Estimation and Counterfactuals in Dynamic Discrete Games Using an Euler-Equations Policy-Iteration Mapping

Victor Aguirregabiria
University of Toronto and CEPR

Arvind Magesan
University of Calgary

June 24, 2014
PRELIMINARY AND INCOMPLETE

Abstract

In this paper, we derive marginal conditions of optimality (i.e., Euler equations) for a general class of dynamic discrete choice structural models, including single-agent models and dynamic games. We show that these equations define a policy iteration mapping in the space of players' strategies (i.e., players' choice probabilities) such that the vector of equilibrium strategies is a fixed point of this mapping. We use this Euler Equations - Policy Iteration (EE-PI) mapping to define two-step (or K-step) GMM estimators of structural parameters, and estimators of counterfactual experiments using the estimated model. The Euler Equations - Policy Iteration mapping and the estimators built on it have important advantages with respect to the standard policy iteration mapping. First, the evaluation of the EE-PI mapping does not involve the computation of infinite period forward presents values, but only expectations of payoffs a few periods forward, i.e., N periods forward in a game with N heterogeneous players. This property results into substantial computational savings. Second, GMM estimators using a sample-based EE-PI mapping are consistent and are not subject to the curse of dimensionality in the computation of present values. These estimators do not suffer of an asymptotic bias associated to the approximation of value functions. We present Monte Carlo experiments to illustrate the computational gains of these methods.

Keywords: Dynamic Discrete structural models; Estimation; Counterfactual experiments; Euler equations; Moment conditions; Approximation bias.

JEL: C13; C35; C51; C61

Corresponding Author: Victor Aguirregabiria. Department of Economics. University of Toronto. 150 St. George Street. Toronto, Ontario `victor.aguirregabiria@utoronto.ca`

Arvind Magesan. Department of Economics. University of Calgary. 2500 University Drive, N.W. Calgary, Alberta `arvindmagesan@gmail.com`

1 Introduction

The development of two-step estimation methods represents a substantial methodological contribution in the econometrics of dynamic discrete choice structural models. The work by Hotz and Miller (1993) and Hotz et al. (1994) represent seminal contributions in this literature. Other contributions in this literature include the sequential policy iteration estimators in Aguirregabiria and Mira (2002), and the extensions of these methods to dynamic games in Aguirregabiria and Mira (2007) and Bajari, Benkard and Levin (2007). Since these methods avoid the solution of the dynamic programming problems that describe individuals' behavior, they make the estimation of richer specifications with larger state spaces computationally feasible. Nevertheless, the implementation of these methods still requires the computation of expectations (present values) defined as integrals or summations over the space of state variables. In applications with continuous state variables or with very large state spaces (e.g., dynamic games with heterogeneous players), the exact solution of expectations or present values is an intractable problem. To deal with this dimensionality problem, applied researchers use approximation techniques such as discretization, Monte Carlo simulation, polynomials, sieves, neural networks, etc. Replacing *true* expected values with approximations introduces an approximation error, and this error induces a statistical bias in the estimation of the parameter of interests. To deal with this dimensionality problem, Aguirregabiria and Magesan (2013) derive marginal conditions of optimality for a general class of single-agent dynamic discrete choice models and show that these Euler equations provide moment conditions that can be used to estimate structural parameters without the need to solve or approximate infinite present values. Given the similarity of these optimality conditions with the ones for dynamic decision models with continuous decision variables (Hansen and Singleton, 1982), we denote these conditions as *Euler Equations* (EEs). Estimators based on moment conditions from these Euler equations are not subject to bias induced by the approximation of value functions. Arcidiacono and Miller (2011 and 2013) propose an approach in the same spirit for a class models with a *finite-dependence* representation.[1]

A common motivation for estimating dynamic structural models is the prediction of agents' behavior and outcomes in scenarios different in some dimension, from the scenario that generated the data. These predictions are useful because they allow the researcher to study the effects of policies that do not occur in the data, e.g., a new tax/subsidy, or a change in some of the structural parameters of the model. These predictions are typically described as *counterfactual experiments*. Despite the development of two-step estimation methods, the computation of counterfactual exper-

[1] The concept of finite dependence was introduced in Altug and Miller (1998). Arcidiacono and Miller (2011, page 1824) provide the following simple and intuitive definition: "when two choice sequences with different initial decisions lead to the same distribution of states after a few periods, we say there is finite dependence". They provide a more general version of that definition in their 2013 working paper (Theorem 6) which we believe is equivalent to the sufficient conditions for the existence of Euler Equations in our paper.

iments remains a serious computational challenge in this literature. The methods available for the estimation of these counterfactuals require the full solution of, or at least the approximation to the dynamic programming problem. In most applications, the exact solution is not computationally feasible, and the approximation methods imply estimation errors that are potentially large and do not converge to zero as the sample size goes to infinity. In general, there is not a method for the estimation of counterfactuals that is both computationally feasible (i.e., it avoids the full solution of the dynamic programming problem) and econometrically consistent. The researcher must trade off these two criteria in deciding which method to use.

In this context, the contribution of this paper is threefold. First, we extend the derivation of Euler Equations (EEs) to dynamic discrete games. This extension is not trivial. In general, standard Euler equations, based on marginal conditions of optimality at two consecutive periods, do not hold for games with heterogeneous players. We derive EEs that involve marginal expected payoffs at $N + 1$ periods, where N is the number of players in the game. Second, we show that these EEs imply a fixed point mapping in space of an individual's decision rule (i.e., conditional choice probabilities) and that this fixed point mapping is a contraction. In contrast to the *standard policy-iteration mapping* in DP problems, this *EE-policy-iteration mapping* does not involve the computation of infinite-period-forward present values, but only the N-period-forward expectations. For this reason, the *EE-policy-iteration mapping* provides a method to compute the solution of the DP problem that is computationally much cheaper than standard policy iterations. Still, the consistent estimation of counterfactuals using the *EE-policy-iteration mapping* is subject to a curse of dimensionality. Third, we define a sample approximation to *EE-policy-iteration mapping.* This sample approximation, which has the contraction mapping property, is defined only at sample points of the state variables, and thus its dimensionality is finite and relatively small. We show that the unique fixed point of this sample *EE-policy-iteration mapping* is a consistent estimator of the true counterfactual optimal decision rule. In contrast, using a sample version of the standard policy iteration mapping does not provide a consistent estimator of the optimal decision rule. The reason is that the sample *EE-policy-iteration mapping* is related to sample moment conditions that are satisfied asymptotically as the sample size goes to infinity while that is not the case for the sample conditions implied by the standard policy-iteration mapping.

Another useful application of the *EE-policy-iteration mapping* is that it can be used to define a sequential procedure to reduce the final sample bias of the two-step estimation methods based on Euler Equations. The idea is that, once a researcher has estimated the parameters of the structural model in the second step (the first step is the estimation of choice probabilities), she can "update" or improve upon the initial estimate of the choice probabilities iterate in the *EE-policy-iteration mapping* to find a fixed point that is consistent with the estimated parameters. With this updated estimate of choice probabilities in hand, the researcher can re-estimate the parameters, and continue in this fashion until the parameter estimates do not change when the probabilities are updated. The main advantage of using this sequential method is that the recalculation of the Conditional Choice Probabilities at each iteration of this method does not require the computation of present values, and that these probabilities only need to be estimated at values of the state variables observed in the sample. Again, this substantially reduces the computational cost of estimating the model.

We illustrate the relative computational gains associated with our approach using two sets of Monte Carlo experiments. In the first set, we show that in the context of a dynamic model of entry and exit, a full solution using the standard policy iteration can take anywhere from 18 to 75 times as long as a full solution using the *EE-policy-iteration mapping* depending on the dimensionality of the state space, implying that models that are computationally infeasible for all practical purposes using standard methods, are feasible using the method we propose. In particular, we show that although the steps to convergence using the standard policy iteration mapping are greater,[2] meaning that the standard policy iteration mapping takes fewer iterations to a fixed point, each iteration of the *EE-policy-iteration mapping* is so computationally inexpensive that it ends up being considerably faster in total time to convergence. Second we examine the finite sample properties of the two estimators by considering a counterfactual policy question in the context of the dynamic model of entry and exit. In particular we study how the two methods do in predicting firm behavior in response to a counterfactual increase in the cost of entry, holding the computation time of the two methods fixed. Holding the computation time fixed requires us to estimate the counterfactual probabilities on a smaller space in the case of the standard mapping, and this can introduce bias and higher variance in the estimates.[3] We show that, holding the computation time of the model fixed across the two approaches, the finite sample properties of the estimator that uses the *EE-policy-iteration mapping* are better than those of the estimator associated with the standard mapping, in the sense that it has lower mean squared error.

The paper is related to a literature that exploits properties of the dynamic discrete decision problem to obtain a representation of the model that does not involve value functions. The key papers here are Hotz and Miller (1993) and Arcidiacono and Miller (2011), who show that models that possess a "finite-dependence" property permit a representation whereby the choice probabilities can be expressed as a function of expected payoffs at a finite number of states, meaning that a researcher need not compute value functions to estimate the structural parameters of the model. Aguirregabiria and Magesan (2013) similarly show that under a "weak stationarity" condition[4] on the dynamic decision process, the model permits an EE representation which depends only on

[2] This is not surprising, as the standard policy iteration mapping is a composite mapping, while the *EE-policy-iteration mapping* is not.

[3] This is a realistic approach, however, as state space reduction methods are often used to make computationally infeasible methods feasible.

[4] Loosely speaking, a dynamic decision problem satisfies the weak stationarity condition if, for any state X today the set of states reachable from decision d and decision d' have at least one common element.

the expected payoffs today and tomorrow. The key contributions of this paper in this regard is that, while the aforementioned papers propose representations that are useful for the estimation of structural parameters, we propose a novel mapping for computing equilibrium probabilities, which can be used to make counterfactual prediction regardless of how the parameters are estimated. To the best of our knowledge, this is the first paper to do so. This is also the first paper that we know of that proposes an *EE-policy-iteration mapping* for reducing the bias associated with two-step estimation methods. Finally, our paper is related to Rust (1997), who proposes a randomization method to reduce the computational cost of solving a dynamic programming problem with continuous state variables.

The rest of the paper is organized as follows. Section 2 presents the model and the derivation of Euler equations, and describes the two-step estimator based on these equations. Section 3 introduces the *Euler Equation policy-iteration mapping* and compares it to the standard policy iteration mapping. Section 4 presents our method to estimate consistently counterfactuals, derives its statistical and computational properties, and compares them with those from the standard methods in the literature. In section 5, we present results from Monte Carlo experiments where we illustrate the advantages our proposed method. We summarize and conclude in section 6.

2 Model and Euler Equations

2.1 Model and basic assumptions

We consider a dynamic programing (DP) model in discrete time, with discrete actions, and mixed continuous/discrete state variables. Every period t, an agent takes a decision a_t to maximize his expected intertemporal payoff $\mathbb{E}_t[\sum_{j=0}^{T-t} \beta^j \ \Pi_t(a_{t+j}, \mathbf{s}_{t+j})]$, where $\beta \in (0,1)$ is the discount factor, T is the time horizon that can be finite or infinite, $\Pi_t(.)$ is the one-period payoff function at period t, and $\mathbf{s}_t \in \mathcal{S}$ is the vector of state variables at period t, which we assume follows a controlled Markov process with transition probability function $f_t(\mathbf{s}_{t+1} \mid a_t, \ \mathbf{s}_t)$. In this paper we consider discrete choice models such that the decision variable a_t belongs to the discrete and finite set $\mathcal{A} = \{0, 1, ..., J\}$. The sequence of value functions $\{V_t(.) : t \geq 1\}$ can be obtained recursively using the Bellman equation:

$$V_t(\mathbf{s}_t) = \max_{a_t \in \mathcal{A}} \left\{ \Pi_t(a_t, \mathbf{s}_t) + \beta \int V_{t+1}(\mathbf{s}_{t+1}) \ f_t(\mathbf{s}_{t+1} \mid a_t, \mathbf{s}_t) \ d\mathbf{s}_{t+1} \right\} \tag{1}$$

The optimal the decision rule, $\alpha_t(.) : \mathcal{S} \rightarrow \mathcal{A}$, is obtained as the arg-max of the expression in brackets. This model includes both stationary and non-stationary models. In the stationary case, the time horizon T is infinite, and payoff and transition probability functions are time-homogenous, and this implies that the value function and the optimal decision rule are also invariant over time.

Following the standard model in this literature (Rust, 1994), we distinguish two sets of state variables: $\mathbf{s}_t = (\mathbf{x}_t, \varepsilon_t)$, where $\mathbf{x}_t$ is the vector of state variables observable to the researcher, and

ε_t represents the unobservables for the researcher. The set of observable state variables $\mathbf{x}_t$ itself is comprised by two types of state variables, exogenous variables $\mathbf{z}_t$ and endogenous variables $\mathbf{y}_t$. They are distinguished by the fact that the transition probability of the endogenous variables depends on the action a_t, while the transition probability of the exogenous variables does not depend on a_t. The vector of unobservables satisfies the assumptions of additive separability (AS) and conditional independence (CI). The one-period payoff function is additively separable in the unobservables: $\Pi_t(a_t, \mathbf{s}_t) = \pi_t(a_t, \mathbf{x}_t) + \varepsilon_t(a_t)$, where $\varepsilon_t \equiv \{\varepsilon_t(a) : a \in A\}$ is a vector of unobservable random variables. And the transition probability (density) function of the state variables factors as: $f(\mathbf{s}_{t+1} \mid a_t, \mathbf{s}_t) = f_x(\mathbf{x}_{t+1} \mid a_t, \mathbf{x}_t)\ dG(\varepsilon_{t+1})$, where $G(.)$ is the CDF of ε_t which is absolutely continuous with respect to Lebesgue measure, strictly increasing and continuously differentiable in all its arguments, and with finite means.

