

Another Look at the Identification of Dynamic Discrete Decision Processes

Victor Aguirregabiria*

Boston University, Department of Economics

270 Bay State Road. Boston, MA 02215. E-mail: vaguirre@bu.edu

April, 2005

Abstract

This paper presents an econometric approach to estimate the behavioral effects of counterfactual policy experiments in the context of dynamic decision models where the current utility function and the distribution of unobservables are nonparametrically specified. Previous studies have shown that the identification of the current utility function in dynamic decision models requires of stronger assumptions than in static decision models. We show in this paper that knowledge of the current utility function (or of a "normalized" utility function) is not necessary to identify counterfactual choice probabilities in dynamic models. To identify these counterfactuals we need the probability distribution of the unobservables and the difference between the present value of choosing always the same alternative and the present value of deviating one period from this strategy. We show that both functions are identified from the factual choice probabilities under similar conditions as in static decision models. Based on this result we propose a nonparametric procedure to estimate the behavioral effects of counterfactual experiments in dynamic decision models. We apply this method to evaluate the effects of an investment subsidy program in the context of a model of machine replacement.

Keywords: Dynamic discrete decision processes; Nonparametric identification; Counterfactual experiments.

JEL: C13, C25

*I want to thank Jim Heckman, Thierry Magnac and John Rust for very useful comments.

1 Introduction

Discrete choice dynamic structural models have proven useful tools for the assessment of public policy initiatives. These econometric models have been applied to the evaluation of different economic policies, both factual and counterfactual, like welfare policies (Sanders and Miller, 1997, Keane and Moffit, 1998, and Keane and Wolpin, 2000), unemployment insurance (Ferrall, 1997), social security and retirement (Berkovec and Stern, 1991, and Rust and Phelan, 1997), patents regulation (Pakes, 1986, and Pakes and Simpsom, 1989), education policies (Eckstein and Zilcha, 1994, Eckstein and Wolpin, 1999, and Keane and Wolpin, 2001), contraceptive choice (Hotz and Miller, 1993), regulation on labor contracts (Aguirregabiria and Alonso-Borrego, 1999, and Rota, 2004), programs on child poverty (Wolpin and Todd, 2003), scrapping subsidies (Adda and Cooper, 2000), or regulation of nuclear plants (Rust and Rothwell, 1995).

A common feature of the econometric models in these applications is the parametric specification of structural functions like agents' utilities, transition probabilities of state variables, and the probability distributions of unobservables.¹ These parametric models contrast with the emphasis on nonparametric specification that we find in other approaches to evaluate public policies. In particular, the literature on evaluation of treatment effects has emphasized the importance of a nonparametric specification of the distribution of unobservables to obtain robust results (see Heckman and Robb, 1985, Manski, 1990, and more recently Heckman and Smith, 1998, and Heckman and Vyltasil, 1999). An important factor that has contributed to the parametric specification of dynamic discrete structural models has been the negative results on the nonparametric identification of the utility function in these models. As shown by Rust (1994) and Magnac and Thesmar (2002), the differences between the utilities of two choice alternatives cannot be identified in dynamic decision models even when the researcher "knows" the time discount factor, the probability distribution of the unobservables, and the transition probabilities of the state variables. This under-identification result contrasts with the identification of utility differences in static (i.e., not forward looking) decision models (see Matzkin, 1992).

This paper takes a different look at the problem of nonparametric identification of dynamic decision models. Instead of looking at the nonparametric identification of the utility function or of utility differences we consider the identification of the behavioral effects of counterfactual policy changes. That is, we study the identification of the choice probabilities associated with hypothetical changes in the utility function. We show that knowledge of

¹An exception is the semiparametric model in Taber (2000) where utilities are parametrically specified but the distribution of the unobservables is nonparametric.

the current utility function or of utility differences is not necessary to identify counterfactual choice probabilities in dynamic decision models. These counterfactuals depend on the distribution of unobservables and on the difference between the value of choosing always the same alternative and the value of deviating one period from that policy. We show that these functions are identified under similar conditions as in static models. Therefore, though agents' preferences cannot be identified, we can identify the behavioral effects associated with changes in these preferences.

Based on this identification result we propose a nonparametric procedure to estimate the behavioral effects of counterfactual experiments in this class of models. The computational cost of this method is equivalent to solving the dynamic programming problem twice (i.e., before and after the policy change). Therefore, from a computational point of view, the method is at least as feasible as a structural parametric approach. In order to analyze the ability of this method to provide informative estimates of policy effects, we implement a Monte Carlo experiment. In this experiment we evaluate the effects of an investment subsidy program in the context of a model of machine replacement. We show that the method provides precise estimates of policy effects when these effects are aggregated over some of the state variables.

This paper is related to the analysis in Heckman and Navarro (2005). These authors present a comprehensive study on semiparametric identification of dynamic discrete choice models that includes a wide range of reduced form models, but also a subclass of the structural models that we consider in this paper. There are important differences between our paper and Heckman and Navarro's analysis of structural models. First, they study semiparametric identification of the complete structure of the model, while we consider nonparametric identification of counterfactual choice probabilities. Second, they consider finite horizon optimal stopping problems and prove identification using backward induction. Instead we consider a general class of models proposed by Rust (1994) that encompasses both infinite horizon and finite horizon problems and not only optimal stopping problems. We cannot use backward induction to prove identification. Instead, we show that our objects of interest are unique fixed points of contraction mappings that depend only on the data.

The rest of the paper is organized as follows. Section 2 presents the model, the basic assumptions and the type of counterfactual policy experiments that we want to evaluate. Section 3 contains the identification results. We describe the estimation procedure in section 4. The Monte Carlo experiment is in section 5. We summarize and conclude in section 6. The proofs of propositions are in the appendix.

2 Model

2.1 Framework and basic assumptions

Time is discrete and indexed by t . Consider an agent who has preferences defined over a sequence of states of the world from period $t = 0$ to $t = T$. A state of the world has two components: a vector of state variables s_t that is given at period t ; and a discrete decision $a_t \in A = \{0, 1, \dots, J\}$ that the agent chooses at period t . The decision at period t affects the evolution of future values of the state variables. The agent's preferences over possible sequences of states of the world can be represented by a utility function $\sum_{j=0}^T \beta^j U_t(a_{t+j}, s_{t+j})$, where $\beta \in [0, 1)$ is the discount factor and $U_t(a_t, s_t)$ is the current utility function at period t .² The agent has uncertainty about future values of state variables. His beliefs about future states can be represented by a sequence of Markov transition probability functions $F_{s,t}(s_{t+1}|a, s)$. These beliefs are *rational* in the sense that they are the true transition probabilities of the state variables. Every period t the agent observes the vector of state variables s_t and chooses his action $a_t \in A$ to maximize the expected utility

$$E \left(\sum_{j=0}^T \beta^j U_t(a_{t+j}, s_{t+j}) \mid a_t, s_t \right) \quad (1)$$

Let $\alpha_t(s)$ and $V_t(s)$ be the optimal decision rule and the value function at period t , respectively. By Bellman principle of optimality the sequence of value functions can be obtained using the recursive expression:

$$V_t(s_t) = \max_{a \in A} \left\{ U_t(a, s_t) + \beta \int V_{t+1}(s_{t+1}) dF_{s,t}(s_{t+1}|a, s_t) \right\} \quad (2)$$

For the rest of the paper we adopt a notation that omits the time subindex from functions and variables. A finite-horizon dynamic decision problem can be represented as an infinite-horizon problem if we just make the utility functions equal to zero for any period $t > T$. Furthermore, we can include the time period t as a state variable of the model and therefore we can omit it as an index in the structural functions and in the optimal decision rule and the value function. Finally, we can also omit the time subindex in the decision and state variables and use (a, s) to represent current values of these variables, and (a', s') for next period values.

From the point of view of the observing researcher there are two types of state variables, $s = (x, \xi)$, where the vector x is observable to the researcher and the vector ξ is unobservable. Without loss of generality we can write the one-period utility as the sum of two components:

$$U(a, s) = u(a, x) + \varepsilon(a, x, \xi), \quad (3)$$

²In fact, the function $U(\cdot, \cdot)$ is an *indirect utility function* that incorporates implicitly the agent's budget constraint and other possible constraints.

where $u(a, x) \equiv E(U(a, s) | a, x)$ and $\varepsilon(a, x, \xi) \equiv U(a, s) - E(U(a, s) | a, x)$. For notational simplicity we use $\varepsilon(a)$ to denote $\varepsilon(a, x, \xi)$. We use also the vector $\varepsilon = \{\varepsilon(a) : a \in A\}$ instead of ξ to represent unobservable state variables. Note that, by definition, the random variables in ε have zero mean and are mean independent of x . We follow Rust (1994) and consider the following assumptions on the distribution of the state variables.

ASSUMPTION 1: The cumulative transition probability of the state variables factors as $F_s(s'|a, s) = F_\varepsilon(\varepsilon') F_x(x'|a, x)$.

Under this assumption the optimal decision rule $\alpha(x, \varepsilon)$ can be described as:

$$\alpha(x, \varepsilon) = \arg \max_{a \in A} \{ v(a, x) + \varepsilon(a) \} \quad (4)$$

where $v(a, x)$ is the expected discounted value of current and future utilities if current decision is a . That is,

$$v(a, x) \equiv u(a, x) + \beta \int \max_{j \in A} \{ v(j, x') + \varepsilon'(j) \} dF_\varepsilon(\varepsilon') dF_x(x'|a, x) \quad (5)$$

The functions $v(0, x)$, $v(1, x)$, ..., $v(J, x)$ are called *conditional choice value functions*.

