Solving Discrete Choice Dynamic Programming Models Using Euler Equations*

Victor Aguirregabiria †
University of Toronto, CEPR

Arvind Magesan ‡
University of Calgary

May 2, 2023

Abstract

This paper shows that the *Finite Dependence* (FD) and *Euler Equations* (EE) representations of discrete choice Markov Decision Processes can be used to construct a computationally efficient algorithm for the solution of this class of dynamic programming (DP) problems. We show that the FD/EE representations imply a fixed point contraction mapping with a Lipschitz constant strictly smaller than the discount factor of the DP problem. This implies substantial computational gains relative to standard solution methods such as value function, relative value function, and policy function iterations. We use numerical experiments to illustrate the computational gains of this method.

Keywords: Dynamic programming discrete choice models; Euler equations; Relative Value Function iteration; Estimation; Approximation bias.

JEL codes: C23, C25, C51.

*We would like to thank comments from Jason Blevins, Rob Bray, Thierry Magnac, Angelo Melino, Bob Miller, Pedro Mira, Jean-Marc Robin, John Rust, Bertel Schjerning, Kunio Tsuyuhara, and from seminar participants at Calgary, North Carolina, Ohio State, Toronto, Western University, the Barcelona GSE Summer Forum (Structural Microeconometrics), the Canadian Economic Association conference, the Microeconometric Network Meeting in Copenhagen, the Society of Economic Dynamics conference and The Banff Empirical Microeconomics Conference 2015.

†Department of Economics, University of Toronto. 150 St. George Street, Toronto, ON, M5S 3G7, Canada, victor.aguirregabiria@utoronto.ca.

‡Department of Economics, University of Calgary. 2500 University Drive, N.W Calgary, AB, Canada, arvindmagesan@gmail.com.

1 Introduction

The Euler equation-GMM approach of Hansen and Singleton (1982) in continuous dynamic decision models, and the *finite dependence* representation of Hotz and Miller (1993) and Arcidiacono and Miller (2011) in dynamic discrete choice models are fundamental contributions to the estimation of dynamic structural models. A main advantage of these estimation methods is that they avoid the curse of dimensionality associated with the computation of present values.[1] In the same spirit, Aguirregabiria and Magesan (2013) provide a representation of dynamic discrete choice models as a continuous decision problem where the decision variables are choice probabilities and derive marginal conditions of optimality similar to the Euler equations in continuous decision problems and to the finite dependence representation in discrete choice models. While these methods significantly reduce the computational burden associated with estimating structural parameters, the end goal of structural work is typically to use an estimated model to study the effect of policies that have never occurred. These counterfactual experiments require the solution of the dynamic programming (DP) problem. However, as far as we know, there are no results in the literature showing computational advantages of the *Finite Dependence / Euler Equations* (FD/EE) representation for solving DP problems.

In this paper, we show that the FD/EE representation of dynamic discrete choice models can be used to construct a fixed point mapping in the space of conditional choice value functions that we denote the *Euler Equation (EE) operator*. This operator is a contraction such that successive iterations in this mapping deliver the unique solution to the DP problem. For an important class of models, so-called *Multi-armed bandit models*, the contraction or *Lipschitz* constant of this mapping is strictly smaller than the discount factor of the DP problem, which significantly speeds up the convergence of fixed point iterations.[2]

We compare the computational properties of the EE iterations algorithm with those of the most commonly used solution methods: Value Function (VF) iterations, Relative Value Function (RVF) iterations, and Policy Function (PF) iterations. The EE operator is a stronger contraction (i.e., it has a smaller Lipschitz constant) than VF and RVF, and has similar contraction properties as the PF operator. In terms of the time required to evaluate the operator a single time, the EE operator is also more efficient than VF, RVF, and PF algorithms, and this difference increases quadratically with the dimension of the state space. We present numerical examples that illustrate how solving the model by using EE iterations implies substantial computational

[1] The computational cost of estimating structural parameters using these methods increases with sample size but not with the dimension of the state space.

[2] The defining feature of multi-armed bandit models is that the only endogenous state variable is the decision variable at the previous period. Many empirically relevant problems have this feature, including market entry/exit, demand with switching costs, occupational choice, migration, and store location, among others.

savings relative to these three alternative methods.

We define a sample-based version of the EE operator and use it to estimate counterfactual *conditional choice probabilities* (CCP).[3] This sample-based EE operator is defined only at sample points of the exogenous state variables, and thus its dimensionality does not increase with the dimension of the state space. We show that this sample-based EE operator is also a contraction with a Lipschitz constant strictly smaller than the discount factor, and its unique fixed point is a root-N consistent estimator of the true solution. The computational cost to obtain counterfactual CCPs using this method does not depend on the dimension of the state space of the exogenous state variables.

We illustrate the computational gains of the EE method using several numerical experiments in the context of a dynamic model of market entry and exit with a high dimension in the exogenous state variables. In the first experiment, we compare the computation time of EE, VF, RVF, and PF algorithms for the exact solution of the model. For moderately sized state spaces (up to 200,000 states), the EE algorithm is over 200 times faster than PF, between 50-70 times faster than VF, and between 5-40 times faster than RVF.[4] These differences increase with the dimension of the state space. Many models that are computationally infeasible for all practical purposes using standard methods are feasible using the EE method. We also use this first set of experiments to study the source of the difference in total computation time across these algorithms. We show that the advantages of the EE algorithm relative to VF and RVF are due to improvements in both time-per-iteration and the number of iterations to convergence. In a second experiment, we use the entry/exit model to study the performance of the sample version of the EE operator relative to more standard methods in predicting the response to a counterfactual increase in the cost of entry, holding the computation time of the different methods fixed. We show that the finite sample properties of the EE estimator are substantially better than those of the standard methods, i.e., mean absolute bias and squared error are between 35 and 60 percentage points smaller in the EE method.

The rest of the paper is organized as follows. Section 2 presents the model and the FD/EE representation. Section 3 contains the main results of this paper. We derive the EE operator and show that it is a contraction with a Lipschitz constant smaller than the discount factor. We also define the sample-based EE operator and explain how it delivers consistent estimates of counterfactual CCPs. Section 4 presents our Monte Carlo experiments. We summarize and conclude in section 5. Proofs of Propositions are in Appendix A.

[3] Our sample-based EE operator is related to the random grid method of Rust (1997), though Rust defines and applies this method to standard value function and policy function operators, and not to Euler equations.

[4] The EE algorithm is also faster than hybrid algorithms that combine PF and VF or RVF iterations.

2 Model

2.1 Framework

We consider a stationary Markov Decision Process (MDP) with a discrete decision variable as in Puterman (1994) or Rust (1994). Each discrete period t, an agent takes a decision a_t to maximize her expected intertemporal payoff, $\mathbb{E}_t\left[\sum_{j=0}^{\infty}\beta^j\ \Pi(a_{t+j},\mathbf{s}_{t+j})\right]$, where $\beta\in(0,1)$ is the discount factor, $\Pi(.)$ is the payoff function, and $\mathbf{s}_t\in\mathcal{S}$ is the vector of state variables that follows a controlled Markov process with a transition probability function $f_s(\mathbf{s}_{t+1}|a_t,\mathbf{s}_t)$. The decision variable a_t belongs to the discrete and finite set $\mathcal{A}=\{0,1,...,J\}$. The value function $V(\mathbf{s}_t)$ can be obtained solving the Bellman equation:

$$V(\mathbf{s}_t)\ =\ \max_{a_t\in\mathcal{A}}\left\{\Pi(a_t,\mathbf{s}_t)+\beta\int V(\mathbf{s}_{t+1})\ f_s(\mathbf{s}_{t+1}|a_t,\mathbf{s}_t)\ d\mathbf{s}_{t+1}\right\} \tag{1}$$

The optimal decision rule, $\alpha(\mathbf{s}_t):\mathcal{S}\to\mathcal{A}$, is obtained as the argument that maximizes the expression in brackets $\{.\}$ in equation (1).

Following the standard model of dynamic discrete choice in the literature (Rust, 1994), we partition the vector of state variables $\mathbf{s}_t$ into two sets of variables, $\mathbf{s}_t=(\mathbf{x}_t,\varepsilon_t)$ where $\mathbf{x}_t$ belongs to a discrete set $\mathcal{X}$ and is observable to the researcher, while $\boldsymbol{\varepsilon}_t\equiv\{\varepsilon_t(a):a\in\mathcal{A}\}$ is a vector of $J+1$ choice-specific state variables that are *unobservable* to the researcher. These state variables satisfy the usual assumptions of *additive separability* and *conditional independence*. Specifically, the payoff function $\Pi(.)$ is given by (i.e., *additive separability*):

$$\Pi(a_t,\mathbf{s}_t)\ =\ \pi(a_t,\mathbf{x}_t)\ +\ \varepsilon_t(a_t) \tag{2}$$

The transition probability function of the state variables factors as (i.e., *conditional independence*) $f_s(\mathbf{s}_{t+1}|a_t,\mathbf{s}_t)\ =\ f\left(\mathbf{x}_{t+1}|a_t,\mathbf{x}_t\right)\ g\left(\boldsymbol{\varepsilon}_{t+1}\right)$, where $g\left(.\right)$ is the density function of $\boldsymbol{\varepsilon}_t$ which is absolutely continuous with respect to the Lebesgue measure, differentiable in all its arguments, and possess finite moments.

Let $V^{\sigma}(\mathbf{x}_t)$ be the *integrated value function*, which is defined as the original value function integrated over the density of the unobservable state variables, i.e., $V^{\sigma}(\mathbf{x}_t)\equiv\int V(\mathbf{x}_t,\boldsymbol{\varepsilon}_t)\ g(\boldsymbol{\varepsilon}_t)\ d\boldsymbol{\varepsilon}_t$. The *integrated Bellman equation* has the following form:

$$V^{\sigma}(\mathbf{x}_t)\ =\ \int\max_{a\in\mathcal{A}}\left\{v(a,\mathbf{x}_t)+\varepsilon_t(a)\right\}\ g\left(\boldsymbol{\varepsilon}_t\right)\ d\boldsymbol{\varepsilon}_t \tag{3}$$

where $v(a,\mathbf{x}_t)$ is the *conditional-choice value function*, or value of choosing alternative a, and is

defined as:

$$v(a, \mathbf{x}_t) \equiv \pi(a, \mathbf{x}_t) + \beta \sum_{\mathbf{x}_{t+1} \in \mathcal{X}} V^{\sigma}(\mathbf{x}_{t+1}) \; f\left(\mathbf{x}_{t+1} | a, \mathbf{x}_t\right) \tag{4}$$

Define the value difference $\widetilde{v}(a, \mathbf{x}_t) \equiv v(a, \mathbf{x}_t) - v(0, \mathbf{x}_t)$, where the selection of alternative 0 as the baseline is without loss of generality. Let $P(a|\mathbf{x}_t) \equiv Pr(\alpha(\mathbf{x}_t, \boldsymbol{\varepsilon}_t) = a|\mathbf{x}_t)$ be the *conditional choice probability* (CCP), i.e., the probability that alternative a is optimal conditional on the observable state $\mathbf{x}_t$. Let $\widetilde{\mathbf{v}}(\mathbf{x}_t) = \{\widetilde{v}(a, \mathbf{x}_t) : a \in \mathcal{A} - \{0\}\}$ be a vector of J value differences, and let $\mathbf{P}(\mathbf{x}_t) \equiv \{P(a|\mathbf{x}_t) : a \in \mathcal{A} - \{0\}\}$ be a vector of choice probabilities. The *optimal choice probability* (OCP) function, $\mathbf{\Lambda}\left(\widetilde{\mathbf{v}}_t(\mathbf{x})\right) \equiv \{\Lambda\left(a, \widetilde{\mathbf{v}}_t(\mathbf{x})\right) : a \in \mathcal{A} - \{0\}\}$, is the following mapping from the space of value differences into the space of choice probabilities:

$$P(a|\mathbf{x}_t) = \Lambda\left(a, \widetilde{\mathbf{v}}(\mathbf{x}_t)\right) \equiv \int 1\left\{\widetilde{v}(a, \mathbf{x}_t) + \varepsilon_t(a) \geq \widetilde{v}(j, \mathbf{x}_t) + \varepsilon_t(j), \;\; \forall j\right\} \;\; g\left(\boldsymbol{\varepsilon}_t\right) d\boldsymbol{\varepsilon}_t \tag{5}$$

where $1\{.\}$ is the indicator function. In vector form, we have $\mathbf{P}(\mathbf{x}_t) = \mathbf{\Lambda}\left(\widetilde{\mathbf{v}}(\mathbf{x}_t)\right)$.

