2 \hat{Q}_{(\theta^0, \mathbf{P}^0)} - \nabla_{\theta\theta}^2 Q_{(\theta^0, \mathbf{P}^0)}^0 = O_p(M^{-1/2})$ and that $\nabla_{\theta\mathbf{P}}^2 \hat{Q}_{(\theta^0, \mathbf{P}^0)} - \nabla_{\theta\mathbf{P}}^2 Q_{(\theta^0, \mathbf{P}^0)}^0 = O_p(M^{-1/2})$. Furthermore, one can write:

$$\nabla_{\theta} \hat{Q}_{(\theta^0, \mathbf{P}^0)} = \sum_{\mathbf{x}, \mathbf{y}} \frac{\mathbb{P}(\mathbf{y}|\mathbf{x})}{P^0(\mathbf{y}|\mathbf{x})} \frac{\partial \Psi(\mathbf{y}|\mathbf{x}, \theta, \mathbf{P})}{\partial \theta} \Big|_{(\theta^0, \mathbf{P}^0)} = \nabla_{\theta} \Psi'_{(\theta^0, \mathbf{P}^0)} \mathbf{D}^0 \mathbf{P} \quad (27)$$

$$\nabla_{\theta} Q_{(\theta^0, \mathbf{P}^0)}^0 = \sum_{\mathbf{x}, \mathbf{y}} \frac{P^0(\mathbf{y}|\mathbf{x})}{P^0(\mathbf{y}|\mathbf{x})} \frac{\partial \Psi(\mathbf{y}|\mathbf{x}, \theta, \mathbf{P})}{\partial \theta} \Big|_{(\theta^0, \mathbf{P}^0)} = \nabla_{\theta} \Psi'_{(\theta^0, \mathbf{P}^0)} \mathbf{D}^0 \mathbf{P}^0 = \mathbf{0}. \quad (28)$$

Therefore, (26) becomes:

$$\begin{aligned} \mathbf{0} &= \nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}'_{(\boldsymbol{\theta}^0, \mathbf{P}^0)} \mathbf{D}^0 (\mathbb{P} - \mathbf{P}^0) + \nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 Q^0_{(\boldsymbol{\theta}^0, \mathbf{P}^0)} (\hat{\boldsymbol{\theta}}_{\text{NPL}} - \boldsymbol{\theta}^0) + \nabla_{\boldsymbol{\theta}\mathbf{P}}^2 Q^0_{(\boldsymbol{\theta}^0, \mathbf{P}^0)} (\hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0) \\ &\quad + O_p \left(\left\| \hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0 \right\|^2 \right). \end{aligned} \quad (29)$$

Solving for $\hat{\boldsymbol{\theta}}_{\text{NPL}} - \boldsymbol{\theta}^0$, provided that $\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 Q^0_{(\boldsymbol{\theta}^0, \mathbf{P}^0)}$ is invertible, one gets:

$$\begin{aligned} \hat{\boldsymbol{\theta}}_{\text{NPL}} - \boldsymbol{\theta}^0 &= - \left[\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 Q^0_{(\boldsymbol{\theta}^0, \mathbf{P}^0)} \right]^{-1} \left\{ \nabla_{\boldsymbol{\theta}\mathbf{P}}^2 Q^0_{(\boldsymbol{\theta}^0, \mathbf{P}^0)} (\hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0) + \nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}'_{(\boldsymbol{\theta}^0, \mathbf{P}^0)} \mathbf{D}^0 (\mathbb{P} - \mathbf{P}^0) \right\} \\ &\quad + O_p \left(\left\| \hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0 \right\|^2 \right) \end{aligned} \quad (30)$$

as stated in Lemma 1 (after multiplying by \sqrt{M}).