Under these assumptions, the optimal decision rule of this DP problem can be represented using the *Conditional Choice Probability* (CCP) function $P_t(a \mid \mathbf{x})$, from $\mathcal{A} \times \mathcal{X} \to [0,1]$. This function represents the probability that given the observable state $\mathbf{x}$ the optimal decision at period t is a. Let $\mathbf{P}_t(\mathbf{x}_t)$ represent the vector with J free CCPs conditional on x_t. For notational simplicity, we occasionally use $\mathbf{P}_t$ to represent $\mathbf{P}_t(\mathbf{x}_t)$ below. Without loss of generality, the probability of choice alternative 0 is excluded from vector $\mathbf{P}_t$. Given a CCP vector $\mathbf{P}_t$, we can define the expected payoff function

$$\Pi_t^P(\mathbf{P}_t, \mathbf{x}_t) \equiv \sum\nolimits_{a=0}^{J} P_t(a \mid \mathbf{x}_t)\ [\pi_t(a, \mathbf{x}_t) + e_t(a, \mathbf{x}_t, \mathbf{P}_t)] \tag{2}$$

and the expected transition probability of the state variables

$$f^P(\mathbf{x}_{t+1} \mid \mathbf{P}_t, \mathbf{x}_t) \equiv \sum\nolimits_{a=0}^{J} P_t(a \mid \mathbf{x}_t)\ f(\mathbf{x}_{t+1} \mid \mathbf{P}_t, \mathbf{x}_t) \tag{3}$$

where $e_t(a, \mathbf{x}_t, \mathbf{P}_t)$ is the expected value of $\varepsilon_t(a)$ conditional on alternative a being chosen under decision rule $\mathbf{P}_t$. Using these definitions, the discrete-choice DP problem can be represented as a continuous-choice DP problem.

PROPOSITION 1. Consider the DP problem where the decision variable at period t is the CCP $\mathbf{P}_t$, the current payoff is $\Pi_t^P(\mathbf{P}_t, \mathbf{x}_t)$, and the transition probability of the state variables is $f^P(\mathbf{x}_{t+1} \mid \mathbf{P}_t, \mathbf{x}_t)$. By definition, the Bellman equation of this DP problem is:

$$V_t^P(\mathbf{x}_t) = \max_{\mathbf{P}_t \in [0,1]^J} \left\{ \Pi_t^P(\mathbf{P}_t, \mathbf{x}_t) + \beta \int V_{t+1}^P(\mathbf{x}_{t+1})\ f^P(\mathbf{x}_{t+1} | \mathbf{P}_t, \mathbf{x}_t)\ d\mathbf{x}_{t+1} \right\} \tag{4}$$

The value functions $V_t^P(.)$ and the optimal decision rule $P_t^*()$ that solve this continuous-choice DP problem are integrals, over the distribution of ε_t, of the value functions and optimal decision rules in the original discrete-choice DP problem, i.e., $V_t^P(\mathbf{x}_t) = \int V_t(\mathbf{x}_t, \varepsilon_t)\ dG(\varepsilon_t)$ and $P_t^*(a|\mathbf{x}_t) = \int 1\{\alpha_t(\mathbf{x}_t, \varepsilon_t) = a\}\ dG(\mathbf{x}_t, \varepsilon_t)$, where $1\{.\}$ is the indicator function.

Proof: In Aguirregabiria and Magesan (2013), Proposition 2(i).

Proposition 1 establishes a very useful representation property of this class of discrete choice models. The discrete-choice DP model has a representation as a continuous-choice DP problem. We show next that we can exploit this representation to derive Euler equations in similar way as in other continuous-choice DP models.

2.2 Euler Equations in single-agent discrete choice models

Given the CCPs at two consecutive periods, $\mathbf{P}_t$ and $\mathbf{P}_{t+1}$, define the two-period-forward transition probability function:

$$f_{(2)}^{P}(\mathbf{x}_{t+2} \mid \mathbf{P}_t, \mathbf{P}_{t+1}, \mathbf{x}_t) = \int f^{P}(\mathbf{x}_{t+2} \mid \mathbf{P}_{t+1}, \mathbf{x}_{t+1}) \; f^{P}(\mathbf{x}_{t+1} \mid \mathbf{P}_t, \mathbf{x}_t) \; d\mathbf{x}_{t+1} \tag{5}$$

The two-period-forward transition probability function is a convolution of the one-period transitions at periods t and $t+1$. Consider the following constrained optimization problem: choose the CCPs $\mathbf{P}_t$ and $\mathbf{P}_{t+1}$ to maximize the sum of the expected and discounted payoffs at periods t and $t+1$ subject to the constraint that the distribution of the state variables at period $t+2$ stays the same as under the optimal solution to the DP problem (4). This constrained optimization problem is formally given by:

$$\begin{array}{ll} \max\limits_{\{\mathbf{P}_t, \mathbf{P}_{t+1}\}} & \Pi_t^P(\mathbf{P}_t, \mathbf{x}_t) + \beta \int \Pi_{t+1}^P(\mathbf{P}_{t+1}, \mathbf{x}_{t+1}) \; f^P(\mathbf{x}_{t+1} | \mathbf{P}_t, \mathbf{x}_t) \; d\mathbf{x}_{t+1} \\ \text{subject to:} & f_{(2)}^{P}(\mathbf{x}_{t+2} \mid \mathbf{P}_t, \mathbf{P}_{t+1}, \mathbf{x}_t) = f_{(2)}^{P}(\mathbf{x}_{t+2} \mid \mathbf{P}_t^*, \mathbf{P}_{t+1}^*, \mathbf{x}_t) \text{ for any } \mathbf{x}_{t+2} \end{array} \tag{6}$$

By construction, the unique solution to this problem is given by the CCP functions $\mathbf{P}_t^*$ and $\mathbf{P}_{t+1}^*$ that solve the DP problem (4) at periods t and $t+1$.[5] Note that for each value of $\mathbf{x}_t$ there is a different constrained optimization problem, and therefore a different solution.

We can solve this problem using Lagrange method. There are two sets of Lagrange marginal conditions of optimality: conditions with respect to $\mathbf{P}_t$,

$$\frac{\partial \Pi_t^P}{\partial \mathbf{P}_t} + \beta \int \Pi_{t+1}^P(\mathbf{P}_{t+1}, \mathbf{x}_{t+1}) \frac{\partial f^P(\mathbf{x}_{t+1} | \mathbf{P}_t, \mathbf{x}_t)}{\partial \mathbf{P}_t} d\mathbf{x}_{t+1} - \int \lambda(\mathbf{x}_{t+2}) \frac{\partial f_{(2)}^P(\mathbf{x}_{t+2})}{\partial \mathbf{P}_t} d\mathbf{x}_{t+2} = 0, \tag{7}$$

and conditions with respect to $\mathbf{P}_{t+1}$,

$$\beta \int \frac{\partial \Pi_{t+1}^P}{\partial \mathbf{P}_{t+1}(\mathbf{x}_{t+1})} f^P(\mathbf{x}_{t+1} | \mathbf{P}_t, \mathbf{x}_t) d\mathbf{x}_{t+1} - \int \lambda(\mathbf{x}_{t+2}) \frac{\partial f_{(2)}^P(\mathbf{x}_{t+2})}{\partial \mathbf{P}_{t+1}(\mathbf{x}_{t+1})} d\mathbf{x}_{t+2} = 0, \tag{8}$$

where $\lambda(\mathbf{x}_{t+2})$ is the Lagrange multiplier associated to the constrain for state $\mathbf{x}_{t+2}$. We define as *Euler Equations* those equations that result from the combination of Lagrange conditions (7) and (8) and that do not include the Lagrange multipliers λ. Euler equations are marginal conditions of optimality for the DP problem that include payoffs and probabilities at two consecutive periods and, very importantly, do not include value functions or Lagrange multipliers. Not every DP

[5] Of course, if the DP problem is stationary, then this solution will be such that $\mathbf{P}_t^* = \mathbf{P}_{t+1}^*$.

problem admits a representation in terms Euler equations. For the rest of this subsection, we present conditions for the derivation of these equations in our model and a detailed description of the derivation.

First of all, it is important to distinguish between endogenous and exogenous state variables. The vector $\mathbf{x}_t$ has two subvectors, $\mathbf{y}_t$ and $\mathbf{z}_t$, such that $f_x(\mathbf{x}_{t+1}|a_t, \mathbf{x}_t) = f_y(\mathbf{y}_{t+1}|a_t, \mathbf{y}_t)\ f_z(\mathbf{z}_{t+1}|\mathbf{z}_t)$. The vector $\mathbf{y}_t$ contains endogenous state variables that evolve according to a *controlled* Markov process that depends on the action a_t. The vector $\mathbf{z}_t$ contains exogenous state variables that follow a stochastic process that does not depend on actions a_t or on the endogenous state variables. We do not impose any restriction on the vector of exogenous state variables, as it can include both continuous and discrete variables, nor do we place restrictions on the transition probability function f_z. However, the derivation of Euler equations requires conditions on the stochastic process of the endogenous state variables.

DEFINITIONS. (i) $\mathcal{Y}$ is the set of vectors in the support of the endogenous state vector $\mathbf{y}_t$. (ii) For a given vector $\mathbf{y}_t$, let $\mathcal{Y}_{(1)}(\mathbf{y}_t) \subseteq \mathcal{Y}$ be the set with all the vectors $\mathbf{y}_{t+1}$ with $f_y(\mathbf{y}_{t+1}|a_t, \mathbf{y}_t) > 0$ for some value of a_t. (iii) Similarly, for given $\mathbf{y}_t$, let $\mathcal{Y}_{(2)}(\mathbf{y}_t) \subseteq \mathcal{Y}$ be the set with all the vectors $\mathbf{y}_{t+2}$ with $\Pr(\mathbf{y}_{t+2}|a_t, a_{t+1}, \mathbf{y}_t) > 0$ for some value of a_t and a_{t+1}. (iv) $\widetilde{f}_y(\mathbf{y}_{t+1}|a, \mathbf{y}_t)$ is the "differential" transition probability $f_y(\mathbf{y}_{t+1}|a_t, \mathbf{y}_t) - f_y(\mathbf{y}_{t+1}|0, \mathbf{y}_t)$, where using choice alternative 0 as the reference or baseline is arbitrary. (v) $\widetilde{\mathbf{F}}(\mathbf{y}_t)$ is a matrix with elements $\widetilde{f}_y(\mathbf{y}_{t+2}|a, \mathbf{y}_{t+1})$ where the columns correspond to all the values $\mathbf{y}_{t+2} \in \mathcal{Y}_{(2)}(\mathbf{y}_t)$ leaving out one, and the rows correspond to all the values $(a, \mathbf{y}_{t+1}) \in [\mathcal{A} - \{0\}] \times \mathcal{Y}_{(1)}(\mathbf{y}_t)$.

The following assumption establishes sufficient conditions for the existence of Euler equations in our model.

ASSUMPTION EE. (1) The vector of endogenous state variables $\mathbf{y}_t$ has discrete (not necessarily finite) support $\mathcal{Y}$. (2) Matrix $\widetilde{\mathbf{F}}(\mathbf{y}_t)$ is full column rank.

As will become more clear below in the derivation of the Euler Equation representation, we leave out one element $\mathbf{y}_{t+2}$ from the set $\mathcal{Y}_{(2)}(\mathbf{y}_t)$ in the construction of the matrix $\widetilde{\mathbf{F}}(\mathbf{y}_t)$ because given (say) the first $|\mathcal{Y}_{(2)}(\mathbf{y}_t)| - 1$ restrictions on the two periods ahead transition function in problem 6, the last restriction is automatically satisfied, because the two periods ahead transition probabilities must sum to one. As such we only need $|\mathcal{Y}_{(2)}(\mathbf{y}_t)| - 1$ linearly independent columns.

We now draw on examples from commonly used models in empirical IO to provide some intuition for Assumption EE-2, before going on to show how it allows for the Euler Equation representation.

EXAMPLE 1 (Rust (1987) Bus Engine Replacement Problem). In this example, the endogenous state variable y_t is the mileage on the bus engine. Suppose that the space of possible mileages is given by the discrete set $\mathcal{Y} = \{0, 1, 2, ...\}$, and that mileage follows a transition rule $y_{t+1} = (1-a_t) * y_t$. The transition probability function is $f(y_{t+1}|1, y_t) = 1\{y_{t+1} = 0\}$ and $f(y_{t+1}|0, y_t) = 1\{y_{t+1} = y_t + 1\}$.

Then, for $y_t = y \geq 0$ we have that $\mathcal{Y}_{(1)}(y_t) = \{0, y+1\}$ and $\mathcal{Y}_{(2)}(y_t) = \{0, 1, y+2\}$. Without loss of generality, we leave out $y+2$ from the set $\mathcal{Y}_{(2)}(y_t)$. Define $\widetilde{f}(y_{t+2}|y_{t+1}) = f(y_{t+2}|1, y_{t+1}) - f(y_{t+2}|0, y_{t+1})$ such that in this model $\widetilde{f}(y_{t+2}|y_{t+1}) = 1\{y_{t+2} = 0\} - 1\{y_{t+2} = y_{t+1} + 1\}$ and the matrix $\widetilde{\mathbf{F}}(y)$ is given by:

$$\widetilde{\mathbf{F}}(y) \quad = \quad \begin{bmatrix} \widetilde{f}(y_{t+2} = 0 \mid y_{t+1} = 0) & \widetilde{f}(y_{t+2} = 1 \mid y_{t+1} = 0) \\ \widetilde{f}(y_{t+2} = 0 \mid y_{t+1} = y+1) & \widetilde{f}(y_{t+2} = 1 \mid y_{t+1} = y+1) \end{bmatrix} = \begin{bmatrix} 1 & -1 \\ 1 & 0 \end{bmatrix} \tag{9}$$

In this example, the matrix $\widetilde{\mathbf{F}}(\mathbf{y}_t)$ clearly has full column rank.