Hotz-Miller Inversion Property (Hotz and Miller, 1993) establishes that the OCP mapping is invertible: there exists an inverse function $\mathbf{\Lambda}^{-1}(.)$ such that $\widetilde{\mathbf{v}}(\mathbf{x}_t) = \mathbf{\Lambda}^{-1}(\mathbf{P}(\mathbf{x}_t))$. The following equation is a Corollary of *Hotz-Miller Inversion Property* (see Arcidiacono and Miller, 2011). For any choice alternative a and state $\mathbf{x}_t$, we have:

$$V^{\sigma}(\mathbf{x}_t) \;\; = \;\; v(a, \mathbf{x}_t) \;\; + \;\; \psi(a, \mathbf{P}(\mathbf{x}_t)) \tag{6}$$

where $\psi(.)$ is a primitive function that only depends on the probability distribution of $\boldsymbol{\varepsilon}_t$. The combination of the integrated Bellman equation in (3) and Hotz-Miller representation in (6) gives us the definition of function $\psi(.)$:[5]

$$\psi(a, \mathbf{P}(\mathbf{x}_t)) \;\; \equiv \;\; \int \max_{j \in \mathcal{A}} \left\{\Lambda^{-1}(j, \mathbf{P}(\mathbf{x}_t)) - \Lambda^{-1}(a, \mathbf{P}(\mathbf{x}_t)) + \varepsilon_t(j)\right\} \; g\left(\boldsymbol{\varepsilon}_t\right) \;\; d\boldsymbol{\varepsilon}_t \tag{7}$$

2.2 Finite dependence / Euler equation representation

The Hotz-Miller inversion property has allowed researchers to develop different representations of optimal behavior that significantly reduce the computational burden associated with the estimation of dynamic discrete choice models. One such representation is the *finite dependence* (FD) representation of Arcidiacono and Miller (2011), who show that if the transition probability of the state variables $\mathbf{x}_t$ satisfies a particular condition, then value differences can be written

[5] For instance, when $\varepsilon_t(a)$ is i.i.d. type I extreme value, we have that $\psi(a, \mathbf{P}_t(\mathbf{x}_t)) = \gamma - \ln P_t(a|\mathbf{x}_t)$, where γ is Euler's constant.

in terms of payoffs and choice probabilities at a finite number of periods. Here we focus on *2-periods finite dependence* as it fully characterizes the solution of multi-armed bandit problems that we study in this paper.

Define the two-periods forward transition probability of the state variables:

$$f^{(2)}\left(\mathbf{x}_{t+2} \mid \mathbf{x}_t, a_t, a_{t+1}\right) = \sum_{\mathbf{x}_{t+1}} f\left(\mathbf{x}_{t+2} \mid \mathbf{x}_{t+1}, a_{t+1}\right) f\left(\mathbf{x}_{t+1} \mid \mathbf{x}_t, a_t\right) \tag{8}$$

DEFINITION. *A controlled transition probability $f\left(\mathbf{x}_{t+1}|\mathbf{x}_t, a_t\right)$ has the 2-period finite dependence property if there is at least a state $\mathbf{x}_t$ and choice paths (j_t, j_{t+1}) and (k_t, k_{t+1}) such that:*

$$f^{(2)}\left(\mathbf{x}_{t+2} \mid \mathbf{x}_t, j_t, j_{t+1}\right) = f^{(2)}\left(\mathbf{x}_{t+2} \mid \mathbf{x}_t, k_t, k_{t+1}\right), \text{ for any } \mathbf{x}_{t+2} \in \mathcal{X} \qquad \blacksquare \tag{9}$$

Following Arcidiacono and Miller (2011), we now use this property to obtain a representation of value differences in terms of payoffs and choice probabilities at periods t and $t+1$. Plugging Hotz-Miller equation (6) into the expression for the conditional choice value function in (4):

$$v(a_t, \mathbf{x}_t) = \pi(a_t, \mathbf{x}_t) + \beta \sum_{\mathbf{x}_{t+1}} \Big(v(a_{t+1}, \mathbf{x}_{t+1}) + \psi(a_{t+1}, \mathbf{P}(\mathbf{x}_{t+1})) \Big) f\left(\mathbf{x}_{t+1}|a_t, \mathbf{x}_t\right) \tag{10}$$

Replacing $v(a_{t+1}, \mathbf{x}_{t+1})$ with its definition from equation (4):

$$\begin{aligned} v(a, \mathbf{x}_t) &= \pi(a, \mathbf{x}_t) + \beta \sum_{\mathbf{x}_{t+1}} \pi(a_{t+1}, \mathbf{x}_{t+1}) f\left(\mathbf{x}_{t+1}|a_t, \mathbf{x}_t\right) \\ &+ \beta \sum_{\mathbf{x}_{t+1}} \psi(a_{t+1}, \mathbf{P}(\mathbf{x}_{t+1}) f\left(\mathbf{x}_{t+1}|a_t, \mathbf{x}_t\right) + \beta^2 \sum_{\mathbf{x}_{t+2}} V(\mathbf{x}_{t+2}) f_2\left(\mathbf{x}_{t+2}|\mathbf{x}_t, a_t, a_{t+1}\right) \end{aligned} \tag{11}$$

Equation (11) holds for any state $\mathbf{x}_t$ and any choice path (a_t, a_{t+1}). Now, consider a state $\mathbf{x}_t$ and two choice paths (j_t, j_{t+1}) and (k_t, k_{t+1}) that satisfy the $2-period$ finite dependence property in equation (9). Combining equation (11) and the $2-period$ finite dependence property, we have the following expression for the value differences:

$$\begin{aligned} v(j_t, \mathbf{x}_t) - v(k_t, \mathbf{x}_t) &= \pi(j_t, \mathbf{x}_t) - \pi(k_t, \mathbf{x}_t) \\ &+ \beta \sum_{\mathbf{x}_{t+1}} \pi(j_{t+1}, \mathbf{x}_{t+1}) f\left(\mathbf{x}_{t+1}|j_t, \mathbf{x}_t\right) - \pi(k_{t+1}, \mathbf{x}_{t+1}) f\left(\mathbf{x}_{t+1}|k_t, \mathbf{x}_t\right) \\ &+ \beta \sum_{\mathbf{x}_{t+1}} \psi(j_{t+1}, \mathbf{P}(\mathbf{x}_{t+1})) f\left(\mathbf{x}_{t+1}|j_t, \mathbf{x}_t\right) - \psi(k_{t+1}, \mathbf{P}(\mathbf{x}_{t+1})) f\left(\mathbf{x}_{t+1}|k_t, \mathbf{x}_t\right) \end{aligned} \tag{12}$$

A key implication of the $2-period$ finite dependence property is that the term associated with the continuation value at period $t+2$, $V(\mathbf{x}_{t+2})$, does not appear in this expression for the value difference. Therefore, equation (12) provides a representation of value differences in terms of payoffs and CCPs at periods t and $t+1$ only. This representation has been used to develop estimation methods – CCP or Hotz-Miller methods – that do not require solving the dynamic programming problem. In this paper, we use this representation to derive a new solution algorithm.

Aguirregabiria and Magesan (2013) establish that discrete choice MDPs within the class of Rust (1994) can be represented as continuous decision models where the decision variable at period t is the vector of CCPs, $\mathbf{P}_t$. They show that the optimal CCPs of the DP problem should also solve a constrained optimization problem that consists of the maximization of expected payoff at periods t and $t+1$, $\mathbb{E}_t\left(\Pi_t + \beta\ \Pi_{t+1}\right)$, under the constraint of keeping fixed the probability distribution of $\mathbf{x}_{t+2}$. The first-order conditions of this constrained optimization problem – i.e., the Lagrange equations – depend on CCPs and payoffs at periods t and $t+1$, and on the Lagrange multipliers of the constraints. The authors show that, under the $2-period$ finite dependence condition, it is possible to solve for the Lagrange multipliers and obtain optimality conditions that depend only on CCPs and payoffs at periods t and $t+1$. These optimality conditions are the so-called Euler equations. Interestingly, these Euler equations have the same form as equation (12) on value differences. For the rest of the paper, we use the term *Euler equations* to denote these $2-period$ optimality conditions.

3 Euler fixed point mapping

Since CCPs depend only on differences between conditional choice values, the right-hand side in the Euler equation (12) depends on primitives of the model and value differences only, such that we can interpret this equation as a mapping from value differences into value differences. We show below that this Euler-equation mapping defines the unique solution to the DP problem. For the rest of the paper, we focus on a particular class of discrete choice MDP: *Multi-armed bandit models.*

DEFINITION. *A Multi-armed bandit model is a discrete choice MDP where the only endogenous state variable is the decision variable at the previous period. That is,* $\mathbf{x}_t = (a_{t-1}, \mathbf{z}_t)$*, where* $\mathbf{z}_t$ *is a vector of exogenous state variables with transition density function* $f_z(\mathbf{z}_{t+1}|\mathbf{z}_t)$. ■

The class of Multi-armed bandit models includes many important economic applications such as market entry/exit, demand with switching costs, occupational choice, migration, and store location, among others.

Multi-armed bandit models have the $2-$*period* finite dependence property for any state $\mathbf{x}_t$ and for any pair of choice paths (j_t, j_{t+1}) and (k_t, k_{t+1}) as long as $j_{t+1} = k_{t+1}$. To characterize the solution of the DP problem, we only need one Euler equation for each $\mathbf{x}_t$ and choice alternative a_t. Without loss of generality, for every action $a \in \mathcal{A} - \{0\}$, we consider the following pair of choice paths at t and $t+1$: $(a, 0)$ and $(0, 0)$. This implies the following system of $J|\mathcal{X}|$ Euler equations:

$$\widetilde{v}(a, \mathbf{x}_t) \;=\; c(a, \mathbf{x}_t) \;+\; \beta \sum_{\mathbf{z}_{t+1}} \Big(S(\widetilde{\mathbf{v}}(a, \mathbf{z}_{t+1})) - S(\widetilde{\mathbf{v}}(0, \mathbf{z}_{t+1})) \Big) f_z\left(\mathbf{z}_{t+1} | \mathbf{z}_t\right) \tag{13}$$

where $c(a, \mathbf{x}_t)$ is the following function of primitives of the model:

$$c(a, \mathbf{x}_t) \;\equiv\; \pi(a, \mathbf{x}_t) - \pi(0, \mathbf{x}_t) + \beta \sum_{\mathbf{z}_{t+1}} \Big(\pi(0, a, \mathbf{z}_{t+1}) - \pi(0, 0, \mathbf{z}_{t+1}) \Big) f_z\left(\mathbf{z}_{t+1} | \mathbf{z}_t\right) \tag{14}$$

And $S(\widetilde{\mathbf{v}}(\mathbf{x}_t))$ is *McFadden's Social Surplus function*:

$$S(\widetilde{\mathbf{v}}(\mathbf{x}_t)) \;\equiv\; \int \max_{j \in \mathcal{A}} \left\{ \widetilde{v}(j, \mathbf{x}_t) + \varepsilon_t(j) \right\} g\left(\boldsymbol{\varepsilon}_t\right) \; d\boldsymbol{\varepsilon}_t \tag{15}$$

We use function $\Gamma_{EE}(a, \mathbf{x}_t, \widetilde{\mathbf{v}})$ to represent the right-hand side of the Euler equation (13) and we denote it the *Euler mapping*. We write this system of equations in vector form to obtain this fixed point representation:

$$\widetilde{\mathbf{v}} \;=\; \mathbf{\Gamma}_{EE}(\widetilde{\mathbf{v}}) \tag{16}$$

where $\widetilde{\mathbf{v}}$ is the vector of value differences for any action in $\mathcal{A} - \{0\}$ and any state in $\mathcal{X}$, and similarly, $\mathbf{\Gamma}_{EE}(\widetilde{\mathbf{v}})$ is the vector of functions $\Gamma_{EE}(a, \mathbf{x}_t, \widetilde{\mathbf{v}})$ for any action and state.

Proposition 1 is the main result of this paper. It establishes that the Euler mapping is a contraction with a Lipschitz constant strictly smaller than the discount factor β. A property of the Euler mapping that plays a key role in the proof of this Proposition is that $\mathbf{\Gamma}_{EE}(a_t, a_{t-1}, \mathbf{z}; \widetilde{\mathbf{v}}) - \mathbf{\Gamma}_{EE}(a_t, 0, \mathbf{z}; \widetilde{\mathbf{v}}) = \Delta(a_t, a_{t-1}, \mathbf{z})$, where $\Delta(a_t, a_{t-1}, \mathbf{z})$ is a primitive of the model that is known to the researcher before solving the model and does not depend on the vector of values $\widetilde{\mathbf{v}}$ where we are evaluating the mapping. This property implies that all the vectors in the sequence $\{\widetilde{\mathbf{v}}^k : k \geq 1\}$, obtained by applying fixed-point iterations from some initial $\widetilde{\mathbf{v}}^0$ (i.e., $\widetilde{\mathbf{v}}^{k+1} = \Gamma_{EE}(\widetilde{\mathbf{v}}^k)$), satisfy the property $\widetilde{v}(a_t, a_{t-1}, \mathbf{z}) - \widetilde{v}(a_t, 0, \mathbf{z}) = \Delta(a_t, a_{t-1}, \mathbf{z})$. Therefore, without loss of generality, we can restrict our analysis of the fixed point mapping Γ_{EE} to the subspace of values $\widetilde{\mathbf{v}}$ that satisfy this restriction. That is, we consider $\Gamma_{EE}(\widetilde{\mathbf{v}})$ on space,

$$\mathcal{V}^R \equiv \left\{ \widetilde{\mathbf{v}} \in \mathbb{R}^{J|\mathcal{X}|} : \widetilde{v}(a_t, a_{t-1}, \mathbf{z}) = \widetilde{v}(a_t, 0, \mathbf{z}) + \Delta(a_t, a_{t-1}, \mathbf{z}) \text{ for any } (a_t, a_{t-1}, \mathbf{z}) \right\} \tag{17}$$

PROPOSITION 1. The Euler mapping Γ_{EE} is a contraction in the complete metric space $(\mathcal{V}^R, \|.\|_\infty)$, and its Lipschitz constant δ is strictly smaller than β. That is, there is a constant $\delta \in (0,1)$ with $\delta < \beta$ such that for any pair $\widetilde{\mathbf{v}}$ and $\widetilde{\mathbf{w}}$ in $\mathcal{V}^R$, we have that $\|\mathbf{\Gamma}_{EE}(\widetilde{\mathbf{v}}) - \mathbf{\Gamma}_{EE}(\widetilde{\mathbf{w}})\|_\infty \leq \delta \|\widetilde{\mathbf{v}} - \widetilde{\mathbf{w}}\|_\infty$. ■

Proof: In Appendix A.