From the point of view of the researcher, who does not observe ε , agents' optimal behavior can be described in terms of the *optimal choice probability function*:

$$P(a|x) \equiv \int I\{\alpha(x, \varepsilon) = a\} dF_\varepsilon(\varepsilon) \quad (6)$$

The *reduced form* of the model consists of this optimal choice probability function and the transition probability function $F_x(x'|a, x)$. The *model structure* consists of the functions $\{u, F_\varepsilon, F_x\}$ and the discount factor β .

2.2 Data and identification of the reduced form

Suppose that there is a population of individuals who behave according to this model. We have a random sample of individuals from this population and we observe $\{a_i, x_i, x'_i\}$ for each individual i . As usual, we study identification with a very large (i.e., infinite) sample. Furthermore, we assume that the sample has variability over the whole support $A \times X \times X$ of the variables $\{a, x, x'\}$. This assumption of full-support variation is needed to identify the reduced form of the model. The assumption implies different data requirements depending on the type of model. In stationary models, a cross-section of individuals is enough to get full support variation. This is also the case in life-cycle models where the agent's age is a state variable, as long as the cross-section includes individuals of all ages and the vector of state variables does not include calendar time and the individual's cohort. If both age and

cohort are state variables, we need panel data or repeated cross-sections to guarantee full support variability. In models with aggregate state variables we typically need panel data with a long time dimension to guarantee that aggregate variables have variation over their full support. However, Altug and Miller (1998) show that under some assumptions on the effects of aggregate variables, we can identify the reduced form with a short panel.

Under the assumption of rational expectations, it is clear that we can identify $F_x(x'|a, x)$ on $A \times X \times X$ from the transition probabilities $\Pr(x'_i = x'|a_i = a, x_i = x)$ in the data. Furthermore, under the assumption of agents' optimal behavior, the optimal probability function $P(a|x)$ can be identified on $A \times X$ from the probabilities $\Pr(a_i = a|x_i = x)$ in the data. Thus, the reduced form probability functions P and F_x are nonparametrically identified on their respective supports.

However, we cannot identify the structural functions $\{u, F_\varepsilon\}$ without further restrictions on the primitives of the model. This is the case both in decision models where agents are forward looking (i.e., $\beta > 0$) and in models where agents are myopic (i.e., $\beta = 0$). In this paper we are not interested in the identification of the complete structure of the model. Instead we consider the identification of individuals' predicted behavior (i.e., optimal probability functions) associated with a counterfactual change in the utility function. Next section describes the type of policy experiments that this paper is concerned with.

2.3 Policy experiments

Suppose that we know the reduced form $\{P, F_x\}$ associated with a structure $\{u, F_\varepsilon, F_x, \beta\}$. We want to evaluate the behavioral effects of a hypothetical policy intervention that modifies the current utility function.³ The current utility function after the intervention is $U^*(.,.)$ such that:

$$U^*(a, s) = U(a, s) + \tau(a, x) \tag{7}$$

The function $\tau(.,.)$ represents the change in the current utility induced by the hypothetical policy. Note that this change can vary across observable state variables and across choice alternatives in a completely unrestricted way. The only restriction is that it does not depend on unobservable state variables.⁴ The researcher does not know neither the factual nor the

³Note that the function $U(.,.)$ is an *indirect utility function*. That is, it incorporates the agent's budget constraint and other possible restrictions. Therefore, the policy intervention does not modify agents' preferences but earnings and expenditures that appear in the budget constraint or in other relevant constraints.

⁴Though we consider that the change in the current utility function does not depend on unobservable state variables, all our results can be easily extended to the case in which the utility change has the following form:

$$U^*(a, s) = U(a, s) + \tau(a, x) + h(x) \varepsilon(a)$$

where both $\tau(a, x)$ and $h(x)$ are known, and $\varepsilon(a)$ is the same unobservable variable as before the policy change.

counterfactual utility functions, but he knows the utility changes $\tau(a, x)$ because these are given by the policy he wants to evaluate.

By definition, we have that $u(a, x) \equiv E(U(a, s) \mid a, x)$ and $u^*(a, x) \equiv E(U^*(a, s) \mid a, x)$. Therefore,

$$u^*(a, s) = u(a, s) + \tau(a, x) \tag{8}$$

Also by definition, $\varepsilon(a) \equiv U(a, s) - u(a, s)$ and $\varepsilon^*(a) \equiv U^*(a, s) - u^*(a, s)$. This implies that $\varepsilon^*(a) = \varepsilon(a)$ and $F_\varepsilon^* = F_\varepsilon$. That is, the policy change only modifies the part of the utility associated with (a, x) and therefore it does not alter the distribution of the unobservable state variables. Note that the transition probability functions of the state variables are also unaffected. Therefore, the counterfactual structure is $\{u^*, F_\varepsilon, F_x, \beta\}$.

Let P^* be the optimal choice probability function associated with the counterfactual structure. The difference between the functions P^* and P represents the behavioral effects of the policy from the point of view of the econometrician. Given our knowledge of $\{P, F_x, \tau\}$ we want to identify nonparametrically the counterfactual choice probabilities.

2.4 Example: Machine replacement and investment subsidies

Dynamic structural models of machine replacement have been considered before by, among others, Rust (1986 and 1987), Sturm (1991), Das (1992), Kennet (1993 and 1994), Rust and Rothwell (1995), Cooper, Haltiwanger and Power (1999), Adda and Cooper (2000), Cho (2002), and Kasahara (2004). Most of these studies use the estimated structural model to evaluate the behavioral effects of a policy change. Kennet (1993) studies how deregulation of the US airline industry affected the number of aircraft engine hours between major overhauls. Rust and Rothwell (1995) analyze the impact on the operation of US nuclear power plants of an increase in the intensity of safety regulation by the US Nuclear Regulatory Commission after an accident on March 1979. Adda and Cooper (2000) evaluate the effects of a policy in France in which the government subsidized the replacement of old cars with new ones. Kasahara (2004) examines the impact on firms' investment in equipment of a temporary increase in import tariffs in Chile.

The model that we present here is similar to the one in Cooper, Haltiwanger and Power (1999). Consider a firm that produces a good using capital and some perfectly flexible inputs. Output depends on the stock of capital K_t , the amount of variable inputs, and a shock to total factor productivity ω_t . The profit function in this model is an indirect profit where we have already solved the optimal amounts of variable inputs. Given the productivity shock and the capital stock, the firm decides whether to replace the existing capital with a new machine or continuing with the same capital for another period. Let $a_t \in \{0, 1\}$ be

the indicator of the machine replacement decision. If the producer decides not to replace, the machine depreciates at rate λ such that $K_t = (1 - \lambda) K_{t-1}$. If the producer chooses replacement, then the capital stock associated with a new machine is $K_t = 1$. Thus, the transition rule for the capital stock is:

$$K_t = a_t + (1 - a_t) (1 - \lambda) K_{t-1} \quad (9)$$

The productivity shock follows an exogenous Markov process with transition probability function $F_\omega(\omega_{t+1}|\omega_t)$.

The profit function has the following form:

$$U_t = y(\omega_t, K_t) - (mc(K_t) + \xi_{1t}) - a_t (rc(K_{t-1}) + \xi_{2t}) \quad (10)$$

$y(., .)$ is the production function. $mc(K_t) + \xi_{1t}$ is the machine maintenance cost, that depends on the age of capital through the function $mc(.)$. The term ξ_{1t} is a random shock in the maintenance cost that is unobservable to the researcher. $rc(K_t) + \xi_{2t}$ is the replacement cost net of the scrapping value of the retired capital. The term ξ_{2t} is a random shock in the replacement cost that is also unobservable to the researcher. Following the notation in section 2.1, we have that $x = (\omega_t, K_{t-1})$, $\varepsilon_t(0) = -\xi_{1t}$, $\varepsilon_t(1) = -\xi_{1t} - \xi_{2t}$, and:

$$u(a_t, x_t) = \begin{cases} y(\omega_t, (1 - \lambda)K_{t-1}) - mc((1 - \lambda)K_{t-1}) & \text{if } a_t = 0 \\ y(\omega_t, 1) - mc(1) - rc((1 - \lambda)K_{t-1}) & \text{if } a_t = 1 \end{cases} \quad (11)$$

Let $Age_t \in \{0, 1, 2, \dots\}$ be the age of the machine at the beginning of period t , before the replacement decision. The evolution of Age_t is: $Age_{t+1} = (1 - a_t) (Age_t + 1)$. There is a one-to-one relationship between capital stock and the age of the machine: i.e., $K_{t-1} = (1 - \lambda)^{Age_t}$. Therefore, we can use either Age_t instead of K_{t-1} as state variable of the model.

Suppose that we are interested in evaluating the effects of a counterfactual policy that modifies firms' replacement costs. This policy tries to promote the retirement of old capital by providing a subsidy that depends on the age of the retired capital. The amount of the subsidy, that coincides with the change in the current profit function associated with the policy, is:

$$\tau(a_t, x_t) = \begin{cases} 0 & \text{if } Age_t < \underline{Age} \\ a_t (\tau_0 - \tau_1 [Age_t - \underline{Age}]) & \text{if } \underline{Age} \leq Age_t \leq \overline{Age} \\ 0 & \text{if } Age_t > \overline{Age} \end{cases} \quad (12)$$

where $\tau_0 > 0$, $\tau_1 > 0$, \underline{Age} and \overline{Age} are parameters that characterize the policy. The subsidy is zero if replacement takes place either too early or too late. For replacement ages within the range $[\underline{Age}, \overline{Age}]$ the subsidy is strictly positive and it decreases linearly with the age of

capital. This type of policy has been common in many countries and it has been motivated as part of an environmental policy to reduce emissions of carbon dioxide.