EXAMPLE 2 (Dynamic model of entry and exit). In this example, the endogenous state variable y_t is the incumbency status of the firm such that $\mathcal{Y} = \{0, 1\}$. The transition function is given by $y_{t+1} = a_t$ such that $f(y_{t+1}|a_t, y_t) = 1\{y_{t+1} = a_t\}$. Unlike the previous example, in this case the sets $\mathcal{Y}_{(1)}(y_t)$ and $\mathcal{Y}_{(2)}(y_t)$ are independent of y_t and $\mathcal{Y}_{(1)}(y_t) = \mathcal{Y}_{(2)}(y_t) = \mathcal{Y}$. Without loss, we leave out $y_{t+2} = 1$ from the set $\mathcal{Y}_{(2)}(y_t)$. The differential transition probability is $\widetilde{f}(y_{t+2}|y_{t+1}) = 1\{y_{t+2} = 1\} - 1\{y_{t+2} = 0\}$, and the transition matrix $\widetilde{\mathbf{F}}(y)$ is just the column vector $(\widetilde{f}(0|0), \widetilde{f}(0|1)) = (-1, -1)'$.

Note that when the set $\mathcal{Y}_{(2)}(y_t)$ has only two elements for any $y_t \in \mathcal{Y}$, Assumption EE-2 is automatically satisfied because $\widetilde{\mathbf{F}}(y)$ has only one column and it is different to zero. Clearly, if the matrix $\widetilde{\mathbf{F}}(\mathbf{y})$ has more columns than rows EE-2 can not be satisfied because the column rank is at most equal to the number of rows. As such, a necessary condition for EE-2 is that $J * |\mathcal{Y}_{(1)}(y_t)| \geq |\mathcal{Y}_{(2)}(y_t)| - 1$. In the case of a binary choice model, this condition is $|\mathcal{Y}_{(1)}(y_t)| \geq |\mathcal{Y}_{(2)}(y_t)| - 1$. What this means is that the set of possible values of the endogenous state variable two periods forward, $\mathcal{Y}_{(2)}(y_t)$, can at most have one more value than the set of possible values of y_t one period forward - the set of possible values of y can only increase by one each time period.

PROPOSITION 2. Under Assumption EE the marginal conditions for the maximization of the Lagrangian function in (6.) imply the following Euler equations. For every value of $\mathbf{x}_t$:

$$\frac{\partial \Pi_t^P(\mathbf{P}_t, \mathbf{x}_t)}{\partial \mathbf{P}_t(\mathbf{x}_t)} + \beta \sum_{\mathbf{x}_{t+1}} \left[\Pi_t^P(\mathbf{P}_{t+1}, \mathbf{x}_{t+1}) - \mathbf{m}(\mathbf{x}_{t+1})' \frac{\partial \mathbf{\Pi}_{t+1}^P(z_{t+1})}{\partial \mathbf{P}_{t+1}(z_{t+1})} \right] \widetilde{f}(y_{t+1}|a, x_t)\ f_z(z_{t+1}|z_t) \quad = \quad 0 \tag{10}$$

where $\partial \mathbf{\Pi}_{t+1}^P(z_{t+1}) / \partial \mathbf{P}_{t+1}(z_{t+1})$ *is a column vector with dimension* $J|\mathcal{Y}_{(1)}(y_t)| \times 1$ *that contains the partial derivatives* $\{\ \partial \Pi_{t+1}^P(y_{t+1}, z_{t+1}) /\ \partial P_{t+1}(a \mid y_{t+1}, z_{t+1})\ \}$ *for every action* $a > 0$ *and every value* $y_{t+1} \in \mathcal{Y}_{(1)}(y_t)$ *that can be reached from* y_t, *and fixed value for* z_{t+1}; *and* $\mathbf{m}(\mathbf{x}_{t+1})$ *is a* $J|\mathcal{Y}_{(1)}(y_t)| \times 1$ *vector such that* $\mathbf{m}(\mathbf{x}_{t+1}) \equiv \mathbf{f}_{t+1}^P{}' \ [\widetilde{\mathbf{F}}_{t+1}{}' \widetilde{\mathbf{F}}_{t+1}]^{-1}\ \widetilde{\mathbf{F}}_{t+1}{}'$, *where* $\mathbf{f}_{t+1}^e$ *is the vector of transition probabilities* $\{f^P(y_{t+2} \mid \mathbf{x}_{t+1}) : y_{t+2} \in \mathcal{Y}_{(2)}(y_t)\}$, *and* $\widetilde{\mathbf{F}}_{t+1}$ *is the matrix* $\widetilde{\mathbf{F}}(y_t)$ *of differential transition probabilities* $\widetilde{f}(y_{t+2}|a, \mathbf{x}_{t+1})$ *defined above.*

Proof: In Aguirregabiria and Magesan (2013), Proposition 3.

EXAMPLE 3 (Machine replacement model). Consider the model in Example 1. Taking into account the form of the matrix $\widetilde{\mathbf{F}}(y)$ in this model, it is possible to show that the Euler equation for this model is:

$$
\begin{array}{c}
\frac{\partial \Pi_t^P}{\partial P_t(1|\mathbf{x}_t)} + \beta\ \mathbb{E}_t \left[\Pi_{t+1}^P(1, z_{t+1}) - \Pi_{t+1}^P(y_t+1, z_{t+1}) \right] \\
+\beta\ \mathbb{E}_t \left[\frac{\partial \Pi_{t+1}^P(1,z_{t+1})}{\partial P_{t+1}(1|1,z_{t+1})} P_{t+1}(0|1, z_{t+1}) - \frac{\partial \Pi_{t+1}^P(y_t+1,z_{t+1})}{\partial P_{t+1}(1|y_t+1,z_{t+1})} P_{t+1}(0|y_t+1, z_{t+1}) \right] \quad = \quad 0
\end{array}
\tag{11}
$$

where we use $\mathbb{E}_t(.)$ to represent in a compact form the expectation over the distribution of $f_z(z_{t+1}|z_t)$. Suppose that the unobservables are i.i.d. extreme value distributed with dispersion parameter σ_ε. Then, the marginal expected profit $\partial \Pi_t^P / \partial P_t(1|\mathbf{x}_t)$ is equal to $\pi(1, \mathbf{x}_t) - \pi(0, \mathbf{x}_t) - \sigma_\varepsilon\ [\ln P_t(1|\mathbf{x}_t) - \ln P_t(0|\mathbf{x}_t)]$, and operating in the previous expression, it is possible to obtain the following Euler equation:

$$
\begin{array}{c}
\left[\pi(1, y_t, z_t) - \pi(0, y_t, z_t) - \sigma_\varepsilon\ \ln\left(\frac{P_t(1|y_t,z_t)}{1-P_t(1|y_t,z_t)} \right) \right] \quad + \\
\beta\ \mathbb{E}_t \left[\pi(1, 1, z_{t+1}) - \pi(1, y_t+1, z_{t+1}) - \sigma_\varepsilon\ \ln\left(\frac{P_{t+1}(1|1,z_{t+1})}{P_{t+1}(1|y_t+1,z_{t+1})} \right) \right] \quad = \quad 0
\end{array}
\tag{12}
$$

EXAMPLE 4. Consider binary choice model of entry/exit in Example 2. The Euler equation for this model is:

$$
\begin{array}{c}
\frac{\partial \Pi_t^P}{\partial P_t(1|\mathbf{x}_t)} + \beta\ \mathbb{E}_t \left(\Pi_{t+1}^P(1, z_{t+1}) - \Pi_{t+1}^P(0, z_{t+1}) \right) \quad + \\
\beta\ \mathbb{E}_t \left(\frac{\partial \Pi_{t+1}^P(0,z_{t+1},P_{t+1})}{\partial P_{t+1}(1|0,z_{t+1})} P_{t+1}(1|0, z_{t+1}) - \frac{\partial \Pi_{t+1}^P(1,z_{t+1},P_{t+1})}{\partial P_{t+1}(1|1,z_{t+1})} P_{t+1}(1|1, z_{t+1}) \right) \quad = \quad 0
\end{array}
\tag{13}
$$

Finally, for the logit version of this model the formula for this Euler equation is:

$$
\begin{array}{c}
\left[\pi(1, y_t, z_t) - \pi(0, y_t, z_t) - \sigma_\varepsilon\ \ln\left(\frac{P_t(1|y_t,z_t)}{1-P_t(1|y_t,z_t)} \right) \right] \quad + \\
\beta\ \mathbb{E}_t \left[\pi(1, 1, z_{t+1}) - \pi_t(1, 0, z_{t+1}) - \sigma_\varepsilon\ \ln\left(\frac{P_{t+1}(1|1,z_{t+1})}{P_{t+1}(1|0,z_{t+1})} \right) \right] \quad = \quad 0
\end{array}
\tag{14}
$$

2.3 Euler Equations in Dynamic Games

2.3.1 Continuous decision/state dynamic games

Consider a two player dynamic game. The two players are represented by the indexes i and j, respectively. The period payoff function of player i is $\pi_i(a_{it}, a_{jt}, y_{it})$, where π_i is a twice continuously differentiable real-valued function, variables $a_{it} \in \mathbb{R}$ and $a_{jt} \in \mathbb{R}$ represent the decisions of players i and j, and $y_{it} \in \mathbb{R}$ is the endogenous state variable for player i. For concreteness, we consider the case where $y_{it+1} = \delta_i\ y_{it} + a_{it}$, where $\delta_i \in (0.1)$ is a parameter. For instance, this framework can represent an investment game played between two firms in an oligopoly industry, where y_{it} represents a firm's stock of a certain capital equipment, a_{it} is the amount of investment (or disinvestment) at period t, and δ_i is the exogenous depreciation rate of capital. For the sake of notational simplicity, we omit here the exogenous state variables. The payoff relevant state variables of the model are (y_{it}, y_{jt}). A Markov Perfect Equilibrium (MPE) in this dynamic game can be characterized in terms of a pair of strategy functions $\{\alpha_i(y_{it}, y_{jt}), \alpha_j(y_{it}, y_{jt})\}$ such that α_i and α_j are functions from $\mathbb{R}^2$ into $\mathbb{R}^1$. Given an arbitrary strategy of player j, say α_j, let $V_i^\alpha(y_{it}, y_{jt})$ be player i's value function, i.e., the value associated to his best response. This value function is the solution to the following Bellman equation:

$$V_i^\alpha(y_{it}, y_{jt}) = \max_{a_{it}\in\mathbb{R}} \{\ \pi_i(a_{it}, \alpha_j(y_{it}, y_{jt}), y_{it}) + \beta\ V_i^\alpha(\delta_i\ y_{it} + a_{it},\ \delta_j\ y_{jt} + \alpha_j(y_{it}, y_{jt}))\ \} \tag{15}$$

For notational convenience, define $\pi_{it}^\alpha(a_{it}, y_{it}, y_{jt}) \equiv \pi_{it}(a_{it}, \alpha_j(y_{it}, y_{jt}), y_{it})$, such that $\dfrac{d\pi_{it}^\alpha}{dy_{it}} \equiv \dfrac{\partial \pi_{it}}{\partial y_{it}} + \dfrac{\partial \pi_{it}}{\partial a_{jt}}\ \dfrac{\partial \alpha_{jt}}{\partial y_{it}}$ and $\dfrac{d\pi_{it}^\alpha}{dy_{jt}} \equiv \dfrac{\partial \pi_{it}}{\partial a_{jt}}\ \dfrac{\partial \alpha_{jt}}{\partial y_{it}}$. For a given player, say i, we can obtain: the marginal condition of optimality with respect to a_{it}, in equation (16)(a); and the envelope conditions with respect to the endogenous state variables y_{it} and y_{jt}, in equations (16)(b) and (16)(c).

$$\begin{aligned}
&(a) & \frac{\partial \pi_{it}}{\partial a_{it}} + \beta\ \frac{\partial V_{it+1}^\alpha}{\partial y_{it+1}} &= 0 \\
&(b) & \frac{\partial V_{it}^\alpha}{\partial y_{it}} &= \frac{d\pi_{it}^\alpha}{dy_{it}} + \beta\ \frac{\partial V_{it+1}^\alpha}{\partial y_{it+1}}\ \delta_i + \beta\ \frac{\partial V_{it+1}^\alpha}{\partial y_{jt+1}}\ \frac{\partial \alpha_{jt}}{\partial y_{it}} \\
&(c) & \frac{\partial V_{it}^\alpha}{\partial y_{jt}} &= \frac{d\pi_{it}^\alpha}{dy_{jt}} + \beta\ \frac{\partial V_{it+1}^\alpha}{\partial y_{jt+1}} \left[\delta_j + \frac{\partial \alpha_{jt}}{\partial y_{jt}}\right]
\end{aligned} \tag{16}$$

We can combine these three conditions to obtain the following equation that does not involve the value function but only marginal period profits.

$$\begin{aligned}
&\frac{\partial \pi_{it}}{\partial a_{it}} + \beta\left[\frac{\partial \pi_{it+1}}{\partial a_{it+1}}\ \delta_i - \frac{d\pi_{it+1}^\alpha}{dy_{it+1}}\right] \quad - \\
&\beta\frac{\partial \alpha_{jt+1}}{\partial y_{it+1}}\left[r_{jt+2}^\alpha\ \left(\frac{\partial \pi_{it+1}}{\partial a_{it+1}} + \beta\left[\frac{\partial \pi_{it+2}}{\partial a_{it+2}}\ \delta_i - \frac{d\pi_{it+2}^\alpha}{dy_{it+2}}\right]\right) - \beta\ \frac{d\pi_{it+2}^\alpha}{dy_{jt+2}}\right] \quad = \quad 0
\end{aligned} \tag{17}$$

where $r_{jt+2}^\alpha \equiv [\delta_j + \partial\alpha_{jt+2}/\partial y_{jt+2}]\,/\partial\alpha_{jt+2}/\partial y_{it+2}$. We denote this equation the Euler equation of player i in this dynamic game. Note that when $\partial\alpha_{jt+1}/\partial y_{it+1} = 0$, there are not dynamic strategic interactions between players and the Euler equation becomes the standard one in a single-agent model, i.e., $\dfrac{\partial \pi_{it}}{\partial a_{it}} + \beta\left[\dfrac{\partial \pi_{it+1}}{\partial a_{it+1}}\ \delta_i - \dfrac{\partial \pi_{it+1}}{\partial y_{it+1}}\right] = 0$. In contrast to the single-agent case where the EE involves marginal payoffs at two consecutive periods, this EE involves marginal payoffs at three periods.