A corollary of Proposition 1 is that successive iterations in the Euler operator is a method to solve this discrete choice MDP. Below we compare this method to the most commonly used methods for solving DP problems.

3.1 Solution algorithms

We describe the different algorithms using vector-form notation. Let $|\mathcal{Z}|$ be the dimension of the space of exogenous state variables. All of the following vectors have dimension $|\mathcal{Z}| \times 1$: $\mathbf{V}(y)$ is the vector of values when the endogenous state is $a_{t-1} = y$; and $\boldsymbol{\pi}(a,y)$, $\widetilde{\mathbf{v}}(a,y)$, $\mathbf{P}(a,y)$ are the vectors of one-period payoffs, differential values, and conditional choice probabilities, respectively, when current action is $a_t = a$ and the endogenous state is $a_{t-1} = y$. Finally, $\mathbf{F}_z$ is the $|\mathcal{Z}| \times |\mathcal{Z}|$ matrix of transition probabilities of the exogenous state variables.

3.1.1 Value function (VF) iterations

The value function operator is a fixed point mapping in the space of the vector of values $\mathbf{V} = \{\mathbf{V}(y) : y \in \mathcal{A}\}$. It is defined as $\mathbf{\Gamma}_{VF}(\mathbf{V}) = \{\mathbf{\Gamma}_{VF}(y, \mathbf{V}) : y \in \mathcal{A}\}$ with

$$\mathbf{\Gamma}_{VF}(y, \mathbf{V}) \quad = \quad \int \max_{j \in A} \left\{ \pi(j,y) + \beta \ \mathbf{F}_z \ \mathbf{V}(j) + \varepsilon_t(j) \right\} \ g(\boldsymbol{\varepsilon}_t) \ d\boldsymbol{\varepsilon}_t \tag{18}$$

The algorithm starts with an initial vector of values $\mathbf{V}_0$. At every iteration $n \geq 1$ it updates the vector using $\mathbf{V}_n = \mathbf{\Gamma}_{VF}(\mathbf{V}_{n-1})$. The computational complexity per iteration derives from the number of multiplications involved in matrix products $\mathbf{F}_z \ \mathbf{V}_{n-1}(j)$ for the $J+1$ choice alternatives. In other words, we need to perform the multiplication $f_z(z^{'}|z) * V(j, \mathbf{z}^{'})$ for each triple $(j, z, z^{'}) \in \mathcal{A} \times \mathcal{Z}^2$. The computational complexity is thus of the order $(J+1) \ |\mathcal{Z}|^2$. The degree of contraction of this mapping, as measured by the Lipschitz constant, is equal to the discount factor β (see Puterman, 1994, and Rust, 1996).

3.1.2 Relative Value function (RVF) iterations

Let $\mathbf{x}_0 \in \mathcal{X}$ be an arbitrary value of the vector of state variables. The *relative value function operator* is a fixed point mapping in the space of the vector of values $\mathbf{V} = \{\mathbf{V}(y) : y \in \mathcal{A}\}$, that

is defined as $\mathbf{\Gamma}_{RVF}(\mathbf{V}) = \{\mathbf{\Gamma}_{RVF}(y, \mathbf{V}) : y \in \mathcal{A}\}$ with

$$\mathbf{\Gamma}_{RVF}(y, \mathbf{V}) \;=\; \int \max_{j \in A} \left\{ \pi(j, y) + \beta \; \mathbf{F}_z \; \left[\mathbf{V}(j) - V(\mathbf{x}_0) \; \mathbf{1}_{|\mathcal{Z}|} \right] + \varepsilon_t(j) \right\} \; g(\boldsymbol{\varepsilon}_t) \; d\boldsymbol{\varepsilon}_t \qquad (19)$$

where $\mathbf{1}_{|\mathcal{Z}|}$ is a column vector of ones. Given an initial $\mathbf{V}_0$, this vector is updated at every iteration $n \geq 1$ using $\mathbf{V}_n = \mathbf{\Gamma}_{RVF}(\mathbf{V}_{n-1})$. The computational complexity per iteration in the case of relative value iteration is the same as the case of value function iterations, and derives from the number of multiplications involved in matrix product $\mathbf{F}_z \; \left[\mathbf{V}(j) - V(\mathbf{x}_0) \; \mathbf{1}_{|\mathcal{Z}|} \right]$ In other words, we need to perform the multiplication $f_z(z^{'}|z) * (V(j, z^{'}) - V(\mathbf{x}_0))$ for each triple $(j, z, z^{'}) \in \mathcal{A} \times \mathcal{Z}^2$. The computational complexity is thus of the order $(J+1) \; |\mathcal{Z}|^2$. The difference between the two algorithms is in their contraction properties. The Lipschitz constant of the RVF operator is $\beta \; \rho(\mathbf{F}_z)$, where $\rho(\mathbf{F}_z) \in (0, 1)$ is the spectral radius of the transition matrix $\mathbf{F}_z$ (Morton and Wecker, 1977; Puterman, 1994, section 6.6; Bray, 2019, Proposition 3).

When the stochastic process of the exogenous state variables has strong time persistence, the spectral radius $\rho(\mathbf{F}_z)$ is close to one and the degree of contraction of RVF is similar to VF. However, when the exogenous state variables are not so persistent, the value of $\rho(\mathbf{F}_z)$ is substantially smaller than one and the RVF algorithm converges to a solution faster than VF iterations.

3.1.3 Policy function (PF), or Newton-Kantorovich, iterations

The policy function operator (Puterman and Brumelle, 1979) is a fixed point mapping in the space of the vector of conditional choice probabilities $\mathbf{P} \in \{\mathbf{P}(a, y) : (a, y) \in [\mathcal{A} - \{0\}] \times \mathcal{A}\}$. It is defined as $\mathbf{\Gamma}_{PF}(\mathbf{P}) = \{\mathbf{\Gamma}_{PF}(a, y, \mathbf{P}) : (a, y) \in [\mathcal{A} - \{0\}] \times \mathcal{A}\}$ with,

$$\mathbf{\Gamma}_{PF}(a, y, \mathbf{P}) = \int 1 \left\{ \pi(a, y) + \beta \; \mathbf{F}_z \; \mathbf{W}^{\mathbf{P}}(a) + \varepsilon_t(a) \geq \pi(j, y) + \beta \; \mathbf{F}_z \; \mathbf{W}^{\mathbf{P}}(j) + \varepsilon_t(j), \; \forall j \right\} \; g(\boldsymbol{\varepsilon}_t) \; d\boldsymbol{\varepsilon}_t \qquad (20)$$

where $\mathbf{W}^{\mathbf{P}}(a)$ is an $|\mathcal{Z}| \times 1$ vector that contains the present discounted values of future payoffs conditional on every possible value of $\mathbf{z}_t$ and on $a_{t-1} = a$, and conditional on the behavior of the agent in future periods follows the vector of choice probabilities $\mathbf{P}$. The vectors of present values $\{\mathbf{W}^{\mathbf{P}}(a) : a \in \mathcal{A}\}$ are obtained solving the system of linear equations: for any $y \in \mathcal{A}$,

$$\mathbf{W}^{\mathbf{P}}(y) \;=\; \sum_{a=0}^{J} \mathbf{P}(a, y) * \left[\pi(a, y) + \beta \; \mathbf{F}_z \; \mathbf{W}^{\mathbf{P}}(a) \right], \qquad (21)$$

where $*$ is the element-by-element product. The linear operator described in the system (21), that delivers the vectors $\mathbf{W}^{\mathbf{P}}(a)$ for a given vector of CCPs $\mathbf{P}$, is denoted the *valuation oper-*

ator. The operator in the right-hand side of equation (20) that returns a new vector of CCPs given the vectors of valuations $\mathbf{W}^{\mathbf{P}}$ is denoted the *policy improvement operator*. By definition, policy function operator in (20) is the composition of the valuation and the policy improvement operators.

The PF algorithm starts with an initial $\mathbf{P}_0$, and at every iteration $n \geq 1$ updates the vector using $\mathbf{P}_n = \mathbf{\Gamma}_{PF}(\mathbf{P}_{n-1})$. The computational complexity of one iteration in the PF operator is given by the complexity of solving the system of linear equations (21). This complexity is determined by the number of objects to be solved in this system, and so is of the order $O(|\mathcal{X}|^3)$. Therefore, one PF iteration is more costly than one VF or RVF iteration, and the difference increases with the dimension of the state space. However, the PF operator is a stronger contraction than VF and RVF operators, and therefore it requires a smaller number of iterations to achieve convergence.

3.1.4 Euler function iterations

The EE operator is a fixed point mapping in the restricted space of value differences $\mathcal{V}^R$ as described above and is defined as:

$$\mathbf{\Gamma}_{EE}(a, y, \widetilde{\mathbf{v}}) \;=\; \mathbf{c}(a, y) \;+\; \beta\, \mathbf{F}_z\Big(S(\widetilde{\mathbf{v}}(a)) - S(\widetilde{\mathbf{v}}(0))\Big) \tag{22}$$

where $\mathbf{c}(a, y)$ is the vector of primitives:

$$\mathbf{c}(a, y) \;\equiv\; \boldsymbol{\pi}(a, y) - \boldsymbol{\pi}(0, y) + \beta\, \mathbf{F}_z\Big(\boldsymbol{\pi}(0, a) - \boldsymbol{\pi}(0, 0)\Big) \tag{23}$$

The algorithm starts with an initial vector $\widetilde{\mathbf{v}}_0 \in \mathcal{V}^R$ and at every iteration $n \geq 1$ updates the vector using $\widetilde{\mathbf{v}}_n = \mathbf{\Gamma}_{EE}(\widetilde{\mathbf{v}}_{n-1})$. The computational complexity of each iteration is of the order $J\ |\mathcal{Z}|^2$, and it comes from the calculation of the matrix products $\mathbf{F}_z\Big(S(\widetilde{\mathbf{v}}(a)) - S(\widetilde{\mathbf{v}}(0))\Big)$. In other words, we need to perform the multiplication $f_z(z^{'}|z) * (S(\widetilde{\mathbf{v}}(a, z^{'})) - S(\widetilde{\mathbf{v}}(0, z^{'})))$ for each triple $(a, \mathbf{z}, \mathbf{z}^{'}) \in \{\mathcal{A} - 0\} \times \mathcal{Z}^2$; note that we do not need to do this multiplication for one of the endogenous states (0). The computational complexity is thus of the order $(J)\ |\mathcal{Z}|^2$. This algorithm has two computational advantages with respect to VF and RVF methods. First, its cost per iteration is smaller: it is of the order $J\ |\mathcal{Z}|^2$ for EE, and $(J+1)\ |\mathcal{Z}|^2$ for VF and RVF. This is because the EE operator needs to calculate expectations of next-period values not for every current choice alternative (as VF and RVF) but for all alternatives except one. This difference increases with the dimension of the state space. But the computational advantage of EE compared to VF and RVF is that it is a stronger contraction such that it converges to the solution using a smaller number of iterations. We illustrate these advantages in section 4.

3.2 Sample-based EE operator and counterfactuals

Suppose that the researcher's dataset consists of panel data of N agents, indexed by i, over T periods with information on agents' actions and state variables, $\{a_{it}\ ,\ \mathbf{x}_{it} : i = 1, 2, ..., N; t = 1, 2, ..., T\}$. We consider a sample where the number of agents N is large and the number of periods T is small, i.e., asymptotic results as $N \to \infty$ and T is fixed. The researcher is interested in using this sample to estimate CCPs associated with a *counterfactual experiment*. We assume that the payoff function is known to the researcher, e.g., the structural parameters of the model have been estimated using FD/EE method such as Arcidiacono and Miller (2011) or Aguirregabiria and Magesan (2013). In this section, we present an estimation method for counterfactual CCPs that uses the EE operator.