Let $P^*(\omega_t, Age_t)$ be the replacement probability function under the new policy. Given this probability function, together with the stochastic process of the productivity shock, we can derive different interesting measures on the effect of the policy, both in the long run and during the transition period to the new steady-state.

(a) *Changes in the age distribution of capital.* The joint (steady-state) distribution of the age of capital and the productivity shock, $H^*(Age, \omega)$, is the solution to the following system of equations:

$$\begin{aligned} \text{For } Age = 0 : \quad H^*(\omega, 0) &= \sum_{j=1}^{\infty} P^*(\omega, j) H^*(\omega, j) \\ \text{For } Age > 0 : \quad H^*(\omega, Age) &= (1 - P^*(\omega, Age - 1)) H^*(\omega, Age - 1) \end{aligned} \quad (13)$$

To obtain the marginal distribution for the age of capital we integrate $H^*(\omega, Age)$ over the productivity shock: $\bar{H}^*(Age) = \int H^*(Age, \omega) d\omega$.

(b) *Change in aggregate investment.* Let I and I^* be the factual and the counterfactual aggregate investments, respectively. We can get these aggregate variables by integrating the replacement probabilities over the steady-state distribution of (ω, Age) .

$$I^* = \sum_{j=1}^{\infty} \int P^*(j, \omega) H^*(j, \omega) d\bar{F}_{\omega}(\omega) \quad (14)$$

3 Identification

This section presents our identification results. In order to illustrate the differences and similarities between dynamic and static decision models in the identification of counterfactual choice probabilities, we start with the identification of static decision models in section 3.1. Then, section 3.2 presents our identification results for dynamic decision models. Both sections consider binary choice models. Section 3.3 explains how these results can be extended to the multinomial case. For notational simplicity, we use $P(x)$ instead of $P(1|x)$ to represent the probability of choosing alternative 1 in the binary choice model.

3.1 Static model

Consider the model in section 2 but now the set of choice alternatives is binary and agents are not forward looking, i.e., $\beta = 0$. The counterfactual choice probabilities in this static model are:

$$\begin{aligned} P^*(x) &= \Pr (u^*(1, x) + \varepsilon(1) \geq u^*(0, x) + \varepsilon(0) \mid x) \\ &= F_{\tilde{\varepsilon}} (\tilde{u}(x) + \tilde{\tau}(x)) \end{aligned} \quad (15)$$

where $\tilde{u}(x) \equiv u(1, x) - u(0, x)$, $\tilde{\tau}(x) \equiv \tau(1, x) - \tau(0, x)$, and $F_{\tilde{\varepsilon}}$ is the distribution function of the random variable $\tilde{\varepsilon} \equiv \varepsilon(0) - \varepsilon(1)$. This expression illustrates that the identification of $P^*(.)$ requires one to identify the functions $F_{\tilde{\varepsilon}}$ and \tilde{u} . The relationship between these functions and the *factual* reduced form probability function $P(.)$ is:

$$P(x) = F_{\tilde{\varepsilon}}(\tilde{u}(x)) \quad (16)$$

Manski (1988) provided a very comprehensive analysis of different identification conditions for the case in which $F_{\tilde{\varepsilon}}$ is nonparametrically specified but \tilde{u} is parametric and linear, i.e., $\tilde{u}(x) = x'\theta$. He also consider the case in which x and $\tilde{\varepsilon}$ are not independently distributed (e.g., quantile independence, index sufficiency). Matzkin (1992) shows that it is possible to identify $F_{\tilde{\varepsilon}}$ and \tilde{u} without imposing any parametric structure in any of the two functions. Let \mathcal{F} be the space of possible \tilde{u} functions, and let $\tilde{u}(X) \subseteq \mathbb{R}$ be the space of real values that the function \tilde{u} can take on X . That is, $\tilde{u}(X) \equiv \{ \tilde{u}(x) \in \mathbb{R} : x \in X \}$. The key condition that Matzkin considers to separate the influence of $F_{\tilde{\varepsilon}}$ and \tilde{u} on P is the following.

Assumption W2 in Matzkin (1992): There is a subset $\bar{X} \subset X$ such that: (i) for any pair of functions $h, h' \in \mathcal{F}$ the difference $h(x) - h'(x)$ is constant over the subset \bar{X} ; and (ii) for every function $h \in \mathcal{F}$ and for every value $r \in \tilde{u}(X)$ there is an $x \in \bar{X}$ such that $h(x) = r$.

Since all functions in \mathcal{F} attain the same values on the subset \bar{X} , different values of the probability $P(x)$ on \bar{X} should be attributed to the distribution function $F_{\tilde{\varepsilon}}$. Furthermore, W2(ii) implies that we can identify $F_{\tilde{\varepsilon}}$ over the whole set $\tilde{u}(X)$.

Matzkin shows that an example where assumption W2 holds is when the functions in \mathcal{F} are homogeneous of degree one. Here I consider a different type of condition under which assumption W2 holds. This condition holds in many applications of discrete choice models where the econometrician observes an outcome variable that is part of the utility function, e.g., individuals' wages, firms' output, etc.

ASSUMPTION 2: The current utility function can be written as:

$$u(a, x) = y(a, w, z) + c(a, z) \quad (17)$$

where $x = (w, z) \in W \times Z$, and $y(., ., .)$ and $c(., .)$ are real valued functions such that: (i) the function $\tilde{y}(w, z) \equiv y(1, w, z) - y(0, w, z)$ is known to the researcher; (ii) w is a continuous random variable and $\tilde{y}(w, z)$ is strictly increasing in w ; and (iii) for any $z \in Z$ there is a $w \in W$ such that $\tilde{y}(w, z) + \tilde{c}(z) = 0$, where $\tilde{c}(z) \equiv c(1, z) - c(0, z)$.

Under this assumption, the set of admissible threshold functions \mathcal{F} is such that for any $h \in \mathcal{F}$, $h(w, z) = \tilde{y}(w, z) + \tilde{c}(z)$, where $\tilde{y}(w, z)$ is known and strictly increasing in w . This

restriction on the space \mathcal{F} is a particular case of Matzkin's assumption W2. To see this, define $\bar{X} = W \times \{z_0\}$ where z_0 is an arbitrary point in the space Z . For any two functions $h, h' \in \mathcal{F}$ and any $(w, z_0) \in \bar{X}$, the function $h(w, z_0) - h'(w, z_0)$ depends only on z_0 and therefore it is constant over the subset \bar{X} . This is Matzkin's assumption W2(i). It is also simple to verify that assumptions 2(i) and 2(ii) imply Matzkin's assumption W2(ii).

As mentioned above, Assumption 2 holds in many applications of discrete choice models where the researcher observes and outcome variable that is part of the utility function. Using this outcome variable, together with observable decision and state variables, it is possible to estimate the function $y(., ., .)$ that is a component of the utility function. For instance, this is the case in models of individuals' labor supply, labor force participation, occupational choice, or schooling where earnings are observable to the researcher and they are a component of the utility function. Other cases in which assumption 2 holds are in models of firms' investment, labor demand, inventory choice or pricing decisions where the firms' sales are observable to the econometrician and they are part of the profit function.

Proposition 1 establishes the nonparametric identification of the structural functions $F_{\tilde{\varepsilon}}$ and \tilde{u} and of the counterfactual probability function P^* in this static model.

PROPOSITION 1: Suppose that Assumptions 1 and 2 hold and that the median of $\tilde{\varepsilon}$ is zero. Then, the function $\tilde{u}(\cdot)$ is nonparametrically identified on X , the distribution function $F_{\tilde{\varepsilon}}$ is nonparametrically identified over the set $\tilde{u}(X)$, and the counterfactual choice probability function $P^(\cdot)$ is identified over the set $X^* = \{x \in X : \tilde{u}(x) + \tilde{\tau}(x) \in \tilde{u}(X)\}$.*

The counterfactual probability function is identified on the set X^* that is included in X . However, there are different cases in which $X^* = X$.

Case I: $\tilde{u}(X) = \mathbb{R}$. This is the case in models where the range of variation of the function $\tilde{y}(\cdot, \cdot)$, or of the function $\tilde{c}(\cdot)$, is the whole real line. Then, $\tilde{u}(X) = \mathbb{R}$ and it is clear that this implies that $X^* = X$.

Case II: $\tilde{u}(X)$ is unbounded from above (below) and $\tilde{\tau}(\cdot)$ is positive (negative) valued. This is typically the case in applications where $y(\cdot, \cdot, \cdot)$ is an earnings equation or a production function.

Case III: $\tilde{u}(X) \subset \mathbb{R}$ but $\tilde{\tau}(\cdot)$ is such that $\sup_{x \in X} \{\tilde{u}(x) + \tilde{\tau}(x)\} \leq \sup_{x \in X} \{\tilde{u}(x)\}$ and $\inf_{x \in X} \{\tilde{u}(x) + \tilde{\tau}(x)\} \geq \inf_{x \in X} \{\tilde{u}(x)\}$. That is, the policy that we want to evaluate is such that it reduces (increases) the utility differential in states where it is large (small).