A.3 Proof of Lemma 2

We first show that:

$$\hat{\boldsymbol{\theta}}_k - \hat{\boldsymbol{\theta}}_* = - \left[\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 \hat{Q}_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} \right]^{-1} \nabla_{\boldsymbol{\theta}\mathbf{P}}^2 \hat{Q}_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} (\hat{\mathbf{P}}_{k-1} - \hat{\mathbf{P}}_*) + O \left(\left\| \hat{\mathbf{P}}_{k-1} - \hat{\mathbf{P}}_* \right\|^2 \right) \quad (31)$$

and

$$\hat{\mathbf{P}}_k - \hat{\mathbf{P}}_* = \nabla \hat{\boldsymbol{\varphi}}_{(\hat{\mathbf{P}}_*)} (\hat{\mathbf{P}}_{k-1} - \hat{\mathbf{P}}_*) + O \left(\left\| \hat{\mathbf{P}}_{k-1} - \hat{\mathbf{P}}_* \right\|^2 \right). \quad (32)$$

The k -th step of the NPL algorithm satisfies:

$$\nabla_{\boldsymbol{\theta}} \hat{Q}_{(\hat{\boldsymbol{\theta}}_k, \hat{\mathbf{P}}_{k-1})} = \mathbf{0} \quad (33)$$

and

$$\hat{\mathbf{P}}_k = \boldsymbol{\Psi} \left(\hat{\boldsymbol{\theta}}_k, \hat{\mathbf{P}}_{k-1} \right). \quad (34)$$

Consider a Taylor expansion of (33) around $(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)$. Noting that $\nabla_{\boldsymbol{\theta}} \hat{Q}_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} = \mathbf{0}$:

$$\mathbf{0} = \nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 \hat{Q}_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} (\hat{\boldsymbol{\theta}}_k - \hat{\boldsymbol{\theta}}_*) + \nabla_{\boldsymbol{\theta}\mathbf{P}}^2 \hat{Q}_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} (\hat{\mathbf{P}}_{k-1} - \hat{\mathbf{P}}_*) + O \left(\left\| \hat{\mathbf{P}}_{k-1} - \hat{\mathbf{P}}_* \right\|^2 \right). \quad (35)$$

Solving for $\hat{\boldsymbol{\theta}}_k - \hat{\boldsymbol{\theta}}_*$ and assuming that $\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 \hat{Q}_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)}$ is invertible, one gets (31).

Consider now a Taylor expansion of (34) around $(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)$. Noting that $\hat{\mathbf{P}}_* = \Psi(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)$:

$$\hat{\mathbf{P}}_k - \hat{\mathbf{P}}_* = \nabla_{\boldsymbol{\theta}} \Psi_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} (\hat{\boldsymbol{\theta}}_k - \hat{\boldsymbol{\theta}}_*) + \nabla_{\mathbf{P}} \Psi_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} (\hat{\mathbf{P}}_{k-1} - \hat{\mathbf{P}}_*) + O\left(\|\hat{\mathbf{P}}_{k-1} - \hat{\mathbf{P}}_*\|^2\right). \quad (36)$$

By plugging (31) in (36):

$$\begin{aligned} \hat{\mathbf{P}}_k - \hat{\mathbf{P}}_* &= \left\{ \nabla_{\mathbf{P}} \Psi_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} - \nabla_{\boldsymbol{\theta}} \Psi_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} \left[\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 \hat{Q}_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} \right]^{-1} \nabla_{\boldsymbol{\theta}\mathbf{P}}^2 \hat{Q}_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} \right\} (\hat{\mathbf{P}}_{k-1} - \hat{\mathbf{P}}_*) \\ &\quad + O\left(\|\hat{\mathbf{P}}_{k-1} - \hat{\mathbf{P}}_*\|^2\right). \end{aligned} \quad (37)$$

Showing that $\nabla \hat{\boldsymbol{\varphi}}_{(\hat{\mathbf{P}}_*)} = \nabla_{\mathbf{P}} \Psi_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} - \nabla_{\boldsymbol{\theta}} \Psi_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} \left[\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 \hat{Q}_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)} \right]^{-1} \nabla_{\boldsymbol{\theta}\mathbf{P}}^2 \hat{Q}_{(\hat{\boldsymbol{\theta}}_*, \hat{\mathbf{P}}_*)}$ would deliver the desired expression in (32). Note that:

$$\nabla \hat{\boldsymbol{\varphi}}_{(\mathbf{P})} = \nabla_{\mathbf{P}} \Psi_{(\hat{\boldsymbol{\theta}}(\mathbf{P}), \mathbf{P})} + \nabla_{\boldsymbol{\theta}} \Psi_{(\hat{\boldsymbol{\theta}}(\mathbf{P}), \mathbf{P})} \nabla \hat{\boldsymbol{\theta}}_{(\mathbf{P})} \quad (38)$$