Given the sample values of the exogenous state variables $\{\mathbf{z}_{it}\}$ and the counterfactual values of the structural parameters, the researcher can construct an empirical counterpart of the EE operator in equation (22). Define the empirical set $\mathcal{Z}_N \equiv \{\mathbf{z} \in \mathcal{Z} : \text{there is an observation } (i,t) \text{ with } \mathbf{z}_{it} = \mathbf{z}\}$, and the empirical transition probability function $\widehat{f}_N(\mathbf{z}'|\mathbf{z}_0)$ defined on $\mathcal{Z}_N \times \mathcal{Z}_N$ into $[0,1]$, such that for any $\mathbf{z}_0 \in \mathcal{Z}_N$, $\widehat{f}_N(\mathbf{z}'|\mathbf{z}_0) = \sum_{i=1}^N 1\{\mathbf{z}_{it+1} = \mathbf{z}' \text{ and } \mathbf{z}_{it} = \mathbf{z}_0\} / \sum_{i=1}^N 1\{\mathbf{z}_{it} = \mathbf{z}_0\}$. Stationarity of the transition probability $f_z(\mathbf{z}_{t+1}|\mathbf{z}_t)$ implies that: (1) the set $\mathcal{Z}_N$ is a random sample from the ergodic set $\mathcal{Z}$; (2) $\mathcal{Z}_N$ converges to $\mathcal{Z}$ as $N \to \infty$; and (3) $\widehat{f}_N(\mathbf{z}'|\mathbf{z}_0)$ converges uniformly to $f_z(\mathbf{z}'|\mathbf{z}_0)$ as $N \to \infty$. Let $\mathbb{E}^{(N)}_{\{\mathbf{z}'|\mathbf{z}\}}[.]$ be a sample-based conditional mean operator from $\mathbb{R}$ into $\mathbb{R}$ such that for any real-valued function $h(\mathbf{z}')$ the operator is defined as:

$$\mathbb{E}^{(N)}_{\{\mathbf{z}'|\mathbf{z}_0\}}\left[h(\mathbf{z}')\right] \equiv \sum_{\mathbf{z}' \in \mathcal{Z}_N} \widehat{f}_N(\mathbf{z}'|\mathbf{z}_0)\ h(\mathbf{z}') \tag{24}$$

The *Sample-based EE operator* $\Gamma^{(N)}_{EE}(\widetilde{\mathbf{v}})$ is defined as the sample counterpart of the EE mapping in Equation (22):

$$\Gamma^{(N)}_{EE}(a, y, \mathbf{z}; \widetilde{\mathbf{v}}) = c^{(N)}(a, y, \mathbf{z}) + \beta\ \mathbb{E}^{(N)}_{\{\mathbf{z}'|\mathbf{z}\}}\Big[S\big(\tilde{v}(a, \mathbf{z}')\big) - S\big(\tilde{v}(0, \mathbf{z}')\big)\Big] \tag{25}$$

This is a fixed point mapping in the space of value differences such that we can obtain a sample-based solution to the DP problem by solving the fixed point problem $\widetilde{\mathbf{v}} = \mathbf{\Gamma}^{(N)}_{EE}(\widetilde{\mathbf{v}})$. Importantly, the dimension of this fixed point mapping is $J(J+1)N$, which can be many orders of magnitude smaller than the dimension of Γ_{EE} when the dimension of $\mathcal{Z}$ is large relative to sample size.

PROPOSITION 2. The Sample-Based EE operator $\mathbf{\Gamma}^{(N)}_{EE}(\widetilde{\mathbf{v}})$ *is a contraction with Lipstchitz constant strictly smaller than* β *and it converges uniformly in probability to the true EE mapping* $\mathbf{\Gamma}_{EE})(\widetilde{\mathbf{v}})$ *as* $N \to \infty$..

Proof: In Appendix A.

Let $\boldsymbol{\theta}$ be the vector of structural parameters characterizing the payoff function π, the density function of the unobservables g, and the transition probability F_z. For notational convenience, we now include explicitly $\boldsymbol{\theta}$ as an argument of the EE operator, i.e., $\boldsymbol{\Gamma}_{EE}\widetilde{\mathbf{v}}, \boldsymbol{\theta})$. Suppose the researcher is interested in estimating the effects on CCPs of a change in the structural parameters from the estimate $\widehat{\boldsymbol{\theta}}$ to an alternative vector $\boldsymbol{\theta}^*$. To estimate the effects of this counterfactual experiment, the researcher needs to solve the DP problem given the structural parameters $\boldsymbol{\theta}^*$. We can represent this solution in terms of the vector of value differences $\widetilde{\mathbf{v}}^*$. The vector $\widetilde{\mathbf{v}}^*$ is defined as the unique solution of the fixed point problem $\widetilde{\mathbf{v}}^* = \boldsymbol{\Gamma}_{EE}(\widetilde{\mathbf{v}}^*, \boldsymbol{\theta}^*)$.

In most empirical applications, the dimension of the state space, and in particular the dimension of $\mathcal{Z}$, is very large such that the exact computation of $\widetilde{\mathbf{v}}^*$ is computationally infeasible. Here we propose an approximation to the solution using the Sample-based EE operator. We approximate $\widetilde{\mathbf{v}}^*$ using $\widetilde{\mathbf{v}}^*_N$, which is defined as the unique fixed point of the Sample-based EE mapping,

$$\widetilde{\mathbf{v}}^*_N = \boldsymbol{\Gamma}^{(N)}_{EE}(\widetilde{\mathbf{v}}^*_N, \boldsymbol{\theta}^*) \equiv \left\{ \Gamma^{(N)}_{EE}(a, y, \mathbf{z}; \widetilde{\mathbf{v}}) : (a, y, \mathbf{z}) \in \mathcal{A} - \{0\} \times \mathcal{A} \times \mathcal{Z}_N \right\} \tag{26}$$

And the corresponding vector of conditional choice probabilities is $\mathbf{P}^*_N = \Lambda(\widetilde{\mathbf{v}}^*_N)$. This approximate solution has the following interesting properties.

(a) Lower computational cost and no curse of dimensionality. The vector $\widetilde{\mathbf{v}}^*_N$ and the mapping $\boldsymbol{\Gamma}^{(N)}_{EE}$ have the same dimension as the sample size. In most empirical applications, this dimension is many orders of magnitude smaller than the dimension of the state space. This reduction in the dimension of the fixed point problem and the stronger contraction property of the EE operator imply substantial computational savings. The Sample-based EE mapping is an Euler equation version of the random operators defined in Rust (1997). Rust shows that these operators succeed in breaking the curse of dimensionality for discrete choice MDP with continuous state variables. This property also applies to our dynamic decision model when the endogenous state variables are discrete and exogenous state variables are continuous.

(b) Consistency. $\widetilde{\mathbf{v}}^*_N$ and $\mathbf{P}^*_N$ are consistent estimators of the true counterfactuals $\widetilde{\mathbf{v}}^*$ and $\mathbf{P}^*$.

PROPOSITION 3. The vector of value differences $\widetilde{\mathbf{v}}^*_N$, which is defined as the fixed point $\widetilde{\mathbf{v}}^*_N = \boldsymbol{\Gamma}^{(N)}_{EE}(\widetilde{\mathbf{v}}^*_N, \boldsymbol{\theta}^*)$, is a root-N consistent and asymptotically normal estimator of $\widetilde{\mathbf{v}}^*$.

Proof: In Appendix A.

4 Numerical experiments

In this section, we present Monte Carlo experiments to illustrate the performance of the *EE iterations algorithm* in terms of computational savings and statistical precision in two problems: the exact solution of the DP problem and the estimation of counterfactual experiments. We evaluate this algorithm in the context of a dynamic model of market entry and exit.

First, we examine the differences in the computational burdens of the four solution algorithms, VF, RVF, PF, and EE. Generally speaking, the total time required to obtain a model solution is comprised of two factors, the amount of time per iteration and the number of iterations. We use the experiments to compare the time per iteration and the number of iterations each method takes to converge. Second, given an estimated model and a counterfactual experiment that consists of an increase in the cost of market entry, we present Monte Carlo experiments to evaluate the finite sample properties of the four algorithms to estimate counterfactual choice probabilities. These four methods consist of finding a fixed point of the sample-based versions of the VF, RVF, PF, and EE operators.

4.1 Design of the experiments

We consider a dynamic model of firm entry and exit decisions in a market. The decision variable a_t is the indicator of being active in a market, such that the action space is $\mathcal{A} = \{0, 1\}$. The endogenous state variable is the lagged value of the decision variable, $y_t = a_{t-1}$, and it represents whether the firm has to pay an entry cost or not. The vector $\mathbf{z}_t$ of exogenous state variables includes firm productivity, and market and firm characteristics that affect variable profit, fixed cost, and entry cost.[6] We specify each of these components in turn.

An active firm earns a profit $\pi(1, \mathbf{x}_t) + \varepsilon_t(1)$ where $\pi(1, \mathbf{x}_t)$ is equal to the variable profit (VP_t) minus fixed cost (FC_t), and minus entry cost (EC_t). The payoff to being inactive is $\pi(0, \mathbf{x}_t) + \varepsilon_t(0)$, where we make the normalization $\pi(0, \mathbf{x}_t) = 0$ for all possible values of $\mathbf{x}_t$. We assume that $\varepsilon_t(0)$ and $\varepsilon_t(1)$ are extreme value type 1 distributed with dispersion parameter $\sigma_\varepsilon = 1$. The variable profit function is $VP_t = [\theta_0^{VP} + \theta_1^{VP} z_{1t} + \theta_2^{VP} z_{2t}] \exp(\omega_t)$ where ω_t is the firm's productivity shock, z_{1t} and z_{2t} are exogenous state variables that affect the firm's price-cost margin in the market; and θ_0^{VP}, θ_1^{VP}, and θ_2^{VP} are parameters. The fixed cost is, $FC_t = \theta_0^{FC} + \theta_1^{FC} z_{3t}$, and the entry cost is, $EC_t = (1 - y_t)[\theta_0^{EC} + \theta_1^{EC} z_{4t}]$, where the term $(1 - y_t)$ indicates that the entry cost is paid only if the firm was not active in the market at the previous period, z_{3t} and z_{4t} are exogenous state variables, and θ's are parameters. The vector of structural

[6] We treat productivity as observable. For instance, using data on firms' output and inputs the researcher can estimate production function parameters and productivity taking into account the selection problem due to endogenous entry and exit decisions, e.g., Olley and Pakes (1996), Ackerberg, Caves, and Frazer (2015).

parameters in the payoff function is $\theta = (\theta_0^{VP}, \theta_1^{VP}, \theta_2^{VP}, \theta_0^{FC}, \theta_1^{FC}, \theta_0^{EC}, \theta_1^{EC})'$. The vector of exogenous state variables $\mathbf{z} = (z_1, z_2, z_3, z_4, \omega)$ has discrete and finite support. Each of the exogenous state variables takes K values. The dimension of the state space $|\mathcal{X}|$ is then $2 * K^5$. Each exogenous state variable follows a discrete-AR(1) process with an intercept parameter γ_0^j, slope parameter γ_1^j, and variance of the innovation shock σ_j^2. Similarly, productivity ω_t follows a discrete-AR(1) process with an intercept parameter γ_0^ω, slope parameter γ_1^ω, and variance of the innovation shock σ_ω^2. We use Tauchen's method to construct the transition probabilities of these discrete state variables (Tauchen, 1986).[7] The DGP used in our numerical and Monte Carlo experiments is summarized in table 1.

Table 1: Parameters in the DGP

Payoff Parameters:	$\theta_0^{VP} = 0.5, \quad \theta_1^{VP} = 1.0, \quad \theta_2^{VP} = -1.0$ $\theta_0^{FC} = 0.5, \quad \theta_1^{FC} = 1.0$ $\theta_0^{EC} = 1.0, \quad \theta_1^{EC} = 1.0$
Stochastic process of state variables:	For $j = 1, 2, 3, 4$, z_{jt} is AR(1) with $\gamma_0^j = 0.0$, $\gamma_1^j = 0.6$ ω_t is AR(1) with $\gamma_0^\omega = 0.2$, $\gamma_1^\omega = 0.9$ Low persistence model: $\sigma_\omega = \sigma_1 = \sigma_2 = \sigma_3 = \sigma_4 = 1$ High persistence model: $\sigma_\omega = \sigma_1 = \sigma_2 = \sigma_3 = \sigma_4 = 0.001$
Discount factor:	$\beta = 0.95$

4.2 Comparing solution methods

We compare VF, RVF, PF, and EE algorithms for six different dimensions of the state space $|\mathcal{X}|$: 64, 486, 2048, 6250, 15552, and 200,000 that correspond to values 2, 3, 4, 5, 6, and 10, respectively, for the number of points K in the support of each exogenous state variable. We use the same starting values to initialize the different algorithms. In principle, the relative performance of these methods may depend on the initial value. To check for this possibility,

[7] Let $\{z_j^{(k)} : k = 1, 2, ..., K\}$ be the support of the state variable z_j, and define the width values $w_j^{(k)} \equiv z_j^{(k+1)} - z_j^{(k)}$. Let $\tilde{z}_{jt}$ be a continuous latent variable that follows the AR(1) process $\tilde{z}_{jt} = \gamma_0^j + \gamma_1^j \tilde{z}_{jt-1} + e_{jt}$, with $e_{jt} \sim$ i.i.d. $N(0, \sigma_j^2)$. Then, the transition probability for the discrete state variable z_{jt} is given by: $\Phi([z_j^{(1)} + (w_j^{(1)}/2) - \gamma_0^j - \gamma_1^j z]/\sigma_j)$ for $z' = z_j^{(1)}$; $\Phi([z_j^{(k)} + (w_j^{(k)}/2) - \gamma_0^j - \gamma_1^j z]/\sigma_j) - \Phi([z_j^{(k-1)} + (w_j^{(k-1)}/2) - \gamma_0^j - \gamma_1^j z]/\sigma_j)$ for $z' = z_j^{(k)}$ with $2 \leq k \leq K-1$; and $1 - \Phi([z_j^{(K-1)} + (w_j^{(K-1)}/2) - \gamma_0^j - \gamma_1^j z]/\sigma_j)$ for $z' = z_j^{(K)}$.

we have implemented this experiment using ten different initial values, the same for all the algorithms. We find very small differences in the relative performance of the algorithms across the different initial values. The results in tables 3 and 4 are the averages over these initial values.