2.3.2 Discrete Choice/State Dynamic Games

Consider now a binary choice version of the previous two player game. The period payoff function of player i is $\pi_i(a_{it}, a_{jt}, y_{it}) + \varepsilon_{it}(a_{it})$, where $a_{it} \in \{0, 1\}$ represents the decision of player i, and

$y_{it+1} = a_{it}$, i.e., a market entry game. $\varepsilon_{it}(0)$ and $\varepsilon_{it}(1)$ are private information state variables that are i.i.d. over players and over time. A MPE in this dynamic game can characterized in terms of a pair of vectors of probabilities $\mathbf{P}_i \equiv \{P_i(\mathbf{y}_t) : \mathbf{y}_t \in \{0,1\}^2\}$ and $\mathbf{P}_j \equiv \{P_j(\mathbf{y}_t) : \mathbf{y}_t \in \{0,1\}^2\}$. Let $\Pi_i^P(\mathbf{P}_i, \mathbf{P}_j, \mathbf{y}_t)$ be the expected payoff of player i (before the realization of ε's) if players behave according to their CCPs $(\mathbf{P}_i, \mathbf{P}_j)$. Given an arbitrary strategy of player j, say $\mathbf{P}_j$, let $V_i^P(\mathbf{y}_t)$ be player i's value function, i.e., the value associated to his best response. This value function is the solution to the following Bellman equation:

$$V_i^P(\mathbf{y}_t) = \max_{\mathbf{P}_{it}} \left\{ \Pi_i^P\left(\mathbf{P}_{it}, \mathbf{P}_j, \mathbf{y}_t\right) + \beta\, \mathbb{E}\left[V_i^P\left(\mathbf{y}_{t+1}\right) \mid \mathbf{P}_{it}, \mathbf{P}_j, \mathbf{y}_t \right] \right\} \tag{18}$$

where

$$\begin{aligned} \mathbb{E}\left[V_i^P\left(\mathbf{y}_{t+1}\right) \mid \mathbf{P}_{it}, \mathbf{P}_j, \mathbf{y}_t \right] &= (1 - P_{it}(\mathbf{y}_t)) \left[(1 - P_j(\mathbf{y}_t))\, V_i^P(0,0) + P_j(\mathbf{y}_t)\, V_i^P(0,1) \right] \\ &+ P_{it}(\mathbf{y}_t)) \left[(1 - P_j(\mathbf{y}_t))\, V_i^P(1,0) + P_j(\mathbf{y}_t)\, V_i^P(1,1) \right] \end{aligned} \tag{19}$$

As in the single-agent discrete choice/state case, the derivation of Euler equations in a dynamic game with discrete choice/state is based on the definition of a constrained optimization problem. The key features of this constrained optimization problem are: (1) the objective function is the expected and discounted payoff during three consecutive periods; (2) the constraint establishes that the distribution of the endogenous state variables at the end of these three periods should be same as in the equilibrium of the dynamic game; and (3) the choice variables are player i's CCPs $\mathbf{P}_{it}$, $\mathbf{P}_{it+1}$, and $\mathbf{P}_{it+2}$. Formally, this constrained optimization problem is (omitting $\mathbf{P}_j$ as an argument for notational simplicity):

$$\max_{\{\mathbf{P}_{it}, \mathbf{P}_{it+1}, \mathbf{P}_{it+2}\}} \quad \Pi_i^P\left(\mathbf{P}_{it}, \mathbf{y}_t\right) + \beta \mathbb{E}\left[\Pi_i^P\left(\mathbf{P}_{it+1}, \mathbf{y}_{t+1}\right) | \mathbf{P}_{it}, \mathbf{y}_t \right] + \beta^2 \mathbb{E}\left[\Pi_i^P\left(\mathbf{P}_{it+2}, \mathbf{y}_{t+2}\right) | \mathbf{P}_{it}, \mathbf{P}_{it+1}, \mathbf{y}_t \right]$$

$$\text{subject to:} \qquad f_{(3)}(\mathbf{y}_{t+3} \mid \mathbf{P}_{it}, \mathbf{P}_{it+1}, \mathbf{P}_{it+2}, \mathbf{y}_t) = f_{(3)}(\mathbf{y}_{t+3} \mid \mathbf{P}_i^*, \mathbf{P}_i^*, \mathbf{P}_i^*, \mathbf{y}_t) \text{ for any } \mathbf{y}_{t+3}$$

where $f_{(3)}(\mathbf{y}_{t+3} \mid \mathbf{P}_{it}, \mathbf{P}_{it+1}, \mathbf{P}_{it+2}, \mathbf{y}_t)$ is the distribution of $\mathbf{y}_{t+3}$ conditional on $\mathbf{y}_t$ and induced by the CCPs $\mathbf{P}_{it}$, $\mathbf{P}_{it+1}$, $\mathbf{P}_{it+2}$ and $\mathbf{P}_j$. This three-period-forward transition probability function is a convolution of the one-period transitions at periods t, $t+1$, and $t+2$.

$$f_{(3)}(\mathbf{y}_{t+3} | \mathbf{P}_{it}, \mathbf{P}_{it+1}, \mathbf{P}_{it+2}, \mathbf{y}_t) \equiv \sum_{\mathbf{y}_{t+2}} f_{(1)}\left(\mathbf{y}_{t+3} | \mathbf{P}_{it+2}, \mathbf{y}_{t+2}\right) \left[\sum_{\mathbf{y}_{t+1}} f_{(1)}\left(\mathbf{y}_{t+2} | \mathbf{P}_{it+1}, \mathbf{y}_{t+1}\right) f_{(1)}\left(\mathbf{y}_{t+1} | \mathbf{P}_{it}, \mathbf{y}_t\right) \right]$$

with

$$f_{(1)}(\mathbf{y}' | \mathbf{P}_i, \mathbf{y}) \equiv (1 - P_i(\mathbf{y}))^{1-y_i'}\, P_i(\mathbf{y})^{y_i'}\, (1 - P_j(\mathbf{y}))^{1-y_j'}\, P_j(\mathbf{y})^{y_j'}$$

By construction, the unique solution to this problem is $\mathbf{P}_{it} = \mathbf{P}_{it+1} = \mathbf{P}_{it+2} = \mathbf{P}_i^*$, where $\mathbf{P}_i^*$ is the best response probability of player i in the infinite horizon dynamic game. The Euler equation for player i is the result of combining three sets of Lagrange marginal conditions of optimality: the marginal conditions with respect to (a) $\mathbf{P}_{it}$, (b) $\mathbf{P}_{it+1}$, and (c) $\mathbf{P}_{it+2}$.

2.4 Relationship between Euler Equations and other CCP representations

Our derivation of Euler equations for DDC models above is related to previous work by Hotz and Miller (1993), Aguirregabiria and Mira (2002), and Arcidiacono and Miller (2011). These papers derive representations of optimal decision rules using CCPs and show how these representations can be applied to estimate DDC models using simple two-step methods that provide substantial computational savings relative to full-solution methods. In these previous papers, we can distinguish three different types of CCP representations of optimal decisions rules: (1) the *present-value representation*; (2) the *terminal-state* representation; and (3) the *finite-dependence* representation.

The *present-value representation* consists of using CCPs to obtain an expression for the expected and discounted stream of future payoffs associated with each choice alternative. In general, given CCPs, the valuation function can be obtained recursively using its definition. And given this valuation function we can construct the agent's optimal decision rule (or best response) at period t given that he believes that in the future he will behave according to the CCPs in the vector $\mathbf{P}$. This *present-value representation* is the CCP approach more commonly used in empirical applications because it can be applied to a general class of dynamic discrete choice models. However, this representation requires the computation of present values and therefore it is subject to the curse of dimensionality. In applications with large state spaces, this approach can be implemented only if it is combined with an approximation method such as the discretization of the state space, or Monte Carlo simulation (e.g., Hotz et al, 1995, and Bajari et al., 2007). In general, these approximation methods introduce a bias in parameter estimates.

The *terminal-state* representation was introduced by Hotz and Miller (1993) and it applies only to optimal stopping problems with a terminal state. The *finite-dependence* representation was introduced by Arcidiacono and Miller (2011) and applies to a particular class of DDC models with the finite dependence property. A DDC model has the finite dependence property if given two values of the decision variable at period t and their respective paths of the state variables after this period, there is always a finite period $t' > t$ (with probability one) where the state variables in the two paths take the same value. The *terminal-state* and the *finite-dependence* CCP representations do not involve the computation of present values, or even the estimation of CCPs at every possible state. This implies substantial computational savings as well as avoiding biases induced by approximation errors.

The system of *Euler-equations* that we have derived in Proposition 3 can be seen also a CCP representation of the optimal decision rule in a DDC model. Our representation shares all the computational advantages of the *terminal-state* and *finite-dependence* representations. However, in contrast to the *terminal-state* and *finite-dependence*, our Euler equation representation applies to a general class of DDC models. We can derive Euler equations for any DDC model where the unobservables satisfy the conditions of additive separability in the payoff function, and conditional independence in the transition of the state variables.

3 Policy Iteration mappings

Consider the Euler equation in the machine replacement model of Examples 1 and 3:

$$\begin{aligned}\left[\pi\left(1, y_t, z_t\right) - \pi\left(0, y_t, z_t\right) - \sigma_\varepsilon \ \ln\left(\frac{P_t(1|y_t,z_t)}{1-P_t(1|y_t,z_t)}\right)\right] \quad &+ \\ \beta \ \mathbb{E}_t\left[\pi\left(1, 1, z_{t+1}\right) - \pi\left(1, y_t+1, z_{t+1}\right) - \sigma_\varepsilon \ \ln\left(\frac{P_{t+1}(1|1,z_{t+1})}{P_{t+1}(1|y_t+1,z_{t+1})}\right)\right] \quad &= \quad 0\end{aligned} \tag{20}$$

Solving for $P_t(1|y_t, z_t)$ we can get the following expression:

$$P_t(1|y_t, z_t) = \frac{\exp\left\{\frac{\pi(1,y_t,z_t)-\pi(0,y_t,z_t)}{\sigma_\varepsilon} + \beta \ \mathbb{E}_t\left[\frac{\pi(1,1,z_{t+1})-\pi(1,y_t+1,z_{t+1})}{\sigma_\varepsilon} - \ln\left(\frac{P_{t+1}(1|1,z_{t+1})}{P_{t+1}(1|y_t+1,z_{t+1})}\right)\right]\right\}}{1+\exp\left\{\frac{\pi(1,y_t,z_t)-\pi(0,y_t,z_t)}{\sigma_\varepsilon} + \beta \ \mathbb{E}_t\left[\frac{\pi(1,1,z_{t+1})-\pi(1,y_t+1,z_{t+1})}{\sigma_\varepsilon} - \ln\left(\frac{P_{t+1}(1|1,z_{t+1})}{P_{t+1}(1|y_t+1,z_{t+1})}\right)\right]\right\}} \tag{21}$$

The right-hand-side of this equation describes a function $\Lambda(\mathbf{x}_t, \mathbf{P}_{t+1})$ from the vector of CCPs at period $t+1$, $\mathbf{P}_{t+1}$, into the probability space, such that $\Lambda(\mathbf{x}_t, .) : [0,1]^{|\mathcal{X}|} \to [0,1]$. Let $\mathbf{\Lambda}(\mathbf{P})$ be the vector-valued function that consists of the collection of the functions $\Lambda(\mathbf{x}_t, .)$ for every value of $\mathbf{x}_t$ in the state space $\mathcal{X}$: $\mathbf{\Lambda}(\mathbf{P}) \equiv \{\Lambda(\mathbf{x}_t, \mathbf{P}) : \text{for any } \mathbf{x}_t \in \mathcal{X}\}$. By definition, $\mathbf{\Lambda}(\mathbf{P})$ is a fixed point mapping in the probability space $[0,1]^{|\mathcal{X}|}$. Using this mapping we can represent in the following vector form the relationship between CCPs at periods t and $t+1$ implied by the EE as follows:

$$\mathbf{P}_t = \mathbf{\Lambda}(\mathbf{P}_{t+1}) \tag{22}$$

In the stationary version of the model (i.e., infinite horizon and time-homogeneous payoff function), the optimal CCPs are time invariant: $\mathbf{P}_t = \mathbf{P}_{t+1}$. Therefore, expression (22) describes the optimal CCPs as a fixed point of the mapping $\mathbf{\Lambda}$. We denote $\mathbf{\Lambda}$ as the *Euler Equation - Policy Iteration mapping (EE-PI) mapping.*

PROPOSITION 3. The EE-PI is a contraction mapping. The optimal vector of CCPs that solves the dynamic decision problem, $\mathbf{P}^*$*, is the only fixed point of the EE-PI mapping.*

Proof: In the Appendix.

A corollary of Proposition 3 is that successive iterations in the EE-PI mapping is a method to solve this discrete choice dynamic programming problem. Below we compare this method to the most commonly used methods of successive approximation to the value function (i.e., iterations in the Bellman equation), and iteration in the standard best respond mapping.