where $\nabla \hat{\boldsymbol{\theta}}_{(\mathbf{P})}$, which is the Jacobian of $\hat{\boldsymbol{\theta}}(\mathbf{P})$ evaluated at \mathbf{P} , is implicitly defined by the Jacobian $\nabla_{\boldsymbol{\theta}} \hat{Q}_{(\hat{\boldsymbol{\theta}}(\mathbf{P}), \mathbf{P})} = \mathbf{0}$. Taking the derivative with respect to \mathbf{P} , one gets:

$$\nabla_{\boldsymbol{\theta}\mathbf{P}}^2 \hat{Q}_{(\hat{\boldsymbol{\theta}}(\mathbf{P}), \mathbf{P})} + \nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 \hat{Q}_{(\hat{\boldsymbol{\theta}}(\mathbf{P}), \mathbf{P})} \nabla \hat{\boldsymbol{\theta}}_{(\mathbf{P})} = \mathbf{0} \quad (39)$$

which, after rearranging gives $\nabla \hat{\boldsymbol{\theta}}_{(\mathbf{P})} = - \left[\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 \hat{Q}_{(\hat{\boldsymbol{\theta}}(\mathbf{P}), \mathbf{P})} \right]^{-1} \nabla_{\boldsymbol{\theta}\mathbf{P}}^2 \hat{Q}_{(\hat{\boldsymbol{\theta}}(\mathbf{P}), \mathbf{P})}$. Substituting in (38) and evaluating at the NPL fixed point gives the expression of $\nabla \hat{\boldsymbol{\varphi}}_{(\hat{\mathbf{P}}_*)}$. Given the expression in (32), a recursive substitution gives:

$$\hat{\mathbf{P}}_K - \hat{\mathbf{P}}_* = \left[\nabla \hat{\boldsymbol{\varphi}}_{(\hat{\mathbf{P}}_*)} \right]^K (\hat{\mathbf{P}}_0 - \hat{\mathbf{P}}_*) + O_p \left(\max_{k \in \{0, 1, 2, \dots, K-1\}} \left[\nabla \hat{\boldsymbol{\varphi}}_{(\hat{\mathbf{P}}_*)} \right]^k \|\hat{\mathbf{P}}_{K-1-k} - \hat{\mathbf{P}}_*\|^2 \right) \quad (40)$$

Suppose that $\rho(\nabla \hat{\boldsymbol{\varphi}}_{(\hat{\mathbf{P}}_*)}) < 1$. Then, $\lim_{K \rightarrow \infty} \|\hat{\mathbf{P}}_K - \hat{\mathbf{P}}_*\| = 0$ provided that $\hat{\mathbf{P}}_0 \in \mathcal{N}(\hat{\mathbf{P}}_*)$.

A.4 Proof of Lemma 3

Some notation must be introduced. Let $\nabla_{\mathbf{P}} \rho_{(\mathbf{P}^0)}^0$ be the derivative of $\rho(\nabla \boldsymbol{\varphi}_{(\mathbf{P})}^0)$ with respect to \mathbf{P} evaluated at \mathbf{P}^0 . Furthermore, let $\nabla \rho_{(\mathbf{P}^0)}^0$ be the derivative of $\rho(\nabla \boldsymbol{\varphi}_{(\mathbf{P})}^0)$ with respect to the elements of $\nabla \boldsymbol{\varphi}_{(\mathbf{P})}^0$ evaluated at \mathbf{P}^0 . Since $\rho(\nabla \boldsymbol{\varphi}_{(\mathbf{P})}^0)$ is a composition, we can write $\nabla_{\mathbf{P}} \rho_{(\mathbf{P}^0)}^0 = \nabla \rho_{(\mathbf{P}^0)}^0 \circ \nabla^2 \boldsymbol{\varphi}_{(\mathbf{P}^0)}^0$ where $\nabla^2 \boldsymbol{\varphi}_{(\mathbf{P}^0)}^0$ is the matrix of the second-order derivatives of $\boldsymbol{\varphi}^0(\mathbf{P})$ with respect to \mathbf{P} , evaluated at \mathbf{P}^0 . For a given \mathbf{P} and some arbitrary function $\boldsymbol{\theta}(\mathbf{P})$, we can write the derivative of $\Psi(\boldsymbol{\theta}(\mathbf{P}), \mathbf{P})$ with respect to \mathbf{P} evaluated at $(\boldsymbol{\theta}(\mathbf{P}), \mathbf{P})$ as:

$$\nabla \Psi_{(\boldsymbol{\theta}(\mathbf{P}), \mathbf{P})} = \nabla_{\mathbf{P}} \Psi_{(\boldsymbol{\theta}(\mathbf{P}), \mathbf{P})} + \nabla_{\boldsymbol{\theta}} \Psi_{(\boldsymbol{\theta}(\mathbf{P}), \mathbf{P})} \nabla \boldsymbol{\theta}(\mathbf{P}). \quad (41)$$

Using this notation, $\nabla \hat{\boldsymbol{\varphi}}(\mathbf{P}) = \nabla \Psi_{(\hat{\boldsymbol{\theta}}(\mathbf{P}), \mathbf{P})}$ and $\nabla \boldsymbol{\varphi}^0(\mathbf{P}) = \nabla \Psi_{(\boldsymbol{\theta}^0(\mathbf{P}), \mathbf{P})}$. Finally, let $\nabla_j^2 \Psi_{(\boldsymbol{\theta}^0(\mathbf{P}^0), \mathbf{P}^0)}$ be the derivative of $\nabla \Psi_{(\boldsymbol{\theta}(\mathbf{P}), \mathbf{P})}$ with respect to the j -th element of the J -dimensional vector $\boldsymbol{\theta}$ evaluated at $(\boldsymbol{\theta}^0(\mathbf{P}^0), \mathbf{P}^0)$.

The random variable $\sqrt{M}(\hat{\rho}_{\text{NPL}} - \rho^0)$ is equivalent to:

$$\sqrt{M}(\hat{\rho}_{\text{NPL}} - \rho^0) = \sqrt{M}(\hat{\rho}_{\text{NPL}} - \rho(\nabla \hat{\boldsymbol{\varphi}}(\mathbf{P}^0))) + \sqrt{M}(\rho(\nabla \hat{\boldsymbol{\varphi}}(\mathbf{P}^0)) - \rho^0). \quad (42)$$

First, consider $\sqrt{M}(\hat{\rho}_{\text{NPL}} - \rho(\nabla \hat{\boldsymbol{\varphi}}(\mathbf{P}^0)))$. A standard delta method argument implies that:

$$\sqrt{M}(\hat{\rho}_{\text{NPL}} - \rho(\nabla \hat{\boldsymbol{\varphi}}(\mathbf{P}^0))) = \nabla_{\mathbf{P}} \rho_{(\mathbf{P}^0)}^0{}' \sqrt{M}(\hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0) + O_p\left(\sqrt{M} \left\| \hat{\mathbf{P}}_{\text{NPL}} - \mathbf{P}^0 \right\|^2\right). \quad (43)$$

Using Proposition 1, we have that:

$$\sqrt{M}(\hat{\rho}_{\text{NPL}} - \rho(\nabla \hat{\boldsymbol{\varphi}}(\mathbf{P}^0))) \xrightarrow{d} -\nabla_{\mathbf{P}} \rho_{(\mathbf{P}^0)}^0{}' [\nabla \phi_{(\mathbf{P}^0)}^0]^{-1} \mathbf{Z}(\mathbf{P}^0). \quad (44)$$

Second, consider $\sqrt{M}(\rho(\nabla \hat{\boldsymbol{\varphi}}(\mathbf{P}^0)) - \rho^0)$. By a standard delta method argument:

$$\begin{aligned} \sqrt{M}(\rho(\nabla \hat{\boldsymbol{\varphi}}(\mathbf{P}^0)) - \rho^0) &= \text{vec} \left[\nabla \rho_{(\mathbf{P}^0)}^0 \right]' \text{vec} \left[\sqrt{M}(\nabla \hat{\boldsymbol{\varphi}}(\mathbf{P}^0) - \nabla \boldsymbol{\varphi}_{(\mathbf{P}^0)}^0) \right] \\ &+ O_p\left(\sqrt{M} \left\| \nabla \hat{\boldsymbol{\varphi}}(\mathbf{P}^0) - \nabla \boldsymbol{\varphi}_{(\mathbf{P}^0)}^0 \right\|^2\right). \end{aligned} \quad (45)$$