First, we analyze the degree of contraction of these algorithms. The number of iterations that a fixed point algorithm needs to achieve convergence is closely related to the degree of contraction of the mapping. To compare the degree of contraction, we calculate an approximation to the Lipschitz constants of the mappings. The Lipschitz constant of a mapping Γ in a complete metric space $(\mathcal{V}, \|.\|)$ is defined as the smallest constant c that satisfies the condition $\|\Gamma(\mathbf{V}) - \Gamma(\mathbf{W})\|/\|\mathbf{V} - \mathbf{W}\| \leq c$ for any two points $\mathbf{V}$ and $\mathbf{W}$ in the domain $\mathcal{V}$. For instance, using the sup-norm, the Lipschitz constant is defined as:

$$L(\Gamma) \equiv \sup_{\mathbf{V},\mathbf{W}\in\mathbb{R}^{|\mathcal{X}|}} \left[\frac{\sup_{\mathbf{x}\in\mathcal{X}} |\Gamma(\mathbf{x}, \mathbf{V}) - \Gamma(\mathbf{x}, \mathbf{W})|}{\sup_{\mathbf{x}\in\mathcal{X}} |V(\mathbf{x}) - W(\mathbf{x})|} \right] \tag{27}$$

Calculating the exact value of the Lipschitz constant for any of the mappings we consider is not a practical option because the dominion of all these mappings is infinite. As such we obtain the following approximation. Let $\{\mathbf{V}_k : k = 0, 1, ..., \mathcal{I}_{\mathbf{V}_0}\}$ be the sequence of values that we obtain by applying successive iterations in the mapping Γ given an initial value $\mathbf{V}_0$, where $\mathcal{I}_{\mathbf{V}_0}$ is the number of iterations to reach convergence. Then, we obtain an approximation (i.e., a lower bound) to the Lipschitz constant of this mapping by considering the ratios $\|\Gamma(\mathbf{V}) - \Gamma(\mathbf{W})\|/\|\mathbf{V} - \mathbf{W}\|$ at the pair of values $\mathbf{V}_k$ and $\mathbf{V}_{k+1}$ generated in the sequence. To obtain a better approximation, we generate sequences from many initial guesses, and take as our approximation to the Lipschitz constant the maximum over all these sequences. That is, the approximation to the Lipschitz constant of mapping Γ is:

$$\widetilde{L}(\Gamma) \equiv \max_{\mathbf{V}_0\in\mathcal{V}_0} \left\{ \max_{k\in\{0,1,...,\mathcal{I}_{\mathbf{V}_0}\}} \frac{\|\Gamma(\mathbf{V}_{k+1}) - \Gamma(\mathbf{V}_k)\|}{\|\mathbf{V}_{k+1} - \mathbf{V}_k\|} \right\} \tag{28}$$

Table 2 reports the Lipschitz constants of the four operators for the different dimensions of the state space, and for two different versions of the model according to the time persistence of the exogenous state variables (see table table 1): a model with *low persistence* that implies $\rho(\mathbf{F}_z) = 0.56$; and a model with *high persistence* that implies a $\rho(\mathbf{F}_z)$ close to one. The Lipschitz constants are very stable across the different dimensions of the state space. For the low persistence model, there is a very substantial difference in the degree of contraction of the VF operator and the other operators. The PF and the EE operators are always the strongest contractions. In the model with low persistence, the RVF operator is a stronger contraction than VF, but it is still far away from the EE and the PF operators. The good contraction properties of the RVF operator

Table 2: Degree of Contraction (Lipschitz Constants) of the mappings

	Low Persistence Model				High Persistence Model			
$\lvert\mathcal{X}\rvert$	**EE**	**PF**	**VF**	**RVF**	**EE**	**PF**	**VF**	**RVF**
64	0.20	0.14	0.95	0.59	0.34	0.56	0.95	0.95
486	0.18	0.13	0.95	0.54	0.34	0.47	0.95	0.95
2,032	0.18	0.11	0.95	0.53	0.31	0.86	0.95	0.95
6,250	0.18	0.10	0.95	0.53	0.32	0.69	0.95	0.95
15,552	0.18	0.09	0.95	0.53	0.28	0.62	0.95	0.95
200,000	0.18	-	0.95	0.53	0.28	-	0.95	0.95

completely disappear in the model with highly persistent state variables. Importantly, that is not the case for the EE operator: the Lipschitz constant of this operator increases when $\rho(\mathbf{F}_z)$ gets close to one but it is still substantially smaller than the Lipschitz constant of VF and RVF operators. This is the main result that we want to emphasize in table 2. The EE operator has strong contraction properties, similar and even better to the ones of the PF operator, regardless of the degree of persistence in the exogenous state variables.

Table 3 presents time per iteration, number of iterations, and total computation time for the different algorithms and dimensions of the state space for the model with low persistence of the state variables, which is a very optimistic scenario for the relative value iterations algorithm. For every dimension of the state space, iterating in the EE mapping is always the most efficient algorithm. The computational gains relative to the standard methods are very substantial, as shown in the columns at the bottom-right of the table reporting the ratio of computing times relative to the EE method. Furthermore, these gains increase with the dimension of the state space.

The PF algorithm is the least efficient algorithm despite that this operator is the stronger contraction and converges to the solution after only 5 iterations. The inefficiency of the PF algorithm comes from the cost of its valuation step that increases cubically with the dimension of the state space. EE is more efficient than VF and RVF both because its time-per-iteration is lower and because it is a stronger contraction such that it requires a smaller number of iterations. These advantages are quite substantial. The time-per-iteration of VF and RVF is more than twice that of the EE. For the largest dimension of the state space (still a modest dimension relative to most empirical applications), the total times of VF and RVF are 85 and 7 times, respectively, more than the EE method.

Table 4 compares the EE and RVF algorithms in the model with high persistence of the

Table 3: Comparison of Solution Methods – Low Persistence Model

Number of states	**Number of Iterations**				**Time per Iteration (in seconds)**			
$\lvert\mathcal{X}\rvert$	EE	PF	VF	RVF	EE	PF	VF	RVF
64 (= 2*2^5)	13	5	351	36	0.001	0.01	< 0.001	< 0.001
486 (= 2*3^5)	13	5	246	31	< 0.001	0.60	< 0.001	< 0.001
2048 (=2*4^5)	13	5	345	30	0.005	16.29	0.01	0.01
6250 (=2*5^5)	13	5	344	30	0.04	150.0	0.08	0.08
15552 (=2*6^5)	13	5	344	30	0.17	916.2	0.46	0.46
200,000 (=2*10^5)	13	$5^{(1)}$	344	30	21.05	198,978	67.64	64.47

Number of states	**Total Time (in seconds)**				**Time Ratios**			
$\lvert\mathcal{X}\rvert$	EE	PF	VF	RVF	$\frac{\text{EE}}{\text{EE}}$	$\frac{\text{PF}}{\text{EE}}$	$\frac{\text{VF}}{\text{EE}}$	$\frac{\text{RVF}}{\text{EE}}$
64 (= 2*2^5)	< 0.001	0.05	0.03	< 0.001	1	60	37.0	4.0
486 (=2*3^5)	0.01	3.02	0.35	0.03	1	300	26.6	3.4
2048 (=2*4^5)	0.05	81.44	2.76	0.27	1	1320	53.1	5.2
6250 (=2*5^5)	0.47	750.0	26.14	2.37	1	1,595	55.9	5.1
15552 (=2*6^5)	2.25	4,581	158.9	13.89	1	2,040	70.7	6.2
200,000 (=2*10^5)	273.6	$\simeq$ 1M$^{(1)}$	23,270	1,934	1	3,630	85.0	7.1

Note (1): Estimated number of iterations, and estimated total time given the observed time-per-iteration.

state variables. As we have shown in table 2 above, in this model the RVF does not have strong contraction properties and its computational properties are practically the same as VF iterations. In contrast, the EE still has strong contraction properties. As a result, table 4 shows that for this model the performance of the EE algorithm is strongly superior to RVF.

Rust (1987, 1988), Powell (2007), Bertsekas (2011), and Iskhakov, Lee, Rust, Schjerning, and Seo (2016) advocate using a hybrid value-policy iteration method. The algorithm starts with value function iterations until a loose convergence criterion is reached. Then, the algorithm switches to policy function iterations. When the switching point is appropriately tuned, this algorithm can be faster than both value function and policy function iterations. We have not reported results from this hybrid method in this experiment. The main reason is that, even for the moderate dimensions of the state space in tables 2 to 4, one single policy iteration takes almost one hundred times longer than the whole time to convergence of the EE method. In other words, even for this moderate dimension, the optimal hybrid algorithm is the pure value function iteration method.

Table 4: Comparison of EE and RVF Solution Methods – High Persistence Model

Number of states	**EE iterations**			**Relative value iterations**			**Ratio total time**
$\lvert\mathcal{X}\rvert$	# iter.	Time-per-iter.	Total time	# iter.	Time-per-iter.	Total time	**RVF / EE**
64 (=2*2^5)	24	< 0.001	0.002	378	< 0.001	0.05	24.5
486 (=2*3^5)	24	< 0.001	0.03	350	< 0.001	0.79	30.4
2048 (=2*4^5)	17	0.01	0.16	341	0.02	6.28	38.8
6250 (=2*5^5)	18	0.06	0.99	322	0.12	37.94	38.0
15552 (=2*6^5)	17	0.43	7.15	333	0.87	289.13	40.5
200,000 (=2*10^5)	16	46.2	721.1	319	98.5	18,818	42.5

Figure 1 presents the time-per-iteration versus the number of points in the state space, both in logarithms, for the different methods. This figure shows that the computational savings per-iteration of using the EE algorithm increase fast with the state space. Figure 2 shows the ratio of total computing time between the RVF and EE algorithms. We can see that the ratio increases very fast with the dimension of the state space – faster than exponentially, as the x-axis is in logarithms,

4.3 Application: Estimation of counterfactuals

We now study how sample-based methods perform in counterfactual policy experiments. Given the large dimension of the state space in actual applications, the exact computation of a solution of the model (under the factual and counterfactual scenarios) is computationally unfeasible, at least using standard solution methods of value function or policy function iterations. The purpose of this section is twofold. First, we show that the EE algorithm makes it feasible solving the model *exactly* when the dimension of the state space is relatively large, i.e., of the order of a few million values. We show that for this dimension, and using standard computing equipment, the EE method solves the model in a few hours while standard methods require computing times between a few weeks (for value function iterations) and more than one year (for policy function iterations). The second purpose of this section is to study the relative performance of these algorithms in empirical applications where the model cannot be solved exactly and we need to use sample-based versions of the operators.

Figure 1: Time-per-iteration vs. Number of States – in Logarithms

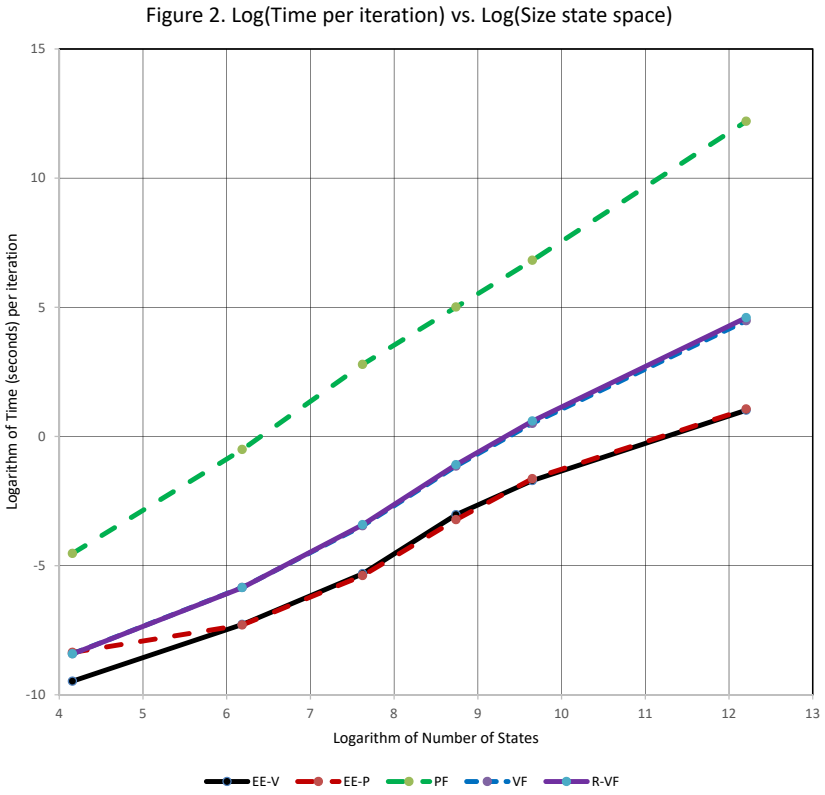

4.3.1 Counterfactual question

The counterfactual policy we consider is an increase in the cost of market entry. The presence of entry costs can generate misallocation in an industry. There may be potential entrants that are more productive than incumbent firms but are not willing to enter in the market and replace the less efficient firms because the entry cost makes this unprofitable. Furthermore, the presence of entry costs makes exit less attractive to incumbent firms, because re-entry is more expensive, and in this way, higher entry costs may discourage low-productivity incumbents from exiting. We are interested in the quantification of the net effect of entry cost on total industry productivity.