4 Estimation of Counterfactuals

4.1 GMM estimation of structural parameters

Suppose that the researcher's dataset consists of panel data of N agents over T periods of time with information on agents' actions and state variables, $\{ a_{it} , \mathbf{x}_{it} : i = 1, 2, ..., N ; t = 1, 2, ..., T\}$. Here we consider the typical sample in a single-agent model where the number of agents N is large, and the number of time periods is typically short. The researcher is interested in using this sample to estimate the structural parameters in preferences, transition probabilities, and the discount factor. Let θ be the vector of structural parameters. We describe here the GMM estimation of these structural parameters using moment restrictions from the Euler equations derived above.

The Euler equations that we have derived imply the following orthogonality conditions: $\mathbb{E}(\xi(a_{it}, \mathbf{x}_{it}, \mathbf{x}_{it+1}; P_{it}, P_{it+1}, \theta) \mid a_{it}, \mathbf{x}_{it}) = 0$, where

$$\xi(a_{it}, \mathbf{x}_{it}, \mathbf{x}_{it+1}; P_{it}, P_{it+1}, \theta) \equiv \frac{\partial \Pi_t^P}{\partial P_t(a_{it}|\mathbf{x}_{it})} + \beta \left[\Pi_{t+1}^P(\mathbf{x}_{it+1}) - \mathbf{m}(\mathbf{x}_{it+1})' \frac{\partial \mathbf{\Pi}_{t+1}^P(z_{it+1})}{\partial \mathbf{P}_{t+1}(z_{it+1})} \right] \frac{\widetilde{f}(y_{it+1}|a_{it}, \mathbf{x}_{it})}{f(y_{it+1}|a_{it}, \mathbf{x}_{it})} \tag{23}$$

Note that this orthogonality condition comes from the Euler equation in Proposition 2, but we have made two changes. First, we have included the expectation $\mathbb{E}(.|a_{it}, \mathbf{x}_{it})$ that replaces the sum $\sum_{\mathbf{x}_{it+1}}$ and the distribution of $\mathbf{x}_{it+1}$ conditional on $(a_{it}, \mathbf{x}_{it})$, i.e., $f(y_{it+1}|a_{it}, \mathbf{x}_{it})$ $f_z(z_{it+1}|z_{it})$. And second, the Euler equation applies to any hypothetical choice, a, at period t, but in the orthogonality condition $\mathbb{E}(\xi(a_{it}, \mathbf{x}_{it}, \mathbf{x}_{it+1}; P_{it}, P_{it+1}, \theta) \mid a_{it}, \mathbf{x}_{it}) = 0$, we consider only the actual/observed choice a_{it}.

Given these conditions, we can construct a consistent an asymptotically normal estimator of θ using a semiparametric two-step GMM. The first step consists in the nonparametric estimation of the CCPs $P_t(a|\mathbf{x}) \equiv \Pr(a_{it} = a \mid \mathbf{x}_{it} = \mathbf{x})$. Let $\widehat{\mathbf{P}}_{t,N} \equiv \{\widehat{P}_t(a|\mathbf{x})\}$ be a vector of nonparametric estimates of CCPs for any choice alternative a and any value of $\mathbf{x}$ observed in the sample. For instance, $\widehat{P}_t(a|\mathbf{x})$ can be a kernel (Nadaraya-Watson) estimator of the regression between $1\{a_i = a\}$ and $\mathbf{x}_{it}$. Note that we do not need to estimate CCPs at states which are not observed in the sample. For simplicity, suppose that the sample includes only two periods, t and $t+1$. Let $\widehat{\mathbf{P}}_{t,N}$ and $\widehat{\mathbf{P}}_{t+1,N}$ be vectors with the nonparametric estimates. In the second step, the GMM estimator of θ is:

$$\hat{\theta}_N = \arg\min_{\theta \in \Theta} \; m_N' \left(\theta, \widehat{\mathbf{P}}_{t,N}, \widehat{\mathbf{P}}_{t+1,N} \right) \; \Omega_N \; m_N \left(\theta, \widehat{\mathbf{P}}_{t,N}, \widehat{\mathbf{P}}_{t+1,N} \right) \tag{24}$$

where $m_N(\theta, \mathbf{P}_t, \mathbf{P}_{t+1})$ is the vector of sample moments:

$$m_N(\theta, \mathbf{P}_t, \mathbf{P}_{t+1}) = \frac{1}{N} \sum_{i=1}^{N} Z(a_{it}, \mathbf{x}_{it}) \; \xi(a_{it}, \mathbf{x}_{it}, \mathbf{x}_{it+1}; P_{it}, P_{it+1}, \theta) \tag{25}$$

$Z(a_{it}, \mathbf{x}_{it})$ is a vector of instruments, i.e., known functions of the observable decision and state variables at period t. As in the case of the static ARUM, this semiparametric two-step GMM estimator is consistent and asymptotically normal under mild regularity conditions. This two-step semiparametric estimator is root-N consistent and asymptotically normal under mild regularity conditions (see Theorems 8.1 and 8.2 in Newey and McFadden, 1994). The variance matrix of this estimator can be estimated using the semiparametric method in Newey (1994), or as recently shown by Ackerberg, Chen, and Hahn (2012) using a computationally simpler parametric-like method as in Newey (1984).

The estimation based on the moment conditions provided by the *Euler-equation*, or *terminal-state*, or *finite-dependence* representations imply an efficiency loss relative to estimation based on *present-value representation.* As shown by Aguirregabiria and Mira (2002, Proposition 4), the two-step pseudo maximum likelihood estimator based on the CCP *present-value* representation is asymptotically efficient (equivalent to the maximum likelihood estimator). This efficiency property is not shared by the other CCP representations. Therefore, there is a trade-off in the choice between CCP estimators based on Euler equations and on present-value representations. The present value representation is the best choice in models that do not require approximation methods. However, in models with large state spaces that require approximation methods, the Euler equations CCP estimator can provide more accurate estimates.

4.2 Pseudo Maximum Likelihood estimator based on EE-PI mapping

Let $\mathcal{Z}_N$ be the set with all the values of the vector of exogenous state variables z that we observe in the sample: $\mathcal{Z}_N \equiv \{z_{it}, z_{it+1} : i = 1, 2, ..., N\}$. Given $\mathcal{Z}_N$, let $\mathbf{P}_N$ be the vector of CCPs at every possible action and endogenous state variable but only at the values of the exogenous state variables observed in the sample: $\mathbf{P}_N \equiv \{P(a|y, z) : (a, y, x) \in \mathcal{A} \times \mathcal{Y} \times \mathcal{Z}_N\}$. We define a sample-based version of EE-PI mapping that is defined on $\mathbf{P}_N$. For our example,

$$\Lambda_N(y_t, z_t; \mathbf{P}; \theta) = \frac{\exp\left\{\frac{\pi(1,y_t,z_t;\theta)-\pi(0,y_t,z_t;\theta)}{\sigma_\varepsilon} + \beta\mathbb{E}_{N,z_t}\left[\frac{\pi(1,1,z_{t+1};\theta)-\pi(1,y_t+1,z_{t+1};\theta)}{\sigma_\varepsilon} - \ln\left(\frac{P(1|1,z_{t+1})}{P(1|y_t+1,z_{t+1})}\right)\right]\right\}}{1+\exp\left\{\frac{\pi(1,y_t,z_t;\theta)-\pi(0,y_t,z_t;\theta)}{\sigma_\varepsilon} + \beta\mathbb{E}_{N,z_t}\left[\frac{\pi(1,1,z_{t+1};\theta)-\pi(1,y_t+1,z_{t+1};\theta)}{\sigma_\varepsilon} - \ln\left(\frac{P(1|1,z_{t+1})}{P(1|y_t+1,z_{t+1})}\right)\right]\right\}} \tag{26}$$

where $\mathbb{E}_{N,z}(.)$ represents the sample counterpart of the conditional expectation over the distribution of z_{t+1} conditional on $z_t = z$, i.e., for an arbitrary function $h(z_{t+1})$, $\mathbb{E}_{N,z}(h(z_{t+1})) = \sum_{i=1}^{N} w_N(z_{it}-z)$ $h(z_{it+1})$, where $w_N(z_{it}-z)$ is a weighting function such as $w_N(z_{it}-z) = 1\{z_{it}-z = 0\}/\sum_{j=1}^{N} 1\{z_{jt}-z = 0\}$ (i.e., frequency estimator), or $w_N(z_{it} - z) = K\left(\frac{z_{it}-z}{b_N}\right)/\sum_{j=1}^{N} K\left(\frac{z_{jt}-z}{b_N}\right)$ (i.e., Kernel estimator). Given $\Lambda_N(y_t, z_t; \mathbf{P}; \theta)$, we can define the *sample-based EE-PI mapping* as:

$$\mathbf{\Lambda}_N(\mathbf{P}; \theta) \equiv \Lambda_N(a, y, z; \mathbf{P}; \theta) : (a, y, x) \in \mathcal{A} \times \mathcal{Y} \times \mathcal{Z}_N\} \tag{27}$$

Note the dimension of the mapping $\mathbf{\Lambda}_N$ is $J * N * |\mathcal{Y}|$ that can be several orders of magnitude smaller than the dimension of $\mathbf{\Lambda}$ is the dimension of $\mathcal{Z}$ is large relative to N.

PROPOSITION 3. The sample-based EE-PI mapping $\mathbf{\Lambda}_N$ is a contraction mapping, and it converges uniformly in $(\mathbf{P};\theta)$ to the true EE-PI mapping $\mathbf{\Lambda}$.

$$\mathbf{\Lambda}_N(\mathbf{P};\theta) \rightarrow_{p-uniformly} \mathbf{\Lambda}(\mathbf{P};\theta)$$

Proof: In the Appendix.

We can define a two-step Pseudo Maximum Likelihood (PML) estimator of θ using this sample-based mapping. For arbitrary $(\mathbf{P}_N;\theta)$, define the pseudo log-likelihood function:

$$Q_N(\mathbf{P}_N;\theta) = \sum_{i=1}^{N} \ln \Lambda_N(a_{it}, y_{it}, z_{it}; \mathbf{P}_N; \theta) \tag{28}$$

Given a nonparametric estimator of $\mathbf{P}_N$, say $\widehat{\mathbf{P}}_N$, the two-step PML is the value of θ that maximizes the pseudo likelihood, i.e.,$\widehat{\theta}_N = \arg\max_\theta Q_N(\widehat{\mathbf{P}}_N;\theta)$.

4.3 Estimation of counterfactuals

Given the sample and the estimator of the structural parameters $\hat{\theta}_N$, the researcher is interested in estimating the vector of agents's CCPs if we perturb the structural parameters from $\hat{\theta}_N$ to an alternative vector θ^*. Let $\mathbf{P}^* \equiv \mathbf{P}(\theta^*)$ be this vector. For single-agent models, this vector is defined as the unique fixed point in $\mathbf{P}$ of the contraction mapping $\mathbf{\Lambda}$:, i.e., $\mathbf{P}(\theta^*) = \mathbf{\Lambda}(\mathbf{P}(\theta^*), \theta^*)$.

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 $\mathbf{P}(\theta^*)$ at every possible state is computationally unfeasible. Here we consider the estimation of $\mathbf{P}(\theta^*)$ at the subset of state points that come from the sample. More specifically, let $\mathcal{Z}_N$ be the set with all the values of the vector of exogenous state variables z that we observe in the sample: $\mathcal{Z}_N \equiv \{z_{it}, z_{it+1} : i = 1, 2, ..., N\}$. The researcher is interested in estimating the vector of $2N$ CCPs:

$$\mathbf{P}^*_{\mathcal{Z}_N} = \mathbf{P}_{\mathcal{Z}_N}(\theta^*) \equiv \{P(a|y, z; \theta^*) : (a, y, x) \in \mathcal{A} \times \mathcal{Y} \times \mathcal{Z}_N\} \tag{29}$$

Note that $\mathbf{P}^*_{\mathcal{Z}_N}$ is a subset of $\mathbf{P}^*$. Also, note that $\mathbf{P}^*_{\mathcal{Z}_N}$ includes all the possible values of the endogenous decision and state variables but only only sample observations for the exogenous state variables. The implicit assumption is that the counterfactual experiment from $\hat{\theta}_N$ to θ^* does not involve any change in the stochastic process of the exogenous state variables f_z.

Here we show how we can use Euler equations to estimate the counterfactual $\mathbf{P}^*_{\mathcal{Z}_N}$. The vector of counterfactual probabilities $\mathbf{P}^*_{\mathcal{Z}_N}$ satisfies the orthogonality conditions from the Euler equations. Our estimator of $\mathbf{P}^*_{\mathcal{Z}_N}$ is the unique fixed point of the empirical EE-PI mapping $\mathbf{\Lambda}_N(\mathbf{P};\theta^*)$, i.e., the estimator $\widehat{\mathbf{P}^*_{\mathcal{Z}_N}}$ is implicitly defined as:

$$\widehat{\mathbf{P}^*_{\mathcal{Z}_N}} = \mathbf{\Lambda}_N\left(\widehat{\mathbf{P}^*_{\mathcal{Z}_N}}; \theta^*\right) \tag{30}$$

In single-agent models, this condition uniquely identifies the counterfactual $\mathbf{P}^*_{\mathcal{Z}_N}$. For the derivation of the statistical properties $\widehat{\mathbf{P}^*_{\mathcal{Z}_N}}$. An important property of this estimator is that it is consistent regardless the dimension of $\mathcal{Z}$. This property is not shared by other estimators based on other approximations of the dynamic decision problem.