3.2 Dynamic model

We now consider the identification of counterfactual choice probabilities when agents are forward looking, i.e., when $\beta > 0$. The optimal choice probability function in this model is:

$$P(x) = F_{\tilde{\varepsilon}}(\tilde{v}(x)) \quad (18)$$

where $\tilde{v}(x) \equiv v(1, x) - v(0, x)$ is the *differential value function*. The counterfactual choice probability function is $P^*(x) = F_{\tilde{\varepsilon}}(\tilde{v}^*(x))$, where $\tilde{v}^*(.)$ is the differential value function after the policy change. We show in this section that the functions $F_{\tilde{\varepsilon}}$, \tilde{v}^* and P^* are identified under similar conditions as in Proposition 1.

However, there are two main differences with the identification in the static case. First, in the dynamic model we cannot identify current utility differences \tilde{u} or any other function that depends only on preferences and not on agent's beliefs. That is, we cannot separate agents' preferences and agents' beliefs. Despite this under-identification of preferences, we can identify the counterfactual choice probabilities. Second, given $F_{\tilde{\varepsilon}}$ the identification of P^* in the static case is pointwise in the space X . That is, given the factual choice probability at point x_0 we can get the counterfactual probability at that point as $P^*(x_0) = F_{\tilde{\varepsilon}}(F_{\tilde{\varepsilon}}^{-1}(P(x_0) + \tilde{\tau}(x_0)))$, where $F_{\tilde{\varepsilon}}^{-1}(.)$ is the inverse function of $F_{\tilde{\varepsilon}}$. In contrast, in the dynamic case we need the whole factual probability function $P(.)$ to identify the counterfactual probability at a single point.

For the sake of clarity, it is convenient to describe our identification results in two steps. First, we show the identification of P^* when the distribution $F_{\tilde{\varepsilon}}$ is known. Second, we prove the joint identification of $F_{\tilde{\varepsilon}}$, \tilde{v}^* and P^* .

3.2.1 Identification with known $F_{\tilde{\varepsilon}}$

Suppose that the probability distribution $F_{\tilde{\varepsilon}}$ is known to the researcher. Then, it is clear that the differential value function $\tilde{v}(.)$ is identified from the factual choice probabilities: i.e., $\tilde{v}(x) = F_{\tilde{\varepsilon}}^{-1}(P(x))$. In the static case, it is obvious that the identification of the differential utility function \tilde{u} implies the identification of the counterfactuals \tilde{u}^* and P^* because $\tilde{u}^* = \tilde{u} + \tilde{\tau}$ and $P^*(x) = F_{\tilde{\varepsilon}}(\tilde{u}(x) + \tilde{\tau}(x))$. However, that is not the case in the dynamic model. Knowledge of the function \tilde{v} is not enough to identify the counterfactual \tilde{v}^* . The reason is that \tilde{v}^* is not just a function of \tilde{v} and τ . To obtain how the policy change affects the differential value function (i.e., to obtain \tilde{v}^*) we need more information than just the factual value function \tilde{v} .

We show here that we can identify separately different components of the differential value \tilde{v} . Given this decomposition we will be able to construct the counterfactual function

\tilde{v}^* . Proposition 2 provides a characterization of the choice probability function that will be very useful to identify and to estimate the counterfactuals.

PROPOSITION 2: The optimal choice probability function P is the unique fixed point of the mapping $\Psi(P)$, where

$$\Psi(P)(x) \equiv F_{\tilde{\varepsilon}} \left(\tilde{\varphi}(x) + \tilde{\delta}(x, P) \right) \quad (19)$$

and (1) $\tilde{\varphi}(x) = \varphi(1, x) - \varphi(0, x)$, where $\varphi(a, x)$ is the value of choosing alternative a today and then select alternative 0 forever in the future; and (2) $\tilde{\delta}(x, P) = \delta(1, x, P) - \delta(0, x, P)$, where $\delta(a, x, P)$ is the value of behaving optimally in the future minus the value of choosing always alternative 0, given that the current choice is a . These functions can be obtained recursively as follows:

$$\varphi(a, x) = u(a, x) + \beta \int \varphi(0, x') dF_x(x'|a, x) \quad (20)$$

and

$$\delta(a, x, P) = \beta \int (G(x', P) + \delta(0, x', P)) dF_x(x'|a, x) \quad (21)$$

where $G(x, P)$ is McFadden's surplus function that is defined as $\int \max \{0; \tilde{v}(x) - \tilde{\varepsilon}\} dF_{\tilde{\varepsilon}}(\tilde{\varepsilon})$, and it can be represented as a function of the optimal choice probability $P(x)$ as

$$G(x, P) = P(x) F_{\tilde{\varepsilon}}^{-1}(P(x)) - \int_{-\infty}^{F_{\tilde{\varepsilon}}^{-1}(P(x))} \tilde{\varepsilon} dF_{\tilde{\varepsilon}}(\tilde{\varepsilon}) . \quad (22)$$

Proposition 2 establishes that we can decompose the differential value function $\tilde{v}(x)$ in two terms: $\tilde{\varphi}(x)$ and $\tilde{\delta}(x, P)$. This decomposition has several important implications. First, to obtain the choice probability function all what we need is the the functions $\tilde{\varphi}$ and $F_{\tilde{\varepsilon}}$ and the discount factor. Given this information we can construct the mapping Ψ and obtain the choice probability function as the fixed point of this mapping. Therefore, all what we need to obtain the counterfactual probability function P^* is the counterfactual function $\tilde{\varphi}^*$. A second important implication of this Proposition is that the function $\tilde{\varphi}$ does not depend on the agent's optimal behavior and therefore on the optimal choice probabilities. We show below that it is quite straightforward to calculate the counterfactual function $\tilde{\varphi}^*$. In particular, this counterfactual is just the sum of the factual function $\tilde{\varphi}$ and a function that only depends on the utility change τ . Therefore, all what we need to identify $\tilde{\varphi}^*$ and P^* is to identify $\tilde{\varphi}$.

We now prove the identification of the function $\tilde{\varphi}$ when the distribution $F_{\tilde{\varepsilon}}$ is known. Given $F_{\tilde{\varepsilon}}$ it is clear from equation (22) that the surplus function is identified from the choice probabilities. Equation (21) defines implicitly $\delta(0, \cdot, P)$ as the unique fixed point of a

contraction mapping. Since the surplus function G is identified, it is clear that $\delta(0, \cdot, P)$ is also identified. Furthermore, $\delta(1, \cdot, P)$ is $\beta \int (G(x') + \delta(0, x', P)) dF_x(x'|1, x)$ and therefore it is also identified. Finally, the optimal choice probability function is such that $P(x) = F_{\tilde{\varepsilon}}(\tilde{\varphi}(x) + \tilde{\delta}(x, P))$. Given that $F_{\tilde{\varepsilon}}(\cdot)$ is invertible and that $\tilde{\delta}(\cdot, P) = \delta(1, \cdot, P) - \delta(0, \cdot, P)$ has been identified, we can identify $\tilde{\varphi}(\cdot)$ as:

$$\tilde{\varphi}(x) = F_{\tilde{\varepsilon}}^{-1}(P(x)) - \tilde{\delta}(x, P) \quad (23)$$

The identified functions $\delta(0, \cdot)$, $\delta(1, \cdot)$, and $\tilde{\varphi}(\cdot)$ depend both on agents' preferences and on agents' beliefs. Can we separately identify preferences and beliefs? No, without further restrictions. An assumption that identifies the utility function is the "normalization" $u(0, x) = 0$ for any $x \in X$. Under this assumption we have that $\tilde{\varphi}(x) = u(1, x)$. This type of "normalization" is innocuous in static models because it does not affect the estimation of counterfactual probabilities, which only depend on utility differences and not on utility levels. However, this normalization is not innocuous in dynamic models. In dynamic models, the counterfactual choice probabilities depend on utility levels and not only on utility differences. Therefore, if the true utility $u(0, x)$ is not zero but we estimate the counterfactual probabilities under the assumption that it is zero, our estimates will be inconsistent and in some cases they can be very seriously biased.

Proposition 3 shows that given the distribution function $F_{\tilde{\varepsilon}}$ we can identify the counterfactual choice probability function from the factual choice probabilities.

PROPOSITION 3: Suppose that the discount factor β , the distribution function $F_{\tilde{\varepsilon}}$, and the optimal choice probability function P are known. Then, the counterfactual choice probability function P^ is identified. In particular, P^* is the unique fixed point of the mapping $\Psi^*(P)$, where*

$$\Psi^*(P)(x) \equiv F_{\tilde{\varepsilon}} \left(\tilde{\varphi}(x) + T(1, x) - T(0, x) + \tilde{\delta}(x, P) \right) \quad (24)$$

The functions $\tilde{\varphi}$ and $\tilde{\delta}$ are the same as in the "factual" mapping $\Psi(P)$ and they are identified from the factual choice probabilities. The functions $T(0, \cdot)$ and $T(1, \cdot)$ only depend on the change in the current utility function τ and they can be obtained using the equation:

$$T(a, x) = \tau(a, x) + \beta \int T(0, x') dF_x(x'|a, x) \quad (25)$$

3.2.2 Identification with unknown $F_{\tilde{\varepsilon}}$

We now consider identification of counterfactual probabilities when the distribution function $F_{\tilde{\varepsilon}}$ is unknown and it is nonparametrically specified. We consider the same specification of the current utility as in Assumption 2 for the identification of the static model. However, in

the dynamic model we need additional conditions for the identification of the distribution of unobservables. The reason is that the state variable w affects the optimal choice probability not only through the current utility function but also through the expected discounted value of future utilities. Knowledge of the function $\tilde{y}(w, z)$ is not sufficient to identify $F_{\tilde{\varepsilon}}$ because the effect of w in the optimal probability function (through future expected utilities) depends on the function $\tilde{c}(\cdot)$ that is unknown. Assumption 3 establishes a condition that guarantees that this effect does not exist and that the effect of w on current and future utilities can be obtained given the functions $\tilde{y}(\cdot, \cdot)$ and $F_{\tilde{\varepsilon}}$.