Suppose that the industry consists of N potential entrants, indexed by i. Competition in this industry is characterized by monopolistic competition, i.e., a single-agent model. The expected

Figure 2: Ratio of RVF Total Computing Time over EE Total Computing Time

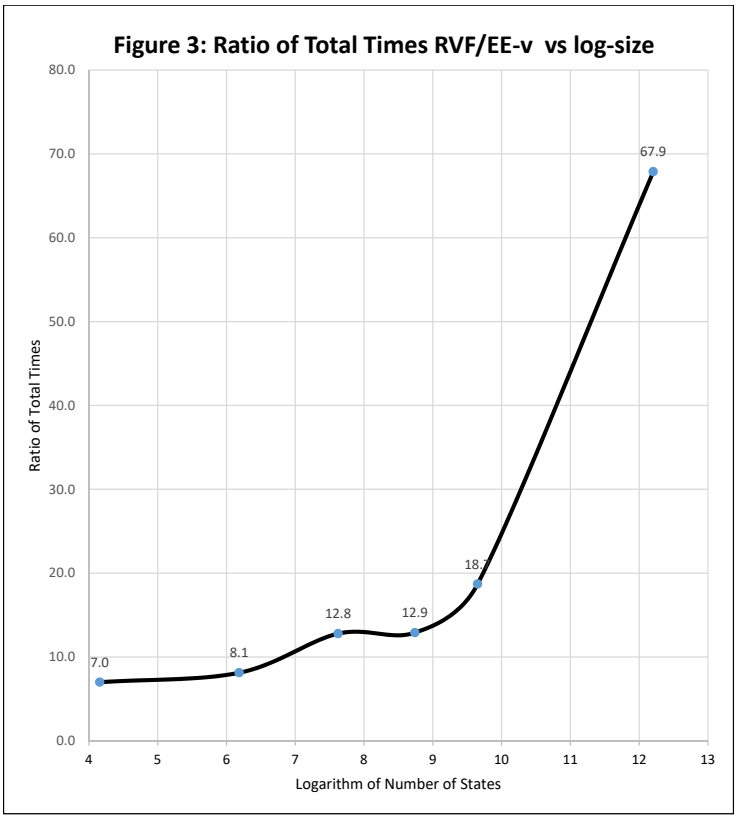

value of the total output produced by firms active in the industry, is:

$$Q^* = \mathbb{E}\left(\sum_{i=1}^{N} a_{it} \ \exp\left(\omega_{it}\right)\right) = N \ \sum_{\mathbf{z},\omega} p(\mathbf{z},\omega) \ \exp\left(\omega\right) \ f^*(\mathbf{z},\omega) \tag{29}$$

where $f^*(\mathbf{z},\omega)$ represents the steady-state or ergodic distribution of the exogenous variables, and $p(\mathbf{z},\omega)$ is the probability that a firm is active when the exogenous state variables take the values $(\mathbf{z},\omega)$, i.e., $p(\mathbf{z},\omega) \equiv \Pr(a_{it} = 1 | \mathbf{z}_t = \mathbf{z}, \omega_{it} = \omega)$. Note that $p(\mathbf{z},\omega)$ is different from the CCP function because the probability $p(\mathbf{z},\omega)$ is not conditional on the firm's incumbent status at previous period. However, by definition, the steady-state condition implies the following

relationship between $p(\mathbf{z},\omega)$ and the CCPs $P(0,\mathbf{z},\omega)$ and $P(1,\mathbf{z},\omega)$:

$$p(\mathbf{z},\omega) = (1 - p(\mathbf{z},\omega))\ P(0,\mathbf{z},\omega) + p(\mathbf{z},\omega)\ P(1,\mathbf{z},\omega) \tag{30}$$

Rearranging we get, $p(\mathbf{z},\omega) = P(0,\mathbf{z},\omega)/[1 - P(1,\mathbf{z},\omega) + P(0,\mathbf{z},\omega)]$. Using this expression to obtain the partial derivative $\partial p(\mathbf{z},\omega)/\partial EC$, where EC represents the entry cost, it is possible to show that the effect of entry cost on $p(\mathbf{z},\omega)$ and Q^* is ambiguous and it depends on the values of the state variables and the structural parameters. An increase in the entry cost has a negative effect on the probabilities of entry $P(0,\mathbf{z},\omega)$ for potential entrants, but it has a positive effect on the probability of staying in the industry $P(1,\mathbf{z},\omega)$ for incumbent firms. Since these effects are of opposite sign, the entry cost has an ambiguous net effect on the steady-state probability that a firm is active, $p(\mathbf{z},\omega)$, and on total output Q^*.

4.3.2 Exact solution using EE algorithm

Suppose that a researcher has estimated this model and is interested in the effect of the entry cost on expected total industry output: i.e., the effect of a change in EC on Q^*. More specifically, the counterfactual experiment we consider is an increase in the entry cost parameter θ_0^{EC} from 1 to 2.5. First, we show that when the dimension of the state space is relatively large (i.e., more than 1 million points), the EE algorithm can be used to solve the model *exactly* under the factual and counterfactual scenarios, while standard solution methods are infeasible.

In this experiment, we consider a DGP with a state space with $|\mathcal{Z}_1| = |\mathcal{Z}_2| = |\mathcal{Z}_3| = |\mathcal{Z}_4| = |\Omega| = 14$, such that the number of points in the complete state space is $|\mathcal{X}| = 2 * 14^5 = 1,075,648$. The values of the structural parameters are the ones given in table 1 above. We have solved the model, under the factual and counterfactual scenarios, by iterating in the EE mapping. Using standard computing equipment, the time-per-iteration was 546 seconds, and the total time (with 13 iterations) was 7,048 seconds, i.e., less than two hours. In contrast, the VF algorithm have a time-per-iteration of 1528 seconds and its total time is 146 hours, more than 6 days.[8]

Table 5 presents predictions from the exact solution of the model using the EE algorithm. We calculate the following average outcomes:

1. Average probability of being active: $p^* = \sum_{\mathbf{z},\omega} p(\mathbf{z},\omega)\ f^*(\mathbf{z},\omega)$.

2. Average probability of entry: $P_0^* = \sum_{\mathbf{z},\omega} P(0,\mathbf{z},\omega)\ f^*(\mathbf{z},\omega)$.

[8]This total time for the VF method is estimated using the condition that the number of iterations to convergence is 344, which is consistent with the results in table 3. The estimated total time for the PF algorithm is more than two years.

Table 5: Factual and Counterfactual Solutions Using EE Algorithm

	(1) Probability Being Active	(2) Entry Probability	(3) Exit Probability	(4) State Persistence	(5) Output
(A) Factual DGP	0.323	0.274	0.580	0.768	0.529
(B) Counterfactual DGP	0.258	0.157	0.513	0.884	0.423
Policy Effect: (B) - (A)	-0.065	-0.117	-0.068	+0.116	-0.106
(Percentage change)	(-20.1%)	(-42.7%)	(-11.7%)	(15.1%)	(-20.0%)

3. Average probability of exit: $1 - P_1^* = \sum_{\mathbf{z},\omega}(1 - P(1, \mathbf{z}, \omega))\ f^*(\mathbf{z}, \omega)$.

4. Average persistence in or out of the market: $\Pr(a_t = y_t) = \sum_{\mathbf{z},\omega}[p(z, \omega)\ P(1, \mathbf{z}, \omega) + (1 - p(z, \omega))\ (1 - P(0, \mathbf{z}, \omega))]\ f^*(\mathbf{z}, \omega)$.

5. Average output per firm: $q^* = \sum_{\mathbf{z},\omega} p(\mathbf{z}, \omega)\ \exp(\omega)\ f^*(\mathbf{z}, \omega)$.

In this experiment, the increase in entry cost reduces the average probability of being active (by 6.5 percentage points), the average probability of entry (by 11.7 percentage points), and the average probability of exit (by 6.8 percentage points). As non-incumbents are more likely to remain outside the market when the cost of entry rises, and incumbents are less likely to exit, we expect an increase in the persistence in the same state, which we see in column (4). Column (5) shows that the increase in the entry cost implies a 20% reduction in industry output.

4.3.3 Estimation of counterfactuals using sample-based operators

The dimension of the state space in this experiment – with more than a million points – is still small relative to the dimensions that we find in actual applications. For instance, in an empirical application of this model, the five exogenous state variables can be continuous variables with substantial variability across firms and over time. Even if the researcher is willing to discretize each of these continuous variables, an accurate representation of the distribution of these variables and their variation over time may require around one hundred grid points per variable. In our model, this implies a state space with approximately 20 billion points. We would like to compare the performance of the different approximation methods in this type of realistic scenario. However, that model cannot be solved exactly, and therefore, we cannot calculate the true approximation errors of the different methods. Instead, we consider that the true DGP is

the model in the previous section with $2 * 14^5$ possible states (approximately, 1 million states) and compare the performance of different approximation methods by fixing the computational time that the researcher is willing to pay to obtain an approximate solution of the model.

Table 6: Average Computation Times Using Sample-Based Operators

	VF	PF	EE
Points in spaces $\mathcal{Z}_{N_{method}}$	100	100	1,735
Number of Iterations	350	5	12.7
Time Per Iteration	$1.7\ 10^{-4}$	$7.2\ 10^{-3}$	$2.0\ 10^{-3}$
Total Time	0.062	0.036	0.026

To make a fair comparison, we keep the time required to solve the fixed points equal for the different solution/approximation methods. As defined above, $\mathcal{Z}$ and $\mathcal{Z}_N$ denote the true set and the sample set of values of the exogenous state variables. Let N_{EE}, N_{PF}, and N_{VF} be three integers smaller than the sample size (number of firms) N. Define $\mathcal{Z}_{N_{EE}}$, $\mathcal{Z}_{N_{PF}}$, and $\mathcal{Z}_{N_{VF}}$ as the sets of sample values of the exogenous state variables for the first N_{EE}, N_{PF}, and N_{VF} firms in the sample, respectively, and over the T sample periods. In the Monte Carlo experiments we solve the model by iterating in the sample-based EE mapping defined on the space $\mathcal{Z}_{N_{EE}}$, and by iterating in the PF and VF mappings defined on the spaces $\mathcal{Z}_{N_{PF}}$ and $\mathcal{Z}_{N_{VF}}$. In order to keep the computation time (roughly) the same across the three methods, we choose $N_{PF} < N_{VF} < N_{EE}$ appropriately. That is, we use a smaller space of the exogenous state variables when we solve the model using PF and VF iterations than when we use EE iterations. To solve the mappings on these reduced spaces, we must define the transition probabilities of the exogenous variables in each case. We use the true transition probabilities, normalized to the reduced space. Table 6 presents the number of points in the spaces $\mathcal{Z}_{N_{EE}}$, $\mathcal{Z}_{N_{VF}}$, and $\mathcal{Z}_{N_{PF}}$ and the average (across the Monte Carlo simulations) computing times of the different methods.

Table 7 presents the results of this Monte Carlo experiment. We estimate the same effects as in table 5. We implement three different iterative methods: EE, VF, and PF. We use 500 Monte Carlo simulations from the DGP and calculate Root Mean Squared Error (RMSE) and the Mean Absolute Bias (MAB) based on these simulations. The results indicate that the sample-based EE algorithm has considerably lower bias and RMSE than both the VF and the PF approximation methods for all the statistics. Systematically over the different parameters, the RMSE for the VF and PF methods are twice as large for the EE algorithm.

Table 7: Monte Carlo Experiment: Effects of Counterfactual Change in Entry Cost

	(1) Probability Being Active	(2) Entry Probability	(3) Exit Probability	(4) State Persistence	(5) Total Output
True Policy Effect	-0.065	-0.117	-0.068	+0.116	-0.106
Mean Absolute Bias					
VF iterations	0.036 (55.9%)	0.016 (13.7%)	0.016 (24.2%)	0.011 (9.4%)	0.070 (66.3%)
PF iterations	0.036 (55.9%)	0.016 (13.7%)	0.016 (24.2%)	0.011 (9.4%)	0.070 (66.3%)
EE iterations	0.024 (36.8%)	0.007 (6.2%)	0.010 (14.4%)	0.008 (7.1%)	0.040 (38.1%)
Root Mean Square Error					
VF iterations	0.044 (67.5%)	0.020 (17.1%)	0.020 (30.1%)	0.014 (11.8%)	0.088 (82.8%)
PF iterations	0.044 (67.5%)	0.020 (17.1%)	0.020 (30.1%)	0.014 (11.8%)	0.088 (82.8%)
EE iterations	0.024 (37.1%)	0.007 (6.3%)	0.010 (14.8%)	0.008 (7.2%)	0.041 (38.6%)

Note: In parentheses, the statistic in percentage over the true value of the parameter.

5 Conclusion

We show that these Finite Dependence / Euler Equations (FD/EE) representation of discrete choice MDP implies a fixed point mapping in the space of conditional choice values (EE operator). This operator is a contraction with a Lipschitz constant strictly smaller than the discount factor of the model We show that solving the DP problem by iterating in this operator provides very substantial computational gains relative to the standard solution methods of value function, relative value function, and policy function iterations.

We define a sample-based version of the EE operator and show that, for any sample, it inherits the same contraction properties as the exact EE operator. The sample-based EE operator is defined only at sample points of the exogenous state variables, and thus its dimensionality does not increase with the dimension of the state space. This operator can be used to obtain a consistent and asymptotically normal estimate of counterfactual CCPs.