PROPOSITION 4. The vector of CCPs $\widehat{\mathbf{P}^*_{\mathcal{Z}_N}}$ implicitly defined as the fixed point $\widehat{\mathbf{P}^*_{\mathcal{Z}_N}} = \mathbf{\Lambda}_N\left(\widehat{\mathbf{P}^*_{\mathcal{Z}_N}}; \theta^*\right)$ is a root-N consistent and asymptotically normal estimator of $\mathbf{P}^*_{\mathcal{Z}_N}$.

We can use a policy iteration algorithm in the mapping $\mathbf{\Lambda}_N(.; \theta^*)$ to compute the counterfactual experiment.[6] This method implies computational savings with respect the standard approaches of value function iteration or iterations in the standard best response mapping. First, note that we do not have to compute infinite-period-forward present values. An evaluation of the standard PI mapping involves two steps: valuation iteration, i.e., given $\mathbf{P}$ we calculate the present values of behaving according to these probabilities forever in the future; and policy improvement iteration, i.e., given these present values, we calculate the best response today to these future valuations. The EE-PI mapping has also this two steps, but the "valuation step" is relatively trivial because it only involves the expectation of one-period forward payoffs. This implies very substantial computational savings. Second, in contrast to our method, other common methods suffer of an approximation bias: an asymptotic bias that does not go zero as the sample size increases and the degree of approximation (e.g., number of Monte Carlo simulations) stays constant.

5 Monte Carlo Experiments

In this section we consider Monte Carlo experiments to study the performance of the *EE-Policy-Iteration* method relative to standard policy iteration methods in the context of a simple dynamic model of entry and exit. We study the performance of the methods in two distinct dimensions. First we examine the differences in the computational burdens that the two policy iteration methods impose, as well as the sources of the differences. The standard policy iteration mapping is a composite mapping, as the policies are expressed in terms of value functions, which are themselves expressed in terms of the policies. The *EE-Policy-Iteration Mapping*, by contrast, is not a composite mapping. We would expect then that the steps to convergence using the standard mapping are larger, meaning it requires fewer iterations than the *EE-Policy-Iteration mapping*, but each step is potentially much more costly, making the relative total time to convergence ambiguous in principle. We use the experiments to study the time per iteration and number of iterations each method takes to convergence to get a better understanding of the computational costs. Second, we use the experiments to evaluate the finite sample properties of the estimators associated with each method.

[6] To save iterations, it seems natural to start iterating at the estimates of the CCPs under thet estimated values of $\hat{\theta}_N$.

The two policy iteration methods offer an interesting bias-variance trade off. The application of the standard mapping typically requires the researcher to use approximation methods (i.e., simulation or state-space reduction) which imply estimation errors that are potentially large. These errors enter in a complicated, non-linear fashion into the probability estimates, and can lead to potentially significant bias. The *EE-Policy-Iteration mapping* is not subject to this bias, but, as the mapping is derived from Euler Equations, it does not exploit all the restrictions that the model places on the data. In this way the estimates associated with the *EE-Policy-Iteration mapping* are prone to being higher variance than those of the standard mapping. We estimate counterfactual probabilities using each method for a large number of simulated datasets, and study this trade-off.

5.1 Design

We consider a simple dynamic model of entry/exit. Let the action space be given by $\mathcal{A} = \{0, 1\}$, where a_t is the indicator of being active in a market or in some particular activity. The endogenous state variable y_t is the lagged value of the decision variable, $y_t = a_{t-1}$, and it represents whether the agent was active at previous period. The vector of *observable* state variables is given by $\mathbf{x}_t = (y_t, \mathbf{z}_t)$ where $\mathbf{z}_t$ is a vector of exogenous state variables. The vector $\mathbf{z}_t$ is itself comprised by several exogenous state variables, including the firm productivity[7], and market and firm characteristics that affect variable profit, fixed cost, or/and entry costs. We specify each of these components in turn.

We assume the following form of the variable profit function: $VP_t = p_t \exp(\lambda\ \omega_t)$, where p_t is the market price p_t and ω_t is a firm productivity shock that varies across firms in the same market. The market price p_t has the following form: $p_t = \alpha_0^{VP} + \alpha_1^{VP} z_{1t} + \alpha_2^{VP} z_{2t}$, where z_{1t} and z_{2t} are exogenous state variables (e.g., z_{1t} is market size and z_{2t} can be interpreted as measuring the competition the firm faces), and α_1^{VP} and α_2^{VP} are parameters. We assume the following form for fixed cost, $FC_t = \alpha_0^{FC} + \alpha_1^{FC} z_{3t}$, and for the entry cost, $EC_t = \alpha_0^{EC} + \alpha_1^{EC} z_{4t}$, where z_{3t} and z_{4t} are exogenous state variables, and α's are parameters. Only firms who were not active last period (i.e., $y_t = 0$) pay the entry cost. An active firm earns a profit $\pi(1, \mathbf{x}_t) + \varepsilon_t(1)$ where $\pi(1, \mathbf{x}_t) = VP_t - FC_t - EC_t * (1 - y_t)$, and 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 $\mathbf{x}_t$. Finally we assume that $\varepsilon_t(0)$ and $\varepsilon_t(1)$ are extreme value type 1 distributed with dispersion parameter σ_ε.

We assume that each of the 4 different market characteristics (z_1, z_2, z_3, z_4) take K possible discrete values, $\mathcal{Z}_j = \{z_k^{(1)}, z_k^{(2)}, ..., z_k^{(K)}\}$ such that $|\mathcal{Z}_j| = K$ for $j = 1, ..., 4$. We also assume that ω is a discrete random variable that takes R possible values. The dimension of the state space $|\mathcal{X}|$ is then $2 * K^4 * R$. We assume the variables in the vector $\mathbf{z}$ are mutually independent. Each

[7] We treat productivity as "observable" but typically in real applications we will not have data on productivity and we will have to recover it from data in a separate step.

z_{kt} follows an discrete-AR(1) process. Let $\tilde{z}_{kt}$ be an continuous "latent" variable that follows the AR(1) process $\tilde{z}_{kt} = \gamma_0^k + \gamma_1^k \tilde{z}_{kt-1} + e_{kt}$, where $e_{kt} \sim$ iid $N(0, \sigma_k^2)$. Then the transition probability for the state variable z_{kt} is given by

$$\Pr(z_{kt+1} = z' | z_{kt} = z) = \begin{cases} \Phi\left(\frac{z' - \gamma_0^k - \gamma_1^k z + (w_k/2)}{\sigma_k}\right) & \text{if} \quad z' = z_k^{(1)} \\ \Phi\left(\frac{z' - \gamma_0^k - \gamma_1^k z + (w_k/2)}{\sigma_k}\right) - \Phi\left(\frac{z' - \gamma_0^k - \gamma_1^k z - (w_k/2)}{\sigma_k}\right) & \text{if} \quad z_k^{(2)} \leq z' \leq z_k^{(K-1)} \\ 1 - \Phi\left(\frac{z' - \gamma_0^k - \gamma_1^k z - (w_k/2)}{\sigma_k}\right) & \text{if} \quad z' = z_k^{(K)} \end{cases}$$

with $w_k = \frac{z_k^{(K)} - z_k^{(1)}}{K-1}$. We assume that productivity ω_{it} also follows a discrete-AR(1) process. The continuous "latent" variable $\tilde{\omega}_t$ follows the AR(1) process $\tilde{\omega}_t = \gamma_0^{\omega} + \gamma_1^{\omega}\ \tilde{\omega}_{t-1} + e_t^{\omega}$, where $e_t^{\omega} \sim$ iid$N(0, \sigma_{e^{\omega}}^2)$, and the definition of the transition probabilities for the discrete variable is the same as for the $z's$ variables. Finally, the transition of the endogenous state variable induced by the CCP is the CCP itself, i.e., $f^P(y_{t+1}|\mathbf{x}_t, P) = P(\mathbf{x}_t) = \Pr(a_t = 1|\mathbf{x}_t)$.

The Euler Equation for this model is the one in Example 2. We can represent this Euler equation in the following compact form:

$$\mathbb{E}_{\{\mathbf{z}_{t+1}|\mathbf{z}_t\}}\left[\widetilde{\pi}^c(\mathbf{x}_t;\theta) + \widetilde{\pi}^f(\mathbf{z}_{t+1};\theta) - \ln\left(\frac{P(\mathbf{x}_t;\theta)}{1 - P(\mathbf{x}_t;\theta)}\right) - \beta \ln\left(\frac{P(1,\mathbf{z}_{t+1};\theta)}{P(0,\mathbf{z}_{t+1};\theta)}\right)\right] \quad = \quad 0$$

where $\widetilde{\pi}^c(\mathbf{x}_t;\theta) \equiv [\pi(1,\mathbf{x}_t) - \pi(0,\mathbf{x}_t)]/\sigma_{\varepsilon}$, and $\widetilde{\pi}^f(\mathbf{z}_{t+1};\theta) \equiv \beta[\pi(1,1,\mathbf{z}_{t+1}) - \pi(1,0,\mathbf{z}_{t+1})]/\sigma_{\varepsilon}$, and θ represents the vector of structural parameters in the payoff function, i.e., $\theta = (\lambda,\ \alpha_0^{VP},\ \alpha_1^{VP},\ \alpha_2^{VP},\ \alpha_0^{FC},\ \alpha_1^{FC},\ \alpha_0^{EC},\ \alpha_1^{EC},\ \sigma_{\varepsilon})'$. $\mathbb{E}_{\{\mathbf{z}_{t+1}|\mathbf{z}_t\}}$ denotes the expectation operator over the distribution of $\mathbf{z}_{t+1}$ conditional on $\mathbf{z}_t$. This EE implies the following EE-PI mapping:

$$\Lambda(\mathbf{x}_t;\mathbf{P};\theta) = \frac{\exp\left\{\widetilde{\pi}^c(\mathbf{x}_t;\theta) + \mathbb{E}_{\{\mathbf{z}_{t+1}|\mathbf{z}_t\}}\left[\widetilde{\pi}^f(\mathbf{z}_{t+1};\theta) - \beta\ln\left(\frac{P(1,\mathbf{z}_{t+1};\theta)}{P(0,\mathbf{z}_{t+1};\theta)}\right)\right]\right\}}{1 + \exp\left\{\widetilde{\pi}^c(\mathbf{x}_t;\theta) + \mathbb{E}_{\{\mathbf{z}_{t+1}|\mathbf{z}_t\}}\left[\widetilde{\pi}^f(\mathbf{z}_{t+1};\theta) - \beta\ln\left(\frac{P(1,\mathbf{z}_{t+1};\theta)}{P(0,\mathbf{z}_{t+1};\theta)}\right)\right]\right\}}$$

5.2 Comparing solution methods

In this subsection we compare the computational burden using the standard PI mapping, $\Psi(\mathbf{x}_t;\mathbf{P})$, and the EE-PI mapping, $\Lambda(\mathbf{x};\mathbf{P})$, both defined over the complete state space. The DGP process used in this exercise is summarized in table 1.

Table 1
Parameters in the DGP

Distribution of ε:	Extreme Value with $\sigma_\varepsilon = 1$
Payoff Parameters:	$\lambda = 1$
	$\alpha_0^{VP} = 0.5$; $\alpha_1^{VP} = 1$; $\alpha_2^{VP} = -0.1$
	$\alpha_0^{FC} = 0.5$; $\alpha_1^{FC} = 1$
	$\alpha_0^{EC} = 1$; $\alpha_1^{EC} = 1$;
Market size (z_1):	z_{1t} is AR(1), $\gamma_0 = 0, \gamma_1 = 0.6, \sigma_e = 1$
Competition (z_2):	z_{2t} is AR(1), $\gamma_0 = 0, \gamma_1 = 0.6, \sigma_e = 1$
Fixed Cost Shock (z_3):	z_{3t} is AR(1), $\gamma_0 = 0, \gamma_1 = 0.6, \sigma_e = 1$
Entry Cost Shock (z_4):	z_{4t} is AR(1), $\gamma_0 = 0, \gamma_1 = 0.6, \sigma_e = 1$
Productivity :	ω_t is AR(1), $\gamma_0 = 0.2, \gamma_1 = 0.9, \sigma_e = 1$
Discount factor	$\beta = 0.95$

We compare the cost associated with solving a fixed point of each method for 6 different dimensions of the state space $|\mathcal{X}|$: 64, 486, 2032, 6250, 15,552. Table 2 presents the time per iteration, number of iterations and total computation time as a function of the state space dimensionality. For each value of the state space dimensionality we started from ten different initial points in searching for the fixed point. The numbers in this table are the average over these.

Table 2
Comparison of Standard and EE Policy Iteration Methods

Number states	**Number Iterations**		**Seconds Per Iteration**		**Total Time (seconds)**		**Time Ratio**
$\lvert\mathcal{X}\rvert$	EE $\Lambda(\mathbf{P})$	ST $\Psi(\mathbf{P})$	EE $\Lambda(\mathbf{P})$	ST $\Psi(\mathbf{P})$	EE $\Lambda(\mathbf{P})$	ST $\Psi(\mathbf{P})$	
64	13.2	5	0.001	0.033	0.019	0.16	18.4
486	13	4.9	0.008	0.891	0.11	4.36	39.7
2032	13	5	0.067	9.984	0.87	49.91	57.3
6250	13	5	0.478	92.813	6.21	464.06	74.8
15552	13	5	3.169	564.798	41.20	2823.99	68.5

A single iteration in the EE-PI mapping is computationally cheaper than one iteration in the standard (ST) PI method. The computational saving, per iteration, of EE-PI relative to ST-PI comes from avoiding the computation of infinite-period forward present values. This computational saving, in CPU time, increases in a convex way with the dimension of the state space, and it becomes very substantial for large state spaces. For state space $|\mathcal{X}| = 15,552$, that is still quite small relative to the dimensions that we find in some applications, the CPU seconds for one EE-PPI iteration is 3.1, while one ST-PI iteration requires 564.7 seconds.