ASSUMPTION 3: The conditional transition probability of $x = (w, z)$ factors as:

$$F_x(x'|a, x) = F_w(w'|w) F_z(z'|a, z) \quad (26)$$

That is, the state variable w is strictly exogenous with respect to the decision variable a . Furthermore, for any w_0, w_1 and w' with $w_0 < w_1$, we have that $F_w(w'|w_0) \geq F_w(w'|w_1)$.

Under Assumptions 2 and 3 we can decompose the function $\varphi(a, x)$ in two components,

$$\varphi(a, w, z) = Y(a, w, z) + C(a, z) \quad (27)$$

where the functions Y and C are implicitly defined by the recursive expressions:

$$\begin{aligned} Y(a, w, z) &= y(a, w, z) + \beta \int Y(0, w', z') dF_w(w'|w) dF_z(z'|a, z) \\ C(a, z) &= c(a, z) + \beta \int C(0, z') dF_z(z'|a, z) \end{aligned} \quad (28)$$

Given that the econometrician knows the outcome function $y(\cdot, \cdot, \cdot)$, the function $Y(\cdot, \cdot, \cdot)$ is identified. Thus, the optimal choice probability function is:

$$P(w, z) \equiv F_{\tilde{\varepsilon}} \left(\tilde{Y}(w, z) + \tilde{C}(z) + \tilde{\delta}(w, z, P) \right) \quad (29)$$

where the function $\tilde{\delta}$ is the same as in Propositions 2 and 3. Furthermore, Assumptions 2 and 3 imply that the function $\tilde{Y}(w, z) + \tilde{\delta}(w, z, P)$ is strictly increasing in w .

If the function $\tilde{\delta}$ were known, then the proof of identification of the probability distribution $F_{\tilde{\varepsilon}}$ would be very similar to the one in Proposition 1 for the static model. However, $\tilde{\delta}$ depends on $F_{\tilde{\varepsilon}}$ that is the object that we want to identify. Therefore, there is a recursive problem: we need to know $\tilde{\delta}$ to identify $F_{\tilde{\varepsilon}}$ but we need to know $F_{\tilde{\varepsilon}}$ to obtain $\tilde{\delta}$. If this recursive problem has a unique fixed point, then the distribution of unobservable state variables is nonparametrically identified.

To show that $F_{\tilde{\varepsilon}}$ is identified we proceed in the following way. The differential value function is $\tilde{v}(w, z) = \tilde{Y}(w, z) + \tilde{C}(z) + \tilde{\delta}(w, z)$. This function is unknown both because \tilde{C}

is unknown and because $\tilde{\delta}$ depends on the distribution $F_{\tilde{z}}$. However, we show that \tilde{v} can be described as the unique fixed point of a mapping that depends only on the data and the discount factor β . This mapping is $\Lambda(\tilde{v}) \equiv \{\Lambda(\tilde{v})(x) : x \in X\}$ with

$$\Lambda(\tilde{v})(x) = \left(\tilde{Y}(w, z) - \tilde{Y}(m(z), z) \right) - \left(\tilde{\delta}(w, z; \tilde{v}) - \tilde{\delta}(m(z), z; \tilde{v}) \right) \quad (30)$$

where: (1) $m(\cdot)$ is a function from Z into W such that $m(z)$ is the value $w \in W$ that makes $P(w, z) = 0.5$; and (2) the function $\tilde{\delta}$ has the same definition as above but we emphasize its dependence on \tilde{v} by including this function as an argument. More importantly, we can represent the surplus function and therefore $\tilde{\delta}$ as a function of \tilde{v} and P instead of as a function of P and $F_{\tilde{z}}$ as in equation (22). That is, under Assumption 3 we have that the surplus function can be written as:

$$G(x, \tilde{v}) = P(w, z) \tilde{v}(w, z) - \int_{-\infty}^w \tilde{v}(u, z) \frac{\partial P(u, z)}{\partial w} du \quad (31)$$

Proposition 4 proves these results, it shows that Λ is a contraction mapping, and it finally establishes that the functions $F_{\tilde{z}}$ and $\tilde{\varphi}$ are nonparametrically identified.

PROPOSITION 4: Suppose that Assumptions 1, 2 and 3 hold and that the discount factor β , the function \tilde{Y} , and the choice probability function P are known. Then, the function $\tilde{\varphi} = \tilde{Y} + \tilde{C}$ is identified on X , and the probability distribution $F_{\tilde{z}}$ is identified on $\tilde{v}(X) \equiv \{\tilde{v}(x) : x \in X\}$.

Proposition 5 shows that the counterfactual probability function is nonparametrically identified and describes the procedure to compute this function.

PROPOSITION 5: Suppose that Assumptions 1, 2 and 3 hold and that the discount factor β , the function \tilde{Y} , and the choice probability function P are known. Then, the counterfactual probability function P^ is identified. More specifically, P^* is the unique fixed point of the mapping Ψ^* defined in Proposition 3, where the functions $\tilde{\varphi}$ and $F_{\tilde{z}}$ that appear in the definition of this mapping have been identified as described in Proposition 4.*

4 Estimation method

This section presents a nonparametric procedure for the estimation of counterfactual choice probabilities that is based on the previous identification results. Suppose that we have a random sample $\{a_i, w_i, z_i : i = 1, 2, \dots, n\}$. First, we use the recursive formula in equation (28) to obtain first the function $Y(0, x)$, and then the functions $Y(1, x)$, and $\tilde{Y}(x)$. We distinguish five steps in this method.

Step 1: Estimation of the choice probability function P . We use a Nadaraya-Watson (kernel) estimator because it guarantees the continuous differentiability of the estimated function. To guarantee the strict monotonicity with respect to w we use the method in Hall and Huang (2001). These authors propose a simple method for monotonicizing kernel-type estimators like the Nadaraya-Watson estimator.

Step 2: Estimation of the function m . The estimated function $\hat{m} : Z \rightarrow W$ is defined as the value of w that solves the equation $\hat{P}(w, z) = 0.5$ for given $z \in Z$. We use Newton's method to find $\hat{m}(z)$. That is, $\hat{m}(z)$ is the limit of the sequence defined by the iterative formula:

$$w_{K+1} = w_K - \frac{\partial \hat{P}(w_K, z)}{\partial w} \left(\hat{P}(w_K, z) - 0.5 \right) \quad (32)$$

Given the strict monotonicity of our estimator $\hat{P}(\cdot, z)$, Newton's method always converges to $\hat{m}(z)$.

Step 3: Estimation of the functions \tilde{v} and \tilde{C} . Given \tilde{Y} , β and the estimated functions \hat{P} and \hat{m} we can construct a consistent estimator of the mapping Λ as defined in equation (30). Then, our estimator of \tilde{v} is the unique fixed point of the contraction mapping $\hat{\Lambda}$. Given this estimator \hat{v} , it is straightforward to obtain a consistent estimator of the function \tilde{C} as:

$$\hat{C}(z) = -\tilde{Y}(\hat{m}(z), z) - \tilde{\delta}(\hat{m}(z), z; \hat{v}) \quad (33)$$

Step 4: Estimation of the distribution function $F_{\tilde{\varepsilon}}$. For any real value $v \in \tilde{v}(X)$ the estimator of $F_{\tilde{\varepsilon}}(v)$ is:

$$\hat{F}_{\tilde{\varepsilon}}(v) = \frac{1}{n} \sum_{i=1}^n \hat{P}(w^*(v, z_i), z_i) \quad (34)$$

where $w^*(v, z)$ is the inverse function of $\hat{v}(w, z)$ with respect to w . That is, $w^*(v, z)$ is a function that provides the value of w such that $\hat{v}(w, z) = v$. We use also Newton's method to find $w^*(v, z)$, i.e., we iterate until convergence in the formula

$$w_{K+1} = w_K - \frac{\partial \hat{v}(w_K, z)}{\partial w} (\hat{v}(w_K, z) - v) \quad (35)$$

Again, the strict monotonicity of $\hat{v}(\cdot, z)$ guarantees that Newton's method always converges to $w^*(v, z)$.

Step 5: Estimation of P^ .* Given the estimated functions $\hat{\varphi} = \tilde{Y} + \hat{C}$ and $\hat{F}_{\tilde{\varepsilon}}$ and the discount factor β , we can construct a consistent estimator of the mapping Ψ^* . Let $\hat{\Psi}^*$ be this estimator. Then, our estimator of the choice probability function P^* is the unique fixed point of the mapping $\hat{\Psi}^*$.

The main computational cost in this procedure comes from the computation of the fixed points of the contraction mappings $\hat{\Lambda}$ and $\hat{\Psi}^*$. This cost is equivalent to solving the dynamic

programming problem twice. It is of the same order of magnitude as estimating a parametric version of the model using the two-step method in Hotz and Miller (1993), or the nested pseudo likelihood algorithm in Aguirregabiria and Mira (2002). In the implementation of this method in next section, we discretize the state space, including w , to solve the fixed point problems. To give an idea of the simplicity of the method, for a model with two state variables, 10,000 cells in the state space, and 1,000 observations, the CPU time of the whole method was less than six seconds using a program written in GAUSS language and an Intel Pentium processor of 2.2MHz. Though the computational cost increases exponentially with the number of cells in the state space, it is clear that we can use this method for any dynamic programming model that we can solve once in a reasonable amount time.