Assuming that $\nabla \Psi_{(\boldsymbol{\theta}(\mathbf{P}), \mathbf{P})}$ is continuous in $\boldsymbol{\theta}(\mathbf{P})$, once again using the delta method leads to:

$$\begin{aligned} \sqrt{M}(\nabla \hat{\boldsymbol{\varphi}}(\mathbf{P}) - \nabla \boldsymbol{\varphi}_{(\mathbf{P})}^0) &= \sum_{j=1}^J \nabla_j^2 \Psi_{(\boldsymbol{\theta}^0(\mathbf{P}), \mathbf{P})} \sqrt{M}(\hat{\boldsymbol{\theta}}_j(\mathbf{P}) - \boldsymbol{\theta}_j^0(\mathbf{P})) \\ &+ O_p\left(\sqrt{M} \left\| \hat{\boldsymbol{\theta}}(\mathbf{P}) - \boldsymbol{\theta}^0(\mathbf{P}) \right\|^2\right). \end{aligned} \quad (46)$$

Then, considering $\mathbf{P} = \mathbf{P}^0$:

$$\sqrt{M}(\nabla \hat{\boldsymbol{\varphi}}(\mathbf{P}^0) - \nabla \boldsymbol{\varphi}_{(\mathbf{P}^0)}^0) \xrightarrow{d} \sum_{j=1}^J \nabla_j^2 \Psi_{(\boldsymbol{\theta}^0(\mathbf{P}^0), \mathbf{P}^0)} L_j(\mathbf{P}^0). \quad (47)$$

Let $\mathbf{G}^0 \equiv -\nabla_{\mathbf{P}} \rho_{(\mathbf{P}^0)}^0{}' [\nabla \phi_{(\mathbf{P}^0)}^0]^{-1}$ and:

$$\mathbf{J}^0 \equiv \text{vec} \left[\nabla \rho_{(\mathbf{P}^0)}^0 \right]' \left[\text{vec} \left[\nabla_1^2 \Psi_{(\boldsymbol{\theta}^0(\mathbf{P}^0), \mathbf{P}^0)} \right], \dots, \text{vec} \left[\nabla_J^2 \Psi_{(\boldsymbol{\theta}^0(\mathbf{P}^0), \mathbf{P}^0)} \right] \right]. \quad (48)$$

By combining equations (42), (44) and (47), we obtain:

$$\sqrt{M} (\hat{\rho}_{\text{NPL}} - \rho^0) \xrightarrow{d} \mathbf{G}^0 \mathbf{Z}(\mathbf{P}^0) + \mathbf{J}^0 \mathbf{L}(\mathbf{P}^0) \quad (49)$$

as required. Since $\mathbf{L}(\mathbf{P}^0)$ and $\mathbf{Z}(\mathbf{P}^0)$ are normally distributed and centered at vectors of 0's, we conclude that $\sqrt{M} (\hat{\rho}_{\text{NPL}} - \rho^0)$ follows a normal distribution centered at 0.

A.5 Proof of Proposition 2

Under the conditions of Proposition 2, we have that for $K > \bar{K}$:

$$\sqrt{M} (\hat{\boldsymbol{\theta}}_K - \boldsymbol{\theta}^0) = \sqrt{M} (\hat{\boldsymbol{\theta}}_{\text{NPL}} - \boldsymbol{\theta}^0) + \mathbb{1}_{\{\hat{\rho}_{\text{NPL}} \geq 1\}} \sqrt{M} (\hat{\boldsymbol{\theta}}_* - \hat{\boldsymbol{\theta}}_{\text{NPL}}). \quad (50)$$

Therefore, as $M \rightarrow \infty$:

$$\sqrt{M} (\hat{\boldsymbol{\theta}}_K - \boldsymbol{\theta}^0) \xrightarrow{d} \begin{cases} \sqrt{M} (\hat{\boldsymbol{\theta}}_{\text{NPL}} - \boldsymbol{\theta}^0) & \text{if } \hat{\rho}_{\text{NPL}} < 1 \\ \sqrt{M} (\hat{\boldsymbol{\theta}}_{\text{NPL}} - \boldsymbol{\theta}^0) + \mathbf{W}_* & \text{if } \hat{\rho}_{\text{NPL}} \geq 1 \end{cases} \quad (51)$$

Finally, the result follows by applying Lemma 3.