We use Monte Carlo experiments to illustrate the computational gains associated with the EE algorithm. In the context of a dynamic model of entry and exit, computing the exact solution on a moderately sized state space using the standard policy iteration can take thousands of times as long as computing the exact solution using the EE mapping, implying that models that are computationally infeasible for all practical purposes using standard methods, are feasible using the method we propose. We then illustrate using Monte Carlo experiments the methods' relative ability to estimate the effect of a counterfactual increase in the cost of entry. We show that, for

a fixed computation time, the finite sample properties of the EE method are better than those of the estimator associated with policy iterations in terms of both mean squared error and mean absolute bias.

A Appendix – Proofs of Propositions

A.1 Williams-Daly-Zachary (WDZ) Theorem

The proof of Proposition 1 exploits the *Williams-Daly-Zachary* Theorem, which we first state and prove. Let $\Lambda(\widetilde{\mathbf{v}}) \equiv \{\Lambda(a, \widetilde{\mathbf{v}}) : a = 1, 2, ..., J\}$ be the Optimal Choice Probability with:

$$\Lambda(a, \widetilde{\mathbf{v}}) = \int 1\{\widetilde{v}_a + \widetilde{\varepsilon}_a \geq \widetilde{v}_j + \widetilde{\varepsilon}_j \quad \forall j\} \; g(\boldsymbol{\varepsilon}) \; d\boldsymbol{\varepsilon} \tag{31}$$

where $S(\widetilde{\mathbf{v}})$ is *McFadden's Social Surplus* function:

$$S(\widetilde{\mathbf{v}}) \equiv \int \max_{j \in \{0,1,...,J\}} [\widetilde{v}_j + \widetilde{\varepsilon}_j] \; g(\boldsymbol{\varepsilon}) \; d\boldsymbol{\varepsilon} \tag{32}$$

WDZ Theorem establishes that for any choice alternative a,

$$\frac{\partial S(\widetilde{\mathbf{v}})}{\partial \widetilde{v}_a} = \Lambda(a, \widetilde{\mathbf{v}}) \tag{33}$$

Proof. Define the function $m(\widetilde{\mathbf{v}} + \widetilde{\boldsymbol{\varepsilon}}) = \max_{j \in \{0,1,...,J\}} [\widetilde{v}_j + \widetilde{\varepsilon}_j]$, such that $S(\widetilde{\mathbf{v}}) = \int m(\widetilde{\mathbf{v}} + \widetilde{\boldsymbol{\varepsilon}}) \; g(\boldsymbol{\varepsilon}) \; d\boldsymbol{\varepsilon}$. It is clear that $\partial m(\widetilde{\mathbf{v}} + \widetilde{\boldsymbol{\varepsilon}})/\partial \widetilde{v}_a = 1\{\widetilde{v}_a + \widetilde{\varepsilon}_a \geq \widetilde{v}_j + \widetilde{\varepsilon}_j \; \forall j\}$. Note that $\partial S(\widetilde{\mathbf{v}})/\partial \widetilde{v}_a$ is equal to $\int \partial m(\widetilde{\mathbf{v}} + \widetilde{\varepsilon})/\partial \widetilde{v}_a \;\; g(\boldsymbol{\varepsilon}) \; d\boldsymbol{\varepsilon}$. Therefore,

$$\frac{\partial S(\widetilde{\mathbf{v}})}{\partial \widetilde{v}_a} = \int 1\{\widetilde{v}_a + \widetilde{\varepsilon}_a \geq \widetilde{v}_j + \widetilde{\varepsilon}_j \; \forall j\} \; g(\boldsymbol{\varepsilon}) \; d\boldsymbol{\varepsilon} = \Lambda(a, \widetilde{\mathbf{v}}) \tag{34}$$

A.2 Proof of Proposition 1

Given a vector of value differences $\widetilde{\mathbf{v}} \equiv (\widetilde{v}_a(y, \mathbf{z})$ for any $a, y, \mathbf{z})$, let $\widetilde{\mathbf{v}}(y, \mathbf{z})$ be the subvector $(\widetilde{v}_a(y, \mathbf{z})$ for any $a)$. Note that $\widetilde{\mathbf{v}}(y, \mathbf{z})$ is the argument of the Social Surplus function $S(.)$. For any state y and vector $\widetilde{\mathbf{v}}(0, \mathbf{z})$, define the "difference" surplus function $\widetilde{S}_y(\widetilde{\mathbf{v}}(0, \mathbf{z}))$ as follows:

$$\widetilde{S}_y(\widetilde{\mathbf{v}}(0, \mathbf{z})) \equiv S(\widetilde{\mathbf{v}}(0, \mathbf{z}) + \Delta(y, \mathbf{z})) - S(\widetilde{\mathbf{v}}(0, \mathbf{z})) = S(\widetilde{\mathbf{v}}(y, \mathbf{z})) - S(\widetilde{\mathbf{v}}(0, \mathbf{z})) \tag{35}$$

By Willians-Daly-Zachary Theorem, we that:

$$\frac{\partial \widetilde{S}_y(\widetilde{\mathbf{v}}(0, \mathbf{z}))}{\partial \widetilde{v}_a(0, \mathbf{z})} = \frac{\partial S(\widetilde{\mathbf{v}}(0, \mathbf{z}) + \Delta(y, \mathbf{z}))}{\partial \widetilde{v}_a(0, \mathbf{z})} - \frac{\partial S(\widetilde{\mathbf{v}}(0, \mathbf{z}))}{\partial \widetilde{v}_a(0, \mathbf{z})} = P(a, y, \mathbf{z}) - P(a, 0, \mathbf{z}) \tag{36}$$

where, with some abuse of notation, $P(a, y, \mathbf{z})$ is the CCP for alternative a given a vector of values $\widetilde{\mathbf{v}}(0, \mathbf{z}) + \Delta(y, \mathbf{z})$, and similarly, $P(a, 0, \mathbf{z})$ is the CCP for alternative a given a vector

of values $\widetilde{\mathbf{v}}(0,\mathbf{z})$. It is clear that the absolute value of this partial derivative is always strictly smaller that one. Define constant M as:

$$M \equiv sup_{a,y,\mathbf{z},\widetilde{\mathbf{v}}} \left| \frac{\partial \widetilde{S}_y(\widetilde{\mathbf{v}}(0,\mathbf{z}))}{\partial \widetilde{v}_a(0,\mathbf{z})} \right| \tag{37}$$

It is clear that $M < 1$.

By definition of mappings Γ_{EE} and $\widetilde{S}_y$, we have that:

$$\Gamma_{EE}(a,y,\mathbf{z},\widetilde{\mathbf{v}}) = c(a,y,\mathbf{z}) + \beta \sum_{\mathbf{z}'} \widetilde{S}_a(\widetilde{\mathbf{v}}(0,\mathbf{z}'))\ f_z(\mathbf{z}'|\mathbf{z}) \tag{38}$$

Therefore, we have that for any $\widetilde{\mathbf{v}}$ and $\widetilde{\mathbf{w}}$ in $\mathcal{V}_R$:

$$\Gamma_{EE}(a,y,\mathbf{z},\widetilde{\mathbf{v}}) - \Gamma_{EE}(a,y,\mathbf{z},\widetilde{\mathbf{w}}) = \beta \sum_{\mathbf{z}'} \left[\widetilde{S}_a(\widetilde{\mathbf{v}}(0,\mathbf{z}')) - \widetilde{S}_a(\widetilde{\mathbf{w}}(0,\mathbf{z}')) \right] f_z(\mathbf{z}'|\mathbf{z}) \tag{39}$$

Applying the Mean Value Theorem (MVT) to function $\widetilde{S}_a$ between $\widetilde{\mathbf{v}}(0,\mathbf{z}')$ and $\widetilde{\mathbf{w}}(0,\mathbf{z}')$, we have that there is a vector of values $\widetilde{\mathbf{v}}^*(0,\mathbf{z}')$ such that $\widetilde{S}_a(\widetilde{\mathbf{v}}(0,\mathbf{z}')) - \widetilde{S}_a(\widetilde{\mathbf{w}}(0,\mathbf{z}')) = \nabla \widetilde{S}_a(\widetilde{\mathbf{v}}^*(0,\mathbf{z}'))$ $[\widetilde{\mathbf{v}}(0,\mathbf{z}') - \widetilde{\mathbf{w}}(0,\mathbf{z}')]$, where $\nabla \widetilde{S}_a$ represents the gradient vector. Thus, we have:

$$\Gamma_{EE}(a,y,\mathbf{z},\widetilde{\mathbf{v}}) - \Gamma_{EE}(a,y,\mathbf{z},\widetilde{\mathbf{w}}) = \beta \sum_{\mathbf{z}'} \nabla \widetilde{S}_a(\widetilde{\mathbf{v}}^*(0,\mathbf{z}'))\, [\widetilde{\mathbf{v}}(0,\mathbf{z}') - \widetilde{\mathbf{w}}(0,\mathbf{z}')]\, f_z(\mathbf{z}'|\mathbf{z}) \tag{40}$$

Taking into account this equation, we have that:

$$\begin{array}{rcl} |\Gamma_{EE}(a,y,\mathbf{z},\widetilde{\mathbf{v}}) - \Gamma_{EE}(a,y,\mathbf{z},\widetilde{\mathbf{w}})| & \leq & \beta \sum_{\mathbf{z}'} \left|\left| \nabla \widetilde{S}_a(\widetilde{\mathbf{v}}^*(0,\mathbf{z}')) \right|\right|\ ||\widetilde{\mathbf{v}}(0,\mathbf{z}') - \widetilde{\mathbf{w}}(0,\mathbf{z}')||\, f_z(\mathbf{z}'|\mathbf{z}) \\ & \leq & \beta M \sum_{\mathbf{z}'} ||\widetilde{\mathbf{v}}(0,\mathbf{z}') - \widetilde{\mathbf{w}}(0,\mathbf{z}')||\, f_z(\mathbf{z}'|\mathbf{z}) \\ & \leq & \beta M\, ||\widetilde{\mathbf{v}}(0) - \widetilde{\mathbf{w}}(0)|| \end{array} \tag{41}$$

where $\widetilde{\mathbf{v}}(0) \equiv (\widetilde{v}_a(0,\mathbf{z})\ \forall a,\mathbf{z})$, and similarly for $\widetilde{\mathbf{w}}(0)$. Finally, note that for any $\widetilde{\mathbf{v}}$ and $\widetilde{\mathbf{w}}$ in $\mathcal{V}_R$:

$$\begin{array}{rcl} ||\widetilde{\mathbf{v}} - \widetilde{\mathbf{w}}|| & = & sup_{a,y,\mathbf{z}} |v_a(y,\mathbf{z}) - w_a(y,\mathbf{z})| \\ & = & sup_{a,y,\mathbf{z}} |v_a(0,\mathbf{z}) + \Delta_a(y,\mathbf{z}) - w_a(0,\mathbf{z}) - \Delta_a(y,\mathbf{z})| \\ & = & sup_{a,\mathbf{z}} |v_a(0,\mathbf{z}) - w_a(0,\mathbf{z})| \\ & = & ||\widetilde{\mathbf{v}}(0) - \widetilde{\mathbf{w}}(0)|| \end{array} \tag{42}$$

Therefore, we have proved that:

$$||\Gamma_{EE}(\widetilde{\mathbf{v}}) - \Gamma_{EE}(\widetilde{\mathbf{w}})|| \leq \beta\ M\ \ ||\widetilde{\mathbf{v}} - \widetilde{\mathbf{w}}|| \qquad \text{with } 0 < M < 1 \qquad \blacksquare \tag{43}$$

A.3 Proof of Proposition 2

(A) $\Gamma_{EE}^{(N)}(\widetilde{\mathbf{v}},\theta)$ is a contraction. This proof is a straightforward extension of the proof of Proposition 1. For a given sample with N cross-sectional observations, denote the space in which the value differences live to be $\mathcal{V}_N$. Additionally, let $\widehat{f}_N(\mathbf{z}'|\mathbf{z})$ be the estimated transition probabilities for the exogenous state vector $\mathbf{z}$. For any $\tilde{v}$ and $\tilde{w}$ in $\mathcal{V}_N$, the sample analog of equation (40) is:

$$\Gamma_{EE}^{(N)}(a,y,\mathbf{z},\widetilde{\mathbf{v}},\theta) - \Gamma_{EE}^{(N)}(a,y,\mathbf{z},\widetilde{\mathbf{w}},\theta) = \beta \sum_{\mathbf{z}'} \nabla \widetilde{S}_a(\widetilde{\mathbf{v}}^*(0,\mathbf{z}'))\,[\widetilde{\mathbf{v}}(0,\mathbf{z}') - \widetilde{\mathbf{w}}(0,\mathbf{z}')]\,\widehat{f}_N(\mathbf{z}'|\mathbf{z}) \quad (44)$$

We then have:

$$\begin{aligned} \left|\Gamma_{EE}^{(N)}(a,y,\mathbf{z},\widetilde{\mathbf{v}},\theta) - \Gamma_{EE}^{(N)}(a,y,\mathbf{z},\widetilde{\mathbf{w}},\theta)\right| &\leq \beta \textstyle\sum_{\mathbf{z}'} \left|\left|\nabla \widetilde{S}_a(\widetilde{\mathbf{v}}^*(0,\mathbf{z}'))\right|\right| \; ||\widetilde{\mathbf{v}}(0,\mathbf{z}') - \widetilde{\mathbf{w}}(0,\mathbf{z}')||\, \widehat{f}_N(\mathbf{z}'|\mathbf{z}) \\ &\leq \beta M \textstyle\sum_{\mathbf{z}'} ||\widetilde{\mathbf{v}}(0,\mathbf{z}') - \widetilde{\mathbf{w}}(0,\mathbf{z}')||\, \widehat{f}_N(\mathbf{z}'|\mathbf{z}) \\ &\leq \beta M\, ||\widetilde{\mathbf{v}}(0) - \widetilde{\mathbf{w}}(0)|| \qquad \text{with } 0 < M < 1 \qquad \blacksquare \end{aligned} \quad (45)$$