We do not derive in this paper the asymptotic distribution of our estimator of P^* . However, this estimator is consistent under standard regularity conditions. The Nadaraya-Watson estimator of P is consistent, and the estimators in steps 2 to 5 (i.e., \hat{m} , $\hat{\Lambda}$, \tilde{v} , \hat{C} , $\hat{F}_{\tilde{z}}$, w^* , $\hat{\Psi}^*$, and \hat{P}^*) are continuous and differentiable functions of the estimator \hat{P} . Therefore, all these estimators are consistent. The derivation of the rate of convergence of $(\hat{P}^* - P^*)$ and of sufficient conditions for its asymptotic normality is a more complicated problem that we do not consider in this problem. In any case, the computation of a consistent estimator of the asymptotic variance is complicated. Furthermore, it is likely that this asymptotic variance is not a good approximation to the finite sample variance for the range of sample sizes that we have in actual applications. Therefore, bootstrap is probably the most convenient approach to estimate the variance of this estimator.

5 Monte Carlo Experiment

This section presents a Monte Carlo experiment where we apply the previous nonparametric method to evaluate the effects of a subsidy to early machine replacement using the model described in section 2.4. The main purpose of this experiment is to analyze the finite sample bias and variance of the estimator using the type data that we can find in actual applications.

5.1 Experiment Design

Table 1 describes the design of the Monte Carlo experiment. The first panel in this table presents the form of the structural functions. These parameters have been chosen to generate an age distribution and a replacement probability function that resembles the empirical results in Cooper, Haltiwanger and Power (1999).⁵ Figures 1 and 2 present the probability

⁵These authors study plants' investment in machinery and equipment using data of US manufacturing plants from the Longitudinal Research Database. They define "machine replacement" as an investment rate

of replacement and the steady-state distribution of age for our model before the policy intervention. The average probability of replacement is 19%, the average age of capital is 2.42 years, and only 2.5% of the replacements occur at ages above 5 years of age.

The second panel in Table 1 presents the policy that we want to evaluate. This policy has the form described in equation (12). It is a subsidy to firms that replace their machine at ages between 3 and 5 years. The maximum subsidy is obtained when replacement occurs at the third year, and the amount of subsidy is 12% of the price of a new machine (i.e., $0.12 * p_{BUY}$). The subsidy decreases when the machine gets older such that it is 8% of the price of a new machine in the fourth year, 4% in the fifth year, and no subsidy for older ages. Figures 3 and 4 present the effect of this policy on the probability of capital replacement and on the age distribution. The probability of replacement decreases at ages lower than 3 years because the new policy creates an incentive to delay replacement in order to get the subsidy. However, the policy increases the probability of retirement at any age greater or equal than 3 years. Overall the policy encourages early machine replacement. The average age of a machine in steady state goes from 2.42 years to 2.07 years. Aggregate productivity in steady state increases by 2.69%.

In the Monte Carlo experiment we use 1,000 replications of a cross section of 1,000 firms. To construct each of these cross sections we take random draws of (ω, Age) from the joint steady-state distribution of these variables. Given these draws we generate the replacement decisions by taking random draws from a Bernoulli with probability $P(\omega, Age)$. Figures 5 and 6 present frequency estimates of the age distribution and the probability of replacement conditional on age for a typical sample in this Monte Carlo experiment. The 95% confidence bands illustrate that conditional on the machine age there is still very significant heterogeneity or sample variability in firms' replacement decisions.

5.2 Some estimation issues

All the functions have been estimated over a discrete and finite set of values of the state variables. The range of values for the age of capital consists of the integers between 0 and the maximum age observed in the sample plus one, i.e., typically the integers between 0 and 11. For the productivity shock we consider a uniform grid with 600 cells over the interval $[-5s_\omega, 5s_\omega]$, where s_ω is the sample standard deviation of the productivity shock. In a sample of 1,000 observations we typically find several observations out of the interval $[-3s_\omega, 3s_\omega]$. However, it is very unlikely to observe values greater in absolute value than four times the standard deviation. Therefore, we do some extrapolation outside the support of ω in the

of 20% or above.

sample. As we explain below, we need this extrapolation to estimate the function m . We exploit the continuity and monotonicity of P with respect to ω to do this extrapolation.

Figure 7 presents the function $\tilde{Y}(\omega, Age)$. This function represents the present value of output if the machine is replaced today and then it is never replaced again in the future, minus the present value of output if the machine is not replaced neither today or in the future. Given the specification of the production technology, the function \tilde{Y} is strictly increasing in both arguments.

Following the description of the estimation procedure in section 4, the first step consists in the estimation of the choice probability function $P(\omega, Age) \equiv E(a_i | \omega_i = \omega, Age_i = Age)$ using a Nadaraya-Watson estimator. For the sake of computational simplicity we use Silverman's rule of thumb to choose the value of the bandwidths for the two conditioning variables. Given the Nadaraya-Watson estimates, we impose monotonicity with respect to ω and Age using the method in Hall and Huang (2001). Figure 8 presents the true and the estimated replacement probability function for three different values of the productivity shock. We also report bootstrapped confidence intervals. There are two important features to comment on these estimates. First, note that the amplitude of the confidence intervals here is smaller than in the frequency estimates in Figure 6. There are three factors that contribute to the narrowing of the confidence intervals: (1) we are conditioning not only on age but also on the productivity shock and this reduces the variance of the residuals; (2) the kernel estimator imposes smoothness and this reduces the variance of the estimator, though it also introduces a finite sample bias; and (3) imposing monotonicity also reduces the variance of the estimator. A second important feature of these estimates is that they are biased for values of age lower than 2 or greater than 6. The true replacement probability is outside the 95% confidence interval for these values of age. The reason of this bias is that our criterion for the choice of bandwidth (i.e., Silverman's rule of thumb) introduces over-smoothing in our estimates. Though we could eliminate this over-smoothing by using a cross-validation method to choose the bandwidth, we prefer to analyze the performance of our estimator when we have certain over-smoothing because this is a common scenario in kernel estimation.

Figure 9 presents the steady-state age distribution induced by the Nadaraya-Watson estimate of the choice probability function.⁶ We also report the true age distribution and bootstrapped confidence intervals. This estimate is also significantly more precise than the raw frequency estimate in Figure 5. Note that the bias in \hat{P} does not have important effects on the estimation of the steady-state age distribution.

In the second step of the procedure, we use \hat{P} to estimate the function m . It is in this step

⁶We use the recursive expression in (13) to obtain the joint steady distribution $H(\omega, Age)$ and then we integrate over ω .

where we need estimates of the P function for values of ω outside the range of variation in a typical sample. The reason is that for $Age > 8$ we need very large values of the productivity shock in order to have $\hat{P}(\omega, Age) = 0.5$. Similarly, for $Age < 3$ the values of ω that solve the equation $\hat{P}(\omega, Age) = 0.5$ are very small. Figure 10 presents the true and the estimated m function for a typical sample.

5.3 Results of the experiment

The main results of the Monte Carlo experiment are presented in Figures 11 to 14 and in Table 2. The two estimated functions that play the most important role in the estimation of the policy effects are \tilde{C} and $F_{\tilde{\varepsilon}}$. Figure 11 presents the true function \tilde{C} and the quantiles 2.5%, 50%, and 97.5% in the Monte Carlo distribution of the estimated function. Though the estimation is precise, there is finite sample bias. This bias might be due to over-smoothing in the estimation of the replacement probability, but it might be also a more general property of our estimator that appears even we use cross-validation in the estimation of P . That is, the estimator \hat{C} is a very nonlinear function of the estimator \hat{P} , and though it is consistent it can be biased for relatively small samples. Figure 12 presents the Monte Carlo distribution of the estimate of the standard deviation of $\tilde{\varepsilon}$. There is also a very clear upward bias in this estimation.

Though, by construction, the upward biases in the estimates of \tilde{C} and $F_{\tilde{\varepsilon}}$ compensate with each other in the estimation of the factual choice probability P , that is not necessarily the case in the estimation of the counterfactual P^* . Figures 13 and 14 show that the estimator of the counterfactual (steady-state) age distribution is not biased. The estimated counterfactual is close to the true one, both in terms of median and of dispersion.

Table 2 presents more aggregate measures of the estimated policy effects on average age and productivity. The factual average age of capital and the average productivity are estimated without any bias, but there is a small upward bias in the estimated counterfactuals. Despite this bias, the estimated effects are very precise and close to the true effects.

6 Summary and Conclusions

This paper presents a nonparametric approach to evaluate the behavioral effects of counterfactual policies using dynamic discrete decision models. Though agents' preferences cannot be nonparametrically identified in this class of models, we show that the behavioral effects of counterfactual changes in preferences are identified under similar conditions as in static models. Our results apply both to finite horizon and infinite horizon decision processes. Based on this identification result we propose a nonparametric procedure to estimate the

behavioral effects of counterfactual experiments in this class of models. The computational cost of this method is equivalent to solving the dynamic programming problem twice (i.e., before and after the policy change). We have analyzed the ability of this method to provide informative estimates of policy effects using a Monte Carlo experiment. In this experiment we evaluate the effects of an investment subsidy program in the context of a model of machine replacement. Using a sample with 1,000 observations we find a small finite sample bias in the estimates of some policy effects. Despite this bias, the method provides precise estimates of the actual policy effects.