A.6 Proof of Proposition 3

Under the conditions stated in Proposition 3, we have that any well-defined NPL algorithm estimator satisfies $\hat{\boldsymbol{\theta}}_K = \tilde{\boldsymbol{\theta}}_{\text{NPL}}$ and $\sqrt{M} (\hat{\boldsymbol{\theta}}_K - \boldsymbol{\theta}^0) = \sqrt{M} (\tilde{\boldsymbol{\theta}}_{\text{NPL}} - \boldsymbol{\theta}^0)$. By definition:

$$\sqrt{M} (\tilde{\boldsymbol{\theta}}_{\text{NPL}} - \boldsymbol{\theta}^0) = \begin{cases} \sqrt{M} (\hat{\boldsymbol{\theta}}_{\text{NPL}} - \boldsymbol{\theta}^0) & \text{if } \sqrt{M} (\hat{\rho}_{\text{NPL}} - \rho^0) < \sqrt{M} (1 - \rho^0) \\ \text{is not well defined} & \text{if } \sqrt{M} (\hat{\rho}_{\text{NPL}} - \rho^0) \geq \sqrt{M} (1 - \rho^0) \end{cases} \quad (52)$$

which corresponds to a truncated distribution. As $M \rightarrow \infty$, the asymptotic distributions of $\sqrt{M} (\hat{\boldsymbol{\theta}}_{\text{NPL}} - \boldsymbol{\theta}^0)$ and $\sqrt{M} (\hat{\rho}_{\text{NPL}} - \rho^0)$ are given in Proposition 1 and in Lemma 3.

References

- AGUIRREGABIRIA, V., AND C. ALONSO-BORREGO (2014): “Labor contracts and flexibility: evidence from a labor market reform in Spain,” *Economic Inquiry*, 52(2), 930–957.
- AGUIRREGABIRIA, V., AND C. HO (2012): “A dynamic oligopoly game of the US airline industry: Estimation and policy experiments,” *Journal of Econometrics*, 168(1), 156–173.
- AGUIRREGABIRIA, V., AND P. MIRA (2002): “Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models,” *Econometrica*, 70(4), 1519–1543.

- (2007): “Sequential estimation of dynamic discrete games,” *Econometrica*, 75(1), 1–53.
- AKSIN, Z., B. ATA, S. EMADI, AND C. SU (2013): “Structural estimation of callers’ delay sensitivity in call centers,” *Management Science*, 59(12), 2727–2746.
- BAJARI, P., C. BENKARD, AND J. LEVIN (2007): “Estimating dynamic model of imperfect competition,” *Econometrica*, 75(1), 1331–1370.
- BARWICK, P., AND P. PATHAK (2015): “The costs of free entry: An empirical study of real estate agents in Greater Boston,” *RAND Journal of Economics*, 46(1), 103–145.
- BARZILAI, J., AND J. BORWEIN (1988): “Two-point step size gradient methods,” *IMA Journal of Numerical Analysis*, 8(1), 141–148.
- BUGNI, F., AND J. BUNTING (2018): “On the iterated estimation of dynamic discrete choice games,” *Working paper*.
- CHEN, J., S. ESTEBAN, AND M. SHUM (2013): “When do secondary markets harm firms?,” *American Economic Review*, 103(7), 2911–2934.
- COLLARD-WEXLER, A. (2013): “Demand fluctuations in the ready-mix concrete industry,” *Econometrica*, 81(3), 1003–1037.
- COPELAND, A., AND C. MONNET (2009): “The welfare effects of incentive schemes,” *The Review of Economic Studies*, 76(1), 93–113.
- DE PINTO, A., AND G. NELSON (2009): “Land use change with spatially explicit data: A dynamic approach,” *Environmental and Resource Economics*, 43(2), 209–229.
- DEARING, A., AND J. BLEVINS (2019): “Efficient and convergent sequential pseudo-likelihood estimation of dynamic discrete games,” *Working paper*.
- EGESDAL, M., Z. LAI, AND C. SU (2015): “Estimating dynamic discrete-choice games of incomplete information,” *Quantitative Economics*, 6(3), 567–597.
- ELICKSON, P., AND S. MISRA (2008): “Supermarket pricing strategies,” *Marketing Science*, 27(5), 811–828.
- FRECHETTE, G., A. LIZZERI, AND T. SALZ (2019): “Frictions in a competitive, regulated market: Evidence from taxis,” *American Economic Review*, 109(8), 2954–2992.
- GAYLE, P., AND X. XIE (2018): “Entry deterrence and strategic alliances,” *Economic Inquiry*, 56(3), 1898–1924.
- HAN, L., AND S. HONG (2011): “Testing cost inefficiency under free entry in the real estate brokerage industry,” *Journal of Business and Economic Statistics*, 29(4), 564–578.
- HOTZ, V., AND R. MILLER (1993): “Conditional choice probabilities and the estimation of dynamic models,” *The Review of Economic Studies*, 60(3), 497–529.
- HUANG, L., AND M. SMITH (2014): “The dynamic efficiency costs of common-pool resource exploitation,” *American Economic Review*, 104(12), 4071–4103.