(B) $\Gamma_{EE}^{(N)}(\widetilde{\mathbf{v}},\theta)$ converges uniformly to $\mathbf{\Gamma}_{EE}(\widetilde{\mathbf{v}},\theta)$. We need to show that, for any value of $(a,\mathbf{x})$, $\sup_{\widetilde{\mathbf{v}},\theta} \left|\mathbf{\Gamma}_{EE}^{(N)}(a,\mathbf{x},\widetilde{\mathbf{v}},\theta) - \mathbf{\Gamma}_{EE}(a,\mathbf{x},\widetilde{\mathbf{v}},\theta)\right| \longrightarrow_p 0$. Or equivalently, for any $\varepsilon > 0$ and any $\delta > 0$, there is an integer $N_0(\varepsilon,\delta)$ such that for every value of $(\widetilde{\mathbf{v}},\theta)$ and any $N > N_0(\varepsilon,\delta)$ we have that,

$$\Pr\left(\left|\mathbf{\Gamma}_{EE,v}^{(N)}(a,\mathbf{x},\widetilde{\mathbf{v}},\theta) - \mathbf{\Gamma}_{EE,v}(a,\mathbf{x},\widetilde{\mathbf{v}},\theta)\right| > \varepsilon\right) < \delta \quad (46)$$

For the rest of this proof and for the sake of notational simplicity, we omit the arguments $(a,\mathbf{x})$ and use γ to represent the vector of parameters $(\widetilde{\mathbf{v}},\theta)$. By definition, we have that $\mathbf{\Gamma}_{EE}(\gamma) = \sum_{\mathbf{z}'\in\mathcal{Z}} h(\mathbf{z}',\gamma)\; f(\mathbf{z}')$ and $\mathbf{\Gamma}_{EE}^{(N)}(\gamma) = \sum_{\mathbf{z}'\in\mathcal{Z}_N} h(\mathbf{z}',\gamma)\; \hat{f}_N(\mathbf{z}')$, where: (a) $h(\mathbf{z}',\gamma)$ is a function of payoffs and choice probabilities that comes from the Euler equation; (b) $h(\mathbf{z}',\gamma)$ is a bounded function such that $h^* \equiv \sup_{\mathbf{z}',\gamma} |h(\mathbf{z}',\gamma)| < \infty$; (c) f is the true population distribution (transition probability) of $\mathbf{z}'$, and $\hat{f}_N$ is the empirical distribution based on the sample; (d) for convenience, and without loss of generality, we consider that $\hat{f}_N(\mathbf{z}') = 0$ for values $\mathbf{z}'$ outside the sample set $\mathcal{Z}_N$, such that we can write $\mathbf{\Gamma}_{EE}^{(N)}(\gamma) = \sum_{\mathbf{z}'\in\mathcal{Z}} h(\mathbf{z}',\gamma)\; \hat{f}_N(\mathbf{z}')$; and (e) $\hat{f}_N$ is a uniformly consistent estimator of f and this implies that, for any $\varepsilon_f > 0$ and any $\delta_f > 0$, there is an integer $N_f(\varepsilon_f,\delta_f)$ such that for any $N > N_f(\varepsilon_f,\delta_f)$ we have that,

$$\Pr\left(\sup_{\mathbf{z}'\in\mathcal{Z}} \left|\hat{f}_N(\mathbf{z}') - f(\mathbf{z}')\right| > \varepsilon_f\right) < \delta_f \quad (47)$$

Given points (a) to (e), we now prove uniform convergence of $\mathbf{\Gamma}_{EE}^{(N)}$ to $\mathbf{\Gamma}_{EE}$. Note that for

any N and γ,

$$\begin{aligned}\left|\mathbf{\Gamma}_{EE}^{(N)}(\gamma) - \mathbf{\Gamma}_{EE}(\gamma)\right| &= \left|\sum_{\mathbf{z}'\in\mathcal{Z}}\left[\hat{f}_N(\mathbf{z}') - f(\mathbf{z}')\right] h(\mathbf{z}',\gamma)\right| \\ &\leq h^* \sum_{\mathbf{z}'\in\mathcal{Z}}\left|\hat{f}_N(\mathbf{z}') - f(\mathbf{z}')\right| \\ &\leq h^* \sup_{\mathbf{z}'\in\mathcal{Z}}\left|\hat{f}_N(\mathbf{z}') - f(\mathbf{z}')\right|\end{aligned} \tag{48}$$

This implies that, for any $\varepsilon > 0$,

$$\begin{aligned}\Pr\left(\left|\mathbf{\Gamma}_{EE}^{(N)}(\gamma) - \mathbf{\Gamma}_{EE}(\gamma)\right| > \varepsilon\right) &\leq \Pr\left(h^* \sup_{\mathbf{z}'\in\mathcal{Z}}\left|\hat{f}_N(\mathbf{z}') - f(\mathbf{z}')\right| > \varepsilon\right) \\ &= \Pr\left(\sup_{\mathbf{z}'\in\mathcal{Z}}\left|\hat{f}_N(\mathbf{z}') - f(\mathbf{z}')\right| > \frac{\varepsilon}{h^*}\right)\end{aligned} \tag{49}$$

Therefore, for any $\varepsilon > 0$ and $\delta > 0$, we can fix $\varepsilon_f = \varepsilon/h^*$, $\delta_f = \delta$, and $N_0(\varepsilon,\delta) = N_f(\varepsilon_f,\delta_f)$, such that for any $N > N_0(\varepsilon,\delta)$ we have that $\Pr\left(\left|\mathbf{\Gamma}_{EE}^{(N)}(\gamma) - \mathbf{\Gamma}_{EE}(\gamma)\right| > \varepsilon\right) < \delta$. ■

A.4 Proof of Proposition 3

(A) Consistency. For notational simplicity, we omit now θ^* as an argument in functions $\mathbf{\Gamma}_{EE}(\widetilde{v},\theta^*)$ and $\mathbf{\Gamma}_{EE}^{(N)}(\widetilde{v},\theta^*)$. The true $\widetilde{v}^*$ is defined as the unique fixed point $\widetilde{v}^* = \mathbf{\Gamma}_{EE}(\widetilde{v}^*)$, and $\widetilde{v}_N^*$ is defined as the unique fixed point $\widetilde{v}_N^* = \mathbf{\Gamma}_{EE}^{(N)}(\widetilde{v}_N^*)$. Given that, (a) $\mathbf{\Gamma}_{EE}(.)$ is a continuous function, and (b) $\mathbf{\Gamma}_{EE}^{(N)}$ converges uniformly in probability to $\mathbf{\Gamma}_{EE,v}$, we have by Slutsky's Theorem that $\widetilde{v}_N^*$ converges in probability to $\widetilde{v}^*$.

(B) Asymptotic Normality. By the fixed point conditions that define the value vectors $\widetilde{v}_N^*$ and $\widetilde{v}^*$, we have that:

$$\begin{aligned}\widetilde{v}_N^* - \widetilde{v}^* &= \mathbf{\Gamma}_{EE}^{(N)}(\widetilde{v}_N^*) - \mathbf{\Gamma}_{EE}(\widetilde{v}^*) \\ &= \sum_{\mathbf{z}\in\mathcal{Z}} h(\mathbf{z},\widetilde{v}_N^*)\ \hat{f}_N(\mathbf{z}) - \sum_{\mathbf{z}\in\mathcal{Z}} h(\mathbf{z},\widetilde{v}^*)\ f(\mathbf{z})\end{aligned} \tag{50}$$

Applying the Mean Value Theorem to the vector function $h(\mathbf{z},\widetilde{v}_N^*)$ at $\widetilde{v} = \widetilde{v}^*$, and using the consistency of $\widetilde{v}_N^*$ (i.e., $\widetilde{v}_N^* = \widetilde{v}^* + o_p(1)$), we have that:

$$h(\mathbf{z},\widetilde{v}_N^*) = h(\mathbf{z},\widetilde{v}^*) + \frac{\partial h(\mathbf{z},\widetilde{v}^*)}{\partial \widetilde{v}'}\left(\widetilde{v}_N^* - \widetilde{v}^*\right) + o_p(1) \tag{51}$$

Substituting this expression into (50) and using the consistency of $\hat{f}_N$ (i.e., $\hat{f}_N(\mathbf{z}) = \hat{f}_N(\mathbf{z}) + o_p(1)$), we obtain:

$$\widetilde{v}_N^* - \widetilde{v}^* = \sum_{\mathbf{z}\in\mathcal{Z}} h(\mathbf{z},\widetilde{v}^*) \left(\hat{f}_N(\mathbf{z}) - f(\mathbf{z})\right) + \left[\sum_{\mathbf{z}\in\mathcal{Z}} f(\mathbf{z}) \frac{\partial h(\mathbf{z},\widetilde{v}^*)}{\partial \widetilde{v}'}\right] (\widetilde{v}_N^* - \widetilde{v}^*) + o_p(1) \tag{52}$$

Solving for $\widetilde{v}_N^* - \widetilde{v}^*$,

$$\widetilde{v}_N^* - \widetilde{v}^* = [\mathbf{I} - \mathbf{D}]^{-1} \mathbf{H} \left(\hat{\mathbf{f}}_N - \mathbf{f}\right) + o_p(1) \tag{53}$$

where: $\mathbf{I}$ is the identity matrix; $\mathbf{D}$ is the matrix $\sum_{\mathbf{z}\in\mathcal{Z}} f(\mathbf{z})\, \partial h(\mathbf{z},\widetilde{v}^*)/\partial \widetilde{v}'$; $\hat{\mathbf{f}}_N$ and $\mathbf{f}$ are the column vectors that contain the probabilities $\hat{f}_N(\mathbf{z})$ and $f(\mathbf{z})$, respectively, for every value of $\mathbf{z}$; and $\mathbf{H}$ is the matrix with columns $[h(\mathbf{z}^{(1)},\widetilde{v}^*)$, $h(\mathbf{z}^{(2)},\widetilde{v}^*)$, ..., $h(\mathbf{z}^{(|\mathcal{Z}|)},\widetilde{v}^*)]$. Under mild regularity conditions, a standard Central Limit Theorem implies that the frequency estimator $\mathbf{f}_{(N)}$ is such that $\sqrt{N}\left(\hat{\mathbf{f}}_N - \mathbf{f}\right)$ converges in distribution to $N(0, \mathbf{V}_f)$. Then, applying the Mann-Wald Theorem we have that:

$$\sqrt{N}\,(\widetilde{v}_N^* - \widetilde{v}^*) \to_d N\left(0,\ [\mathbf{I} - \mathbf{D}]^{-1} \mathbf{H}\ \mathbf{V}_f\ \mathbf{H}'\ [\mathbf{I} - \mathbf{D}']^{-1}\right) \qquad \blacksquare \tag{54}$$

References

Ackerberg, D., K. Caves, and G. Frazer (2015): "Identification properties of recent production function estimators," *Econometrica*, 83(6), 2411–2451.

Aguirregabiria, V., and A. Magesan (2013): "Euler equations for the estimation of dynamic discrete choice structural models," in *Structural Econometric Models (Advances in Econometrics)*, vol. 31, pp. 3–44. Emerald Group Publishing Limited.

Arcidiacono, P., and R. Miller (2011): "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity," *Econometrica*, 79(6), 1823–1867.

Bertsekas, D. (2011): "Approximate policy iteration: A survey and some new methods," *Journal of Control Theory and Applications*, 9(3), 310–335.

Bray, R. L. (2019): "Strong convergence and dynamic economic models," *Quantitative Economics*, 10(1), 43–65.

Hansen, L., and K. Singleton (1982): "Generalized instrumental variables estimation of nonlinear rational expectations models," *Econometrica*, 50(5), 1269–1286.

Hotz, V., and R. Miller (1993): "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 60(3), 497–529.

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.

Morton, T. E., and W. E. Wecker (1977): "Discounting, ergodicity and convergence for Markov decision processes," *Management Science*, 23(8), 890–900.

Olley, G., and A. Pakes (1996): "The Dynamics of Productivity in the Telecommunications Equipment," *Econometrica*, 64(6), 1263–1297.

Powell, W. (2007): *Approximate Dynamic Programming: Solving the curses of dimensionality*, vol. 703. John Wiley & Sons.

Puterman, M. (1994): *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons. New Jersey.

Puterman, M., and S. Brumelle (1979): "On the convergence of policy iteration in stationary dynamic programming," *Mathematics of Operations Research*, 4(1), 60–69.

Rust, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, (5), 999–1033.

——— (1988): "Maximum likelihood estimation of discrete control processes," *SIAM journal on control and optimization*, 26(5), 1006–1024.

——— (1994): "Structural estimation of Markov decision processes," *Handbook of Econometrics*, 4, 3081–3143.

——— (1996): "Numerical dynamic programming in economics," *Handbook of computational economics*, 1, 619–729.

——— (1997): "Using randomization to break the curse of dimensionality," *Econometrica*, pp. 487–516.

Tauchen, G. (1986): "Finite state markov-chain approximations to univariate and vector autoregressions," *Economics Letters*, 20(2), 177–181.