APPENDIX

PROOF OF PROPOSITION 1.

[1] Given the form of the utility function, the optimal choice probability is:

$$P(w, z) = F_{\tilde{\varepsilon}} (\tilde{y}(w, z) + \tilde{c}(z))$$

The function $P(., .)$ is identified on X . For $z \in Z$, let $m(z)$ be the value $w \in W$ that solves the equation $P(w, z) = 0.5$. Assumption 2(iii) implies that $m(z)$ exists and is unique for any $z \in Z$. Therefore, identification of $P(., .)$ on X implies the identification of $m(.)$ on Z . Given that the median of $\tilde{\varepsilon}$ is zero, we have that for any $z \in \bar{Z}$:

$$\tilde{y}(m(z), z) + \tilde{c}(z) = 0$$

or $\tilde{c}(z) = -\tilde{y}(m(z), z)$. Therefore, the function $\tilde{c}(.)$ is identified on Z , and the function $\tilde{u}(.) = \tilde{y}(., .) + \tilde{c}(.)$ is identified on X .

[2] Now, we prove the identification of $F_{\tilde{\varepsilon}}$. For $z \in Z$, define $\tilde{u}(W \times \{z\}) \subseteq \mathbb{R}$ as the range of possible values that the function $\tilde{u}(.)$ can take over the space $W \times \{z\}$. For $z \in Z$ and $e \in \tilde{u}(W \times \{z\})$, let $w(e, z)$ be the value $w \in W$ that solves the equation $\tilde{u}(w(e, z), z) = e$. Since $\tilde{u}(., .)$ is strictly increasing in w , the value $w(e, z)$ exists and it is unique. Furthermore, it is identified given $\tilde{u}(., .)$. Thus, we can identify $F_{\tilde{\varepsilon}}$ on $\tilde{u}(W \times \{z\})$ as:

$$F_{\tilde{\varepsilon}}(e) = F_{\tilde{\varepsilon}} (\tilde{u}(w(e, z), z)) = P (w(e, z), z)$$

We can use different values of z to identify $F_{\tilde{\varepsilon}}$ over the whole set $\tilde{u}(W \times Z)$.

[3] The counterfactual choice probability function is $P^*(x) = F_{\tilde{\varepsilon}} (\tilde{u}(x) + \tilde{\tau}(x))$. Given that we have identified $\tilde{u}(.)$ on X and $F_{\tilde{\varepsilon}}(.)$ on $\tilde{u}(X)$, it is clear that we can obtain $P^*(x)$ at any value x such that $\tilde{u}(x) + \tilde{\tau}(x) \in \tilde{u}(X)$.

PROOF OF PROPOSITION 2. First, we show that the differential value function $\tilde{v}(x)$ can be written as $\tilde{\varphi}(x) + \tilde{\delta}(x, P)$.

[1] Remember that by definition:

$$v(a, x) \equiv u(a, x) + \beta \int \max_{j \in A} \{ v(j, x') + \varepsilon'(j) \} dF_{\varepsilon}(\varepsilon') dF_x(x'|a, x)$$

Given the definition of the surplus function in the enunciation of this Proposition, we have that:

$$\int \max_{j \in A} \{ v(j, x') + \varepsilon'(j) \} dF_{\varepsilon}(\varepsilon') = v(0, x') + G(x')$$

Solving this expression in the first equation that defines the conditional choice value function we have that:

$$v(a, x) = u(a, x) + \beta \int v(0, x') dF_x(x'|a, x) + \beta \int G(x') dF_x(x'|a, x)$$

We can apply the same decomposition to the value $v(0, x')$ that appears in this expression. If we do this, we get:

$$\begin{aligned} v(a, x) &= u(a, x) + \beta \int u(0, x') dF_x(x'|a, x) + \beta^2 \int v(0, x'') dF_x(x''|0, x') dF_x(x'|a, x) \\ &+ \beta \int G(x') dF_x(x'|a, x) + \beta^2 \int G(x'') dF_x(x''|0, x') dF_x(x'|a, x) \end{aligned}$$

If we continue applying the decomposition to $v(0, x'')$, $v(0, x''')$, and so on, we get:

$$\begin{aligned} v(a, x) &= u(a, x) + \sum_{j=1}^{\infty} \beta^j \left[\int u(0, x^j) \left(\prod_{i=2}^j dF_x(x^i|0, x^{i-1}) \right) dF_x(x'|a, x) \right] \\ &+ \sum_{j=1}^{\infty} \beta^j \left[\int G(x^j) \left(\prod_{i=2}^j dF_x(x^i|0, x^{i-1}) \right) dF_x(x'|a, x) \right] \end{aligned}$$

The second term in the right hand side is the present value of choosing alternative 0 forever in the future given that the current choice and the current state are a and x , respectively. The third term in the right hand side is the difference between the value of choosing always the optimal alternative and the value of choosing always alternative 0. Then, given the definitions of $\varphi(a, x)$ and $\delta(a, x)$ in the enunciate of the Lemma, it is clear that $v(a, x) = \varphi(a, x) + \delta(a, x)$ where:

$$\varphi(a, x) = u(a, x) + \sum_{j=1}^{\infty} \beta^j \left[\int u(0, x^j) \left(\prod_{i=2}^j dF_x(x^i|0, x^{i-1}) \right) dF_x(x'|a, x) \right]$$

and

$$\delta(a, x) = \sum_{j=1}^{\infty} \beta^j \left[\int G(x^j) \left(\prod_{i=2}^j dF_x(x^i|0, x^{i-1}) \right) dF_x(x'|a, x) \right]$$

Given these expressions for $\varphi(a, x)$ and $\delta(a, x)$, it is straightforward to show that:

$$\varphi(a, x) = u(a, x) + \beta \int \varphi(0, x') dF_x(x'|a, x)$$

and

$$\delta(a, x) = \beta \int (G(x') + \delta(0, x')) dF_x(x'|a, x)$$

Taking into account the definition of the surplus function as $\int \max\{0; \tilde{v}(x) - \tilde{\varepsilon}\} dF_{\tilde{\varepsilon}}(\tilde{\varepsilon})$, and that $\tilde{v}(x) = F_{\tilde{\varepsilon}}^{-1}(P(x))$, we have that this surplus function can be written as:

$$\begin{aligned} G(P(x)) &= P(x) (\tilde{v}(x) - E(\tilde{\varepsilon}|\tilde{\varepsilon} \leq \tilde{v}(x))) \\ &= P(x) F_{\tilde{\varepsilon}}^{-1}(P(x)) - \int_{-\infty}^{F_{\tilde{\varepsilon}}^{-1}(P(x))} \tilde{\varepsilon} dF_{\tilde{\varepsilon}}(\tilde{\varepsilon}) \end{aligned}$$

[2] Thus, $v(a, x) = \varphi(a, x) + \delta(a, x, P)$. This implies that we the expression $P(x) = F_{\tilde{\varepsilon}}(\tilde{v}(x))$ can be rewritten as:

$$P(x) = F_{\tilde{\varepsilon}}(\tilde{\varphi}(x) + \tilde{\delta}(x, P))$$

Therefore, the optimal choice probability function P is a fixed point of the mapping $\Psi(P)$ where $\Psi(P)(x) \equiv F_{\tilde{\varepsilon}}(\tilde{\varphi}(x) + \tilde{\delta}(x, P))$. This is a particular case of the fixed point probability mapping in Aguirregabiria and Mira (2002). Proposition 1(i) in Aguirregabiria and Mira (2002) shows that this mapping has a unique fixed point.

PROOF OF PROPOSITION 3. By Proposition 2, the counterfactual probability function P^* is the unique fixed point of the mapping $\Psi^*(P)$, where

$$\Psi^*(P)(x) \equiv F_{\tilde{\varepsilon}} \left(\tilde{\varphi}^*(x) + \tilde{\delta}^*(x, P) \right)$$

where $\tilde{\varphi}^*(x) = \varphi^*(1, x) - \varphi^*(0, x)$ and $\tilde{\delta}^*(x, P) = \delta^*(1, x, P) - \delta^*(0, x, P)$ are associated with the counterfactual utility function $u^*(a, x)$.

[1] **Identification of $\tilde{\delta}^*$.** Taking into account the definition of the functions $\delta(0, \cdot, \cdot)$ and $\delta(1, \cdot, \cdot)$ in Proposition 2, we can see that these functions depend on the probability distribution $F_{\tilde{\varepsilon}}$ and on the discount factor, but they do not depend on the current utility function. Therefore, $\tilde{\delta}^*(\cdot, \cdot) = \tilde{\delta}(\cdot, \cdot)$ and this function is identified as we have shown before.