The number of iterations to obtain convergence is larger for the EE-PI method than for the standard PI: 13 versus 5. This is because the ST-PI is a stronger contraction than the EE-PI. To compare the degree of contraction of the two mappings, we calculate the Lipschitz Constants of the two PI mappings are $c_\Psi(\mathbf{P})$ and $c_\Lambda(\mathbf{P})$ defined as:

$$\begin{aligned} c_\Psi(\mathbf{P}) &\equiv \sup_{\mathbf{x}} \left\{ \frac{||\Psi^3(\mathbf{x};\mathbf{P}) - \Psi^2(\mathbf{x};\mathbf{P})||}{||\Psi(\mathbf{x};\mathbf{P}) - \mathbf{P}||} \right\} \\ c_\Lambda(\mathbf{P}) &\equiv \sup_{\mathbf{x}} \left\{ \frac{||\Lambda^3(\mathbf{x};\mathbf{P}) - \Lambda^2(\mathbf{x};\mathbf{P})||}{||\Lambda(\mathbf{x};\mathbf{P}) - \mathbf{P}||} \right\} \end{aligned} \tag{31}$$

Table 3 reports the values of these constants, averaged over a random grid of values of $\mathbf{P}$, for different dimensions of the state space.

Table 3
Lipschitz Constants of the two PI mappings

Dimension $\|\mathcal{X}\|$	$\mathbf{c}_\Lambda$	$\mathbf{c}_\Psi$
64	0.085	0.007
486	0.053	0.004
2032	0.051	0.003
6250	0.059	0.004
15552	0.060	0.004

Very importantly, the EE-PI method shares a well-known property of the standard PI method (see Rust, 1996): the number of iterations to convergence is very stable with respect to the dimension of the state space. In fact, in our numerical example the number of iterations of the two methods remain constant at 5 and 13 iterations respectively. This property has two important implications. First, for a large enough state space, the EE-PI method is computationally more efficient than the ST-PI method: their respective number of iterations stay constant but the cost of one iteration in ST-PI increases faster with $|\mathcal{X}|$ than the cost of one EE-PI iteration. A second implication has to do with the comparison between EE-PI and value function (or Bellman equation) iteration as solution methods. Since the source of this computational savings per-iteration in EE-PI is similar to the one in value function iteration (i.e., they both require computing only one-period forward expectations), one could be tempted to claim that the computational advantages (and disadvantages) of EE-PI are similar to those of value function iteration. However, a key difference between these two solution methods is in the behavior of the number of iterations to convergence when the state space increases. For the method of value function iterations, it is well known that the number of iterations monotonically increases with the dimension of the state space. Therefore, for large state space value function iteration can be substantially more expensive than EE-PI as a solution method.

This intuition is confirmed in tables 4 and 5. In table 4 we see compare the computational costs of the value iteration method with the EE-PI method. Finding a fixed point in the space of value functions requires hundreds of iterations while the EE-PI mapping requires only 13. Each value function iteration is less costly than an EE iteration regardless of the size of the state space, but the number of value iterations makes the total time required to find a fixed point using value iterations as much as 20 times as large as the total time using the EE-PI mapping. In other words, although each value iteration is less costly than each EE-PI iteration, the steps towards convergence using the EE-PI mapping are significantly larger than the steps towards convergence using value iterations. This is made clear in table 5 when we compare the Lipschitz constants of the two mappings.

Table 4
Comparison of EE Policy Iteration and Standard Value Iteration Methods

Number states	**Number Iterations**		**Seconds Per Iteration**		**Total Time (seconds)**		**Time Ratio**
$\|\mathcal{X}\|$	EE $\Lambda(\mathbf{P})$	VF	EE $\Lambda(\mathbf{P})$	VF	EE $\Lambda(\mathbf{P})$	VF	
64	13.2	310.3	0.001	0.0003	0.019	0.094	4.95
486	13	271.4	0.008	0.004	0.11	1.009	9.173
2032	13	280.2	0.067	0.046	0.87	12.756	14.662
6250	13	294.9	0.478	0.420	6.21	123.831	19.94
15552	13	295.9	3.169	2.877	41.20	851.599	20.67

Table 5
Lipschitz Constants of the EE-PI and Value function mappings

Dimension $\lvert\mathcal{X}\rvert$	$\mathbf{c}_{\Lambda}$	$\mathbf{c}_{VF}$
64	0.085	0.319
486	0.053	0.124
2032	0.051	0.192
6250	0.059	0.268
15552	0.060	0.284

5.3 Counterfactual experiment

As we are ultimately interested in using our framework to estimate counterfactual behavior, we now study how the two methods perform relative to one another in answering an economically relevant counterfactual policy question in the context of the simple dynamic entry-exit model presented above. The counterfactual policy we consider is an increase in the cost of 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. However the presence of entry costs makes exit less attractive to incumbent firms, and in this way higher entry costs may discourage productive incumbents from exiting. The net effect of eliminating or reducing entry costs on industry output and productivity is not unambiguous. We are interested in the quantification of this effect.

Suppose that the industry consists of N potential entrants, indexed by i. Competition in this industry is characterized by monopolistic competition, i.e., single-agent model. The expected value of the total output produced by firms active in the industry is:

$$Q(\mathbf{z}_t) = \mathbb{E}\left(\sum_{i=1}^{N} a_{it} \ \exp\left(\lambda \ \omega_{it}\right) \ \mid \mathbf{z}_t \right) = N \ \int q(\mathbf{z}_t, \omega) \ \exp\left(\lambda \ \omega\right) \ f_{\omega}^{*}(\omega) \ d\omega$$

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

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

$$q(\mathbf{z}, \omega) = (1 - q(\mathbf{z}, \omega))\ P(0, \mathbf{z}, \omega) + q(\mathbf{z}, \omega)\ P(1, \mathbf{z}, \omega) \tag{32}$$

Rearranging we get:

$$q(\mathbf{z}, \omega) = \frac{P(0, \mathbf{z}, \omega)}{1 - P(1, \mathbf{z}, \omega) + P(0, \mathbf{z}, \omega)} \tag{33}$$

Notice that the effect of an increase in the entry cost on $q(\mathbf{z}, \omega)$ is ambiguous. This is because:

$$\frac{\partial q(\mathbf{z}, \omega)}{\partial EC} = \frac{\frac{\partial P(0,\mathbf{z},\omega)}{\partial EC}(1 - P(1, \mathbf{z}, \omega)) + \frac{\partial P(1,\mathbf{z},\omega)}{\partial EC} P(0, \mathbf{z}, \omega)}{\Big(1 - P(1, \mathbf{z}, \omega) + P(0, \mathbf{z}, \omega)\Big)^2} \tag{34}$$

and $\frac{\partial P(0,\mathbf{z},\omega)}{\partial EC} < 0$ (a higher entry cost makes entry less attractive for new entrants) and $\frac{\partial P(1,\mathbf{z},\omega)}{\partial EC} > 0$ (a higher entry cost makes exit less attractive for incumbents). If increasing the entry cost causes disproportionately more productive incumbents to stay in the market, higher entry cost may actually increase output.[8] We are interested in using this model to study the effect of the entry cost on expected total industry output: i.e., the effect of a change in EC on $Q(\mathbf{z}_t)$.

More specifically, the counterfactual experiment we consider is an increase in the entry cost parameter α_0^{EC} from 1 to 2.5. We consider a state space with $|\mathcal{Z}_1| = |\mathcal{Z}_2| = |\mathcal{Z}_3| = |\mathcal{Z}_4| = |\Omega| = 10$, such that the number of points in the complete state space is $|\mathcal{X}| = 2 * 10^5$. This is a realistic size for interesting applications in economics. All other settings are as in table 1.

We simulate many sets of data, and compare the average (over the samples) answer to the counterfactual question delivered by each method. For the purposes of illustration, Table 4 presents the average (over Monte Carlo simulations) of several statistics that describe the factual and counterfactual data generating processes: average probability of activity of being active, $\sum_{it} a_{it}/NT$; probability of exit, $\sum_{it}(1-a_{it})a_{it-1}/\sum_{it} a_{it-1}$; probability of entry, $\sum_{it} a_{it}(1-a_{it-1})/\sum_{it}(1-a_{it-1})$; and measure of persistence in incumbent status, $\hat{cov}(a_{it}, a_{it-1})/\hat{var}(a_{it})$.

Table 6: Summary Statistics Describing the DGP

	Probability of Activity	Exit Probability	Entry Probability	Persistence
Factual DGP	0.389	0.472	0.284	0.244
Counterfactual DGP	0.310	0.374	0.168	0.457

[8] Specifically, the sign of $\frac{\partial^2 q(\mathbf{z},\omega)}{\partial EC \partial \omega}$ is also ambiguous. That is, it is not clear if an increase in the entry cost effects the steady state probability of observing a low productivity firms more than the steady state probability of observing a high productivity firm. In principle these effects could have opposite sign. This makes it unclear how *average* output depends on the entry cost. Although the number of firms increases, the new firms may be disproportionately less productive.

These descriptive statistics are consistent with what we expect. When entry costs increase, potential entrants are less likely to enter, and potential exiters (incumbent firms) are less likely to exit. The activity decision is more persistent when entry costs are higher - firms who are currently out of the market are more likely to remain so, and similarly with incumbents. The net effect of the counterfactual is for activity to decrease.

To make a fair comparison of the two policy iteration mappings, we would like to keep the time required to solve the fixed points roughly equal. We now describe in some detail how we do so. Define the full space of the exogenous vectors $(\mathbf{z}, \omega)$ to be $\mathcal{Z}$, and let $\mathcal{Z}_N$ denote the set of points that we observe in a given sample of size N: $\mathcal{Z}_N = \{z_1, ..., z_N\}$. Let $c^E \in (0,1]$ and $c^S \in (0,1]$ be constants, and define $N^E \equiv int_+(c_E N)$ and $N^S \equiv int_+(c_S N)$ where $int_+(x)$ represents *the smallest integer greater than or equal to x*. That is, N^E and N^S are a fraction of the markets in the data. Further, define:

$$\mathcal{Z}_N^E \equiv \{\mathbf{z}_{it} : t = 1, 2;\ i = 1, 2..., N^E\}$$
$$\mathcal{Z}_N^S \equiv \{\mathbf{z}_{it} : t = 1, 2;\ i = 1, 2..., N^S\}$$

That is, $\mathcal{Z}_N^E$ is the set of exogenous vectors $\mathbf{z}$ observed over both periods in the first N^E firms, and similarly for $\mathcal{Z}_N^S$. In the Monte Carlo experiments we will solve the EE-mapping on the space defined by $\mathcal{Z}_N^E$ and we will solve the ST-mapping on the space defined by $\mathcal{Z}_N^S$. In order to keep the computation time roughly constant across the two, we will choose $c^S < c^E$. That is, we will use a smaller number of markets to define the space on which we solve the standard mapping than we do for the EE mapping.

To solve the mappings on these different spaces we must define a transition function for each. In the case of the EE-mapping we have for any value of the exogenous state vector today $\mathbf{z}$ and tomorrow $\mathbf{z}'$ such that $\mathbf{z}, \mathbf{z}' \in \mathcal{Z}_N^E$ we define:

$$\widehat{f}_N^E(\mathbf{z}'|\mathbf{z}) \equiv \frac{\sum_{i=1}^{N} K\left(\frac{\mathbf{z}_{it} - \mathbf{z}}{b_N}\right) 1\{\mathbf{z}_{it+1} = \mathbf{z}'\}}{\sum_{m=1}^{M} K\left(\frac{\mathbf{z}_{it} - \mathbf{z}}{b_N}\right)}$$

and similarly for $\widehat{f}_N^S(\mathbf{z}'|\mathbf{z})$. $\widehat{f}_N^E$ implies a matrix of dimension $|\mathcal{Z}_N^E| \times |\mathcal{Z}_N^E|$ while $\widehat{f}_N^S$ implies a matrix of dimension $|\mathcal{Z}_N^S| \times |\mathcal{Z}_N^S|$. Notice that we use all the simulated data to estimate both $\widehat{f}_N^E$ and $\widehat{f}_N^S$.

The objects that we are interested in for the purposes of answering our counterfactual question are: (1) the average probability of being active as a function of ω:

$$E[q_0(\mathbf{z}, \omega)|\omega] = \int q_0(\mathbf{z}, \omega)\ f_{\mathbf{z}}^*(\mathbf{z}) d\mathbf{z}\ ;$$

(2) the treatment effect at the (population) average market type $\tilde{\mathbf{z}}$:

$$TE(\tilde{\mathbf{z}}) = \Delta(\tilde{\mathbf{z}}) = \int \Big(q_{cf}(\tilde{\mathbf{z}}, \omega) - q_0(\tilde{\mathbf{z}}, \omega)\Big) \exp\left(\lambda\ \omega\right)\ f_{\omega}^*(\omega)\ d\omega\ ;$$

and (3) the average (over market types) treatment effect:

$$ATE = \int \Delta(\mathbf{z})\ f^*_{\mathbf{z}}(\mathbf{z}) d\mathbf{z}$$

To study the relative ability of the two methods to answer counterfactual questions, we calculate these objects using CP estimates from both the EE-PI mapping and the ST-PI mapping. Before displaying the results, we discuss in some detail how we calculate these objects. For the calculation of the average probability of activity as a function of productivity, we do the following. Consider first the Euler Equation estimates. Let $(\mathbf{z}^E, \omega^E)$ be a representative element on the grid $\mathcal{Z}^E_N$. To calculate the EE estimate of $E[q_0(\mathbf{z}, \omega)|\omega]$ we first obtain

$$\hat{q}^{EE}(\mathbf{z}, \omega) = \frac{\hat{P}^{EE}(0, \mathbf{z}, \omega)}{1 - \hat{P}^{EE}(1, \mathbf{z}, \omega) + \hat{P}^{EE}(0, \mathbf{z}, \omega)}$$

for each $(\mathbf{z}^E, \omega^E) \in \mathcal{Z}^E_N$, where $\hat{P}^{EE}$ represents the Euler Equation estimate of the choice probabilities. Let $\hat{f}^*_E(\mathbf{z})$ represent the ergodic distribution of the vector of market characteristics $\mathbf{z}$ on the space $\mathbf{z}^E_N$. Then, we calculate:

$$\hat{E}_{EE}[q(\mathbf{z}, \omega)|\omega] = \int \hat{q}^{EE}(\mathbf{z}, \omega) \hat{f}^*_E(\mathbf{z}) d\mathbf{z}$$

and similarly we can calculate the analogous object given the ST-PI mapping, $\hat{E}_{ST}[q(\mathbf{z}, \omega)|\omega]$. [9]

With the estimate of the treatment effect $\Delta(\mathbf{z})$ at each vector of market characteristics $\mathbf{z}$ in hand, the EE and ST estimates of the Average Treatment Effect can be obtained simply by integrating over the space of market characteristics using $f^*_E(\mathbf{z})$ and $f^*_S(\mathbf{z})$. The estimates of the treatment effect at the mean market type $\tilde{\mathbf{z}}$ require us to evaluate the EE and CP mappings at a point that is potentially not in the sample.