- ISKHAKOV, F., J. LEE, J. RUST, B. SCHJERNING, AND K. SEO (2016): “Comment on “Constrained optimization approaches to estimation of structural models”,” *Econometrica*, 84(1), 365–370.
- KANO, K. (2013): “Menu costs and dynamic duopoly,” *International Journal of Industrial Organization*, 31(1), 102–118.
- KASAHARA, H., AND K. SHIMOTSU (2012): “Sequential estimation of structural models with a fixed point constraint,” *Econometrica*, 80(5), 2303–2319.
- KOSOROK, M. (2008): *Introduction to empirical processes and semiparametric inference*. Springer.
- LA CRUZ, W., J. MARTINEZ, AND M. RAYDAN (2006): “Spectral residual method without gradient information for solving large-scale nonlinear systems of equations,” *Mathematics of Computation*, 75(255), 1429–1448.
- LIN, H. (2015): “Quality choice and market structure: A dynamic analysis of nursing home oligopolies,” *International Economic Review*, 56(4), 1261–1290.
- LIN, Z., AND H. XU (2017): “Estimation of social-influence-dependent peer pressure in a large network game,” *The Econometrics Journal*, 20(3), S86–S102.
- LIU, X., AND J. ZHOU (2017): “A social interaction model with ordered choices,” *Economics Letters*, 161(1), 86–89.
- MOGENSEN, P. (2015): “Monte Carlo studies and the convergence of NPL for games,” *Working paper*.
- 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,” *The Review of Economic Studies*, 75(3), 901–928.
- (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(5), 999–1033.
- SU, C. (2014): “Estimating discrete-choice games of incomplete information: Simple static examples,” *Quantitative Marketing and Economics*, 12(2), 167–207.
- SU, C., AND K. JUDD (2012): “Constrained optimization approaches to estimation of structural models,” *Econometrica*, 80(5), 2213–2230.
- SWEETING, A. (2013): “Dynamic product positioning in differentiated product industries: The effect of fees for musical performance rights on the commercial radio industry,” *Econometrica*, 81(5), 1763–1803.

- TOMLIN, B. (2014): “Exchange rate fluctuations, plant turnover and productivity,” *International Journal of Industrial Organization*, 35(1), 12–28.
- VARADHAN, R., AND P. GILBERT (2009): “BB: An R package for solving a large system of nonlinear equations and for optimizing a high-dimensional nonlinear objective function,” *Journal of Statistical Software*, 32(4), 1–26.
- VARADHAN, R., AND C. ROLAND (2008): “Simple and globally convergent methods for accelerating the convergence of any EM algorithm,” *Scandinavian Journal of Statistics*, 35(2), 335–353.
- VITORINO, M. (2012): “Empirical entry games with complementarities: An application to the shopping center industry,” *Journal of Marketing Research*, 49(2), 175–191.