[9] Note that an important difference between the two estimates $\hat{E}_{EE}[q(\mathbf{z}, \omega)|\omega]$ and $\hat{E}_{ST}[q(\mathbf{z}, \omega)|\omega]$ is the probability distributions$\hat{f}^*_E(\mathbf{z})$ and $\hat{f}^*_S(\mathbf{z})$ used to integrate over the space of market characteristics. As $\mathcal{Z}^E_N$ will typically be much larger than $\mathcal{Z}^S_N$ (in order to keep the computation times the same) the differences in the finite sample properties of the estimates $\hat{E}_{EE}[q(\mathbf{z}, \omega)|\omega]$ and $\hat{E}_{ST}[q(\mathbf{z}, \omega)|\omega]$ will depend, to some extent, on the differences in the estimates $\hat{f}^*_E(\mathbf{z})$ and $\hat{f}^*_S(\mathbf{z})$. In particular, we have more observations for each probability we estimate in $\hat{f}^*_S(\mathbf{z})$ than we do in $\hat{f}^*_E(\mathbf{z})$, and this should contribute to higher precision for the ST-PI estimates. Of course the fact that the grid $\mathcal{Z}^E_N$ is closer to the true grid than $\mathcal{Z}^S_N$ contributes to lower bias in the case of the EE-PI mapping. One way to control for these differences, should we wish to do so, is to replace the estimated probabilities with the probabilities $f^*_E(\mathbf{z})$ and $f^*_S(\mathbf{z})$ where:

$$f^*_E(\mathbf{z}) = \frac{f^*(\mathbf{z})}{\sum_{z' \in \mathcal{Z}^E_N} f^*(\mathbf{z})} \text{for any} \mathbf{z} \in \mathcal{Z}^E_N$$
$$f^*_S(\mathbf{z}) = \frac{f^*(\mathbf{z})}{\sum_{z' \in \mathcal{Z}^S_N} f^*(\mathbf{z})} \text{for any} \mathbf{z} \in \mathcal{Z}^S_N$$

That is, we normalize the true ergodic distribution so that it is a well-defined probability distribution on the respective spaces $\mathcal{Z}^E_N$ and $\mathcal{Z}^S_N$.

5.3.1 Results

Table 5 presents the average and maximum (over the vector of exogenous variables) Root Mean Squared Error (RMSE) and the Mean Absolute Bias (MAB) of the two methods based on $1,000$ Monte Carlo simulations from the DGP.

Table 7: Results from Monte Carlo Experiments

Parameter	**Method**	**Average root MSE**	**Average MAB**
$E[q_0(\mathbf{z},\omega)\|\omega]$	**ST-mapping**	0.120	0.095
	EE-mapping	0.032	0.026
$TE(\tilde{\mathbf{z}})$	**ST-mapping**	0.086	0.066
	EE-mapping	0.045	0.036
ATE	**ST-mapping**	0.059	0.039
	EE-mapping	0.010	0.008

The results quite clearly indicate that the EE-method is considerably lower bias and MSE than the ST-method in all cases. Moreover, the EE-PI took 1/5 as much time as the ST-PI did to yield a fixed point (although the amount of time required to calculate a fixed point is very small in both cases), meaning not only is the EE mapping do a better job of making counterfactual predictions in a statistical sense, it does so quicker than the ST mapping does. While it is not so surprising that the EE-mapping has lower bias, it is perhaps surprising how much lower the variance is. In fact we can see that the proportion of the MSE due to variance in the estimates is much larger in the case of the ST-mapping. This may be due to the fact that, as we use a much smaller number of markets to define the grid on which the ST-CPs are estimated, we expect more variance in the grid over simulations, which would imply more variance in the estimates. This should be a concern for researchers applying the standard policy iteration mapping using state space reduction methods. Note as well, both methods are significantly better at predicting the average treatment effect than the treatment effect at the average. This is not so surprising. If the estimators over predict on average at some points and under predict at others, some of this will be averaged away in the ATE, but of course not for an estimate at one single point (the average).

6 Conclusion

TBW

APPENDIX

[1] Euler equation dynamic game with continuous decision variable.

From equation (a), we have that $\dfrac{\partial V_{it}^{\alpha}}{\partial y_{it}} = \dfrac{-1}{\beta}\dfrac{\partial \pi_{it-1}}{\partial a_{it-1}}$ and $\beta\dfrac{\partial V_{it+1}^{\alpha}}{\partial y_{it+1}} = -\dfrac{\partial \pi_{it}}{\partial a_{it}}$. Plugging these expressions into equation (b), we obtain:

$$(b') \quad \frac{-1}{\beta}\frac{\partial \pi_{it-1}}{\partial a_{it-1}} = \frac{d\pi_{it}^{\alpha}}{dy_{it}} - \frac{\partial \pi_{it}}{\partial a_{it}}\ \delta_i + \beta\ \frac{\partial V_{it+1}^{\alpha}}{\partial y_{jt+1}}\ \frac{\partial \alpha_{jt}}{\partial y_{it}}$$

Solving for $\beta\dfrac{\partial V_{it+1}^{\alpha}}{\partial y_{jt+1}}$ in $(b\prime)$, we get:

$$(b'') \quad \beta\ \frac{\partial V_{it+1}^{\alpha}}{\partial y_{jt+1}} = -\left[\beta\frac{\partial \alpha_{jt}}{\partial y_{it}}\right]^{-1}\left(\frac{\partial \pi_{it-1}}{\partial a_{it-1}} + \beta\left[\frac{\partial \pi_{it}}{\partial a_{it}}\ \delta_i - \frac{d\pi_{it}^{\alpha}}{dy_{it}}\right]\right)$$

Plugging equation (b'') in both the right-hand-side and the left-hand-side of equation (c), we get:

$$(c') \quad \frac{-1}{\beta}\left[\beta\frac{\partial \alpha_{jt-1}}{\partial y_{it-1}}\right]^{-1}\left(\frac{\partial \pi_{it-2}}{\partial a_{it-2}} + \beta\left[\frac{\partial \pi_{it-1}}{\partial a_{it-1}}\ \delta_i - \frac{d\pi_{it-1}^{\alpha}}{dy_{it-1}}\right]\right)$$

$$= \frac{d\pi_{it}^{\alpha}}{dy_{jt}} - \left[\beta\frac{\partial \alpha_{jt}}{\partial y_{it}}\right]^{-1}\left(\frac{\partial \pi_{it-1}}{\partial a_{it-1}} + \beta\left[\frac{\partial \pi_{it}}{\partial a_{it}}\ \delta_i - \frac{d\pi_{it}^{\alpha}}{dy_{it}}\right]\right)\left[\delta_j + \frac{\partial \alpha_{jt}}{\partial y_{jt}}\right]$$

Re-arranging terms and dating the equation two periods forward, we can re-write equation (c') as:

$$\frac{\partial \pi_{it}}{\partial a_{it}} + \beta\left[\frac{\partial \pi_{it+1}}{\partial a_{it+1}}\ \delta_i - \frac{d\pi_{it+1}^{\alpha}}{dy_{it+1}}\right] -$$

$$\beta\frac{\partial \alpha_{jt+1}}{\partial y_{it+1}}\left[r_{jt+2}^{\alpha}\ \left(\frac{\partial \pi_{it+1}}{\partial a_{it+1}} + \beta\left[\frac{\partial \pi_{it+2}}{\partial a_{it+2}}\ \delta_i - \frac{d\pi_{it+2}^{\alpha}}{dy_{it+2}}\right]\right) - \beta\ \frac{d\pi_{it+2}^{\alpha}}{dy_{jt+2}}\right] = 0$$

where $r_{jt+2}^{\alpha} \equiv \left[\delta_j + \partial\alpha_{jt+2}/\partial y_{jt+2}\right]/\partial\alpha_{jt+2}/\partial y_{it+2}$.

References

[1] Ackerberg, D., X. Chen, and J. Hahn (2012): "A Practical Asymptotic Variance Estimator for Two-Step Semiparametric Estimators," The Review of Economics and Statistics, 94(2): 481–498.

[2] Aguirregabiria, V. (1997): "Estimation of dynamic programming models with censored decision variables," *Investigaciones Economicas*, 21 (2), 167-208.

[3] Aguirregabiria, V., and Magesan, A. (2013): "Euler Equations for the Estimation of Dynamic Discrete Choice Structural Models," *Advances in Econometrics*, Volume 31 (*Structural Microeconometrics*), E. Choo and M. Shum, eds. 3-44.

[4] Aguirregabiria, V. and P. Mira (2002): "Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models," *Econometrica*, 70, 1519-1543.

[5] Aguirregabiria, V. and P. Mira (2007): "Sequential estimation of dynamic discrete games," *Econometrica*, 75, 1-53.

[6] Altug, S. and R. Miller (1998): "The effect of work experience on female wages and labour supply," Review of Economic Studies, 65, 45-85.

[7] Arcidiacono, P. and R. Miller (2011): "CCP Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity," *Econometrica*, 79, 1823-1867.

[8] Bajari, P., L. Benkard and J. Levin (2007): "Estimating dynamic models of imperfect competition," *Econometrica*, 75, 1331-1370.

[9] Berry, S., and A. Pakes (2000): "Estimation from the Optimality Conditions for Dynamic Controls," manuscript. Department of Economics. Yale University.

[10] Cooper, R., J. Haltiwanger, and J. Willis (2010): "Euler Equation Estimation for Discrete Choice Models: A Capital Accumulation Application," NBER Working Paper 15675.

[11] Hansen, L. P., and K. J. Singleton (1982): "Generalized instrumental variables estimation of nonlinear rational expectations models," *Econometrica*, 50, 1269-1286.

[12] Hotz, J., and R.A. Miller (1993): "Conditional choice probabilities and the estimation of dynamic models," *Review of Economic Studies*, 60, 497-529.

[13] Hotz, J., R.A. Miller, S. Sanders, and J. Smith (1994): "A simulation estimator for dynamic models of discrete choice," *Review of Economic Studies*, 61, 265-89.

[14] Manski, C. F. (1993): "Dynamic choice in social settings: Learning from the experiences of others," *Journal of Econometrics*, 58 (1), 121-136.

[15] McFadden, D. (1981): "Econometric Models of Probabilistic Choice," in C. Manski and D. McFadden (eds.), Structural Analysis of Discrete Data with Econometric Applications, MIT Press, Cambridge, MA.

[16] McFadden, D. (1989): "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration," *Econometrica*, 57 (5), 995-1026.

[17] Newey, W.K. (1984): "A Method of Moments Interpretation of Sequential Estimators," Economics Letters 14, pp. 201 –206.

[18] Newey, W.K. (1994): "The Asymptotic Variance of Semiparametric Estimators," Econometrica, 62, 1349 –1382.

[19] Newey, W.K. and D. F. McFadden (1994): "Large sample estimation and hypothesis testing", in R.F. Engle III and D.F. McFadden (eds.), The Handbook of Econometrics, vol. 4. North-Holland, Amsterdam.

[20] Pakes, A. (1994): "Dynamic structural models, problems and prospects," in C. Sims (ed.) Advances in Econometrics. Sixth World Congress, Cambridge University Press.

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

[22] Pakes, A. and D. Pollard (1989). "Simulation and the Asymptotics of Optimization Estimators," Econometrica, 57(5), 1027-57.

[23] Pesendorfer, M. and P. Schmidt-Dengler (2008): "Asymptotic Least Squares Estimators for Dynamic Games," The Review of Economic Studies. Forthcoming.

[24] Rust, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, 55, 999-1033.

[25] Rust, J. (1992): "Do People Behave According to Bellman's Principle of Optimality?" Working Paper E-92-10. The Hoover Institute. Stanford University.

[26] Rust, J. (1994): "Structural estimation of Markov decision processes," in R. E. Engle and McFadden (eds.) *Handbook of Econometrics Volume 4*, North-Holland.

[27] Rust, J. (1996): "Numerical dynamic programming in economics," Handbook of computational economics, 1, 619-729.

[28] Rust, J. (1997): "Using Randomization to Break the Curse of Dimensionality," Econometrica, 65 (3), 487-516.

[29] Stern, S. (1997). Simulation-based estimation. Journal of Economic Literature, 35, 2006–2039.

[30] Stokey, N., R. Lucas, and E. Prescott (1989): "Recursive methods in economic dynamics," Harvard University Press.

[31] Su, C-L., K. Judd (2012): "Constrained Optimization Approaches to Estimation of Structural Models," Econometrica, 80(5), 2213-2230.