[2] **Identification of $\tilde{\varphi}^*(\cdot)$:** Taking into account the expressions for $\varphi(a, x)$ in the proof of Proposition 2 we have that:

$$\varphi^*(a, x) = u^*(a, x) + \sum_{j=1}^{\infty} \beta^j \left[\int u^*(0, x^j) \left(\prod_{i=2}^j dF_x(x^i | 0, x^{i-1}) \right) dF_x(x' | a, x) \right]$$

Given that $u^*(a, x) = u(a, x) + \tau(a, x)$ we have that:

$$\begin{aligned} \varphi^*(a, x) &= u(a, x) + \sum_{j=1}^{\infty} \beta^j \left[\int u(0, x^j) \left(\prod_{i=2}^j dF_x(x^i | 0, x^{i-1}) \right) dF_x(x' | a, x) \right] \\ &+ \tau(a, x) + \sum_{j=1}^{\infty} \beta^j \left[\int \tau(0, x^j) \left(\prod_{i=2}^j dF_x(x^i | 0, x^{i-1}) \right) dF_x(x' | a, x) \right] \\ &= \varphi(a, x) + T(a, x) \end{aligned}$$

where $T(a, x)$ is the term associated with the values of the function $\tau(\cdot, \cdot)$. Therefore,

$$\tilde{\varphi}^*(x) = \tilde{\varphi}(x) + T(1, x) - T(0, x)$$

The function $T(0, x)$ can be obtained as the fixed point of the contraction mapping:

$$T(0, x) = \tau(0, x) + \beta \int T(0, x') dF_x(x' | 0, x)$$

Given $T(0, \cdot)$, we can obtain $T(1, x)$ as $\tau(1, x) + \beta \int T(0, x') dF_x(x'|1, x)$. Since the function $\tau(\cdot, \cdot)$ is known, it is clear that $T(0, \cdot)$ and $T(1, \cdot)$ are identified.

[3] Identification of P^* . Thus, P^* is the unique fixed point of the mapping $\Psi^*(P)$, where:

$$\Psi^*(P)(x) \equiv F_{\tilde{\varepsilon}} \left(\tilde{\varphi}(x) + T(1, x) - T(0, x) + \tilde{\delta}(x, P) \right)$$

We have shown that the functions $\tilde{\varphi}(\cdot)$, $T(1, \cdot)$, $T(0, \cdot)$, and $\tilde{\delta}(\cdot, \cdot)$ are identified. Thus, given $F_{\tilde{\varepsilon}}$, the counterfactual probability function is identified.

PROOF OF PROPOSITION 4. The proof proceeds in two steps. First, we show that the distribution function $F_{\tilde{\varepsilon}}$ is uniquely identified given $P(\cdot)$ and the differential value function $\tilde{v}(\cdot)$. Second, we show that, for given P , $\tilde{\varphi}$ and β , the differential value function \tilde{v} is the unique fixed point of a contraction mapping. Therefore, we conclude that there is a unique distribution function $F_{\tilde{\varepsilon}}$ that is consistent with given P , $\tilde{\varphi}$ and β .

[1] Identification of $F_{\tilde{\varepsilon}}$ for given P and \tilde{v} . Suppose that we know $P(x)$ and $\tilde{v}(x)$ for every $x \in X$. Define the set $\tilde{v}(X) \equiv \{\tilde{v}(x) : x \in X\}$. Under Assumption 3, the function \tilde{v} is continuous and strictly increasing in w . Therefore, there is an inverse function $w^*(v, z)$ such that, for any $v \in \tilde{v}(X)$ and for any $z \in Z$, we have that $\tilde{v}(w^*(v, z), z) = v$. The model implies that $P(w, z) = F_{\tilde{\varepsilon}}(\tilde{v}(w, z))$. Therefore, it is clear that for any $v \in \tilde{v}(X)$ we can obtain $F_{\tilde{\varepsilon}}(v)$ as $P(w^*(v, z), z)$.

[2] Identification of \tilde{v} for given P , $\tilde{\varphi}$ and β . The differential value function \tilde{v} solves the following fixed-point problem:

$$\tilde{v}(x) = \tilde{\varphi}(x) + \sum_{j=1}^{\infty} \beta^j \left[\int G(x^j; \tilde{v}) \left(\prod_{i=2}^j dF_x(x^i|0, x^{i-1}) \right) (dF_x(x'|1, x) - dF_x(x'|0, x)) \right]$$

where $G(x; \tilde{v})$ is the surplus function $\int \max_{j \in A} \{0; \tilde{v}(x) - \tilde{\varepsilon}\} dF_{\tilde{\varepsilon}}(\tilde{\varepsilon})$. Given this description of the surplus, the function \tilde{v} is a fixed point of a mapping that depends on $\tilde{\varphi}$, β and $F_{\tilde{\varepsilon}}$. However, we can also represent \tilde{v} as a fixed point of a mapping that depends on $\tilde{\varphi}$, β and P , but not on $F_{\tilde{\varepsilon}}$. First, notice that $G(x) = P(x) \tilde{v}(x) - \int_{-\infty}^{\tilde{v}(x)} \tilde{\varepsilon} dF_{\tilde{\varepsilon}}(\tilde{\varepsilon})$. Second, using the relationship $P(w, z) = F_{\tilde{\varepsilon}}(\tilde{v}(w, z))$ and the continuity and strict monotonicity of $P(w, z)$ with respect to w , we have that $\int_{-\infty}^{\tilde{v}(w, z)} \tilde{\varepsilon} dF_{\tilde{\varepsilon}}(\tilde{\varepsilon}) = \int_{-\infty}^w \tilde{v}(u, z) \frac{\partial P(u, z)}{\partial w} du$. Therefore, the surplus function can be written as:

$$G(x; \tilde{v}) = P(x) \tilde{v}(x) - \int_{-\infty}^w \tilde{v}(u, z) \frac{\partial P(u, z)}{\partial w} du$$

And this implies that \tilde{v} can be described as a fixed point of a mapping that depends on $\tilde{\varphi}$, β and P , but not on $F_{\tilde{\varepsilon}}$.

Let $\Lambda(\tilde{v}) \equiv \{\Lambda(\tilde{v})(x) : x \in X\}$ be this mapping. Now, we show that Λ is a contraction mapping and therefore it has a unique fixed point. To prove this we use Blackwell's sufficient conditions for a contraction (see Theorem 3.3 in Stockey and Lucas, 1989). These sufficient conditions are monotonicity and discounting.

(a) *Monotonicity:* We should prove that for any two functions \tilde{v}^0 and \tilde{v}^1 such that $\tilde{v}^1(x) - \tilde{v}^0(x) \geq 0$ for any $x \in X$, then $\Lambda(\tilde{v}^1)(x) - \Lambda(\tilde{v}^0)(x) \geq 0$ for any $x \in X$. Using the definition of the mapping Λ above, a sufficient condition for the second inequality is that $G(x; \tilde{v}^1) - G(x; \tilde{v}^0) \geq 0$ for any $x \in X$. Note that:

$$G(x; \tilde{v}^1) - G(x; \tilde{v}^0) = P(x) (\tilde{v}^1(x) - \tilde{v}^0(x)) - \int_{-\infty}^w (\tilde{v}^1(u, z) - \tilde{v}^0(u, z)) \frac{\partial P(u, z)}{\partial w} du$$

Solving by parts the integral, it is straightforward to show that:

$$G(x; \tilde{v}^1) - G(x; \tilde{v}^0) = \int_{-\infty}^w P(u, z) du \geq 0$$

(b) *Discounting:* We should prove that there exists some constant $\lambda \in [0, 1)$ such that for any function \tilde{v} , any constant c , and any $x \in X$ we have that $\Lambda(\tilde{v} + c)(x) \leq \Lambda(\tilde{v})(x) + \lambda c$. We start obtaining $G(x; \tilde{v} + c)$.

$$\begin{aligned} G(x; \tilde{v} + c) &= P(x) (\tilde{v}(x) + c) - \int_{-\infty}^w (\tilde{v}(u, z) + c) \frac{\partial P(u, z)}{\partial w} du \\ &= G(x; \tilde{v}) + c P(x) - c \int_{-\infty}^w \frac{\partial P(u, z)}{\partial w} du = G(x; \tilde{v}) \end{aligned}$$

Therefore, there is discounting in the surplus function. Furthermore, given the definition of Λ , it is clear that $\Lambda(\tilde{v} + c)(x) = \Lambda(\tilde{v})(x)$, i.e., there is discounting in the mapping Λ .

PROOF OF PROPOSITION 5. By Proposition 3, P^* is the unique fixed point of the mapping Ψ^* . This mapping depends on the known functions $T(1, \cdot) - T(0, \cdot)$, on the discount factor β , and on the functions $\tilde{\varphi}$ and $F_{\tilde{\varepsilon}}$. By Proposition 4, the functions $\tilde{\varphi}$ and $F_{\tilde{\varepsilon}}$ are identified given β , \tilde{Y} , and P . Therefore, the mapping Ψ^* and its unique fixed point P^* are identified.

References

- [1] Adda, J. and R. Cooper (2000): "Balladurette and Juppette: A Discrete Analysis of Scrapping Subsidies," *Journal of Political Economy*, 108, 778-806.
- [2] Aguirregabiria, V. (2005): "Nonparametric Identification of Behavioral Responses to Counterfactual Policy Interventions in Dynamic Discrete Decision Processes," *Economics Letters*, forthcoming.
- [3] Aguirregabiria, V. and C. Alonso-Borrego (1999): "Labor Contracts and Flexibility: Evidence from a Labor Market Reform in Spain," manuscript.
- [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] Altug, S. and R. Miller (1998): "The Effect of Work Experience on Female Wages and Labour Supply," *Review of Economic Studies*, 65, 45-85.
- [6] Berkovec, J. and S. Stern (1991): "Job Exit Behavior of Older Men," *Econometrica*, 59, 189-210.
- [7] Cooper, R., J. Haltiwanger and L. Power (1999): "Machine Replacement and the Business Cycle: Lumps and Bumps," *The American Economic Review*, 89, 921-946.
- [8] Cho, S. (2002): "An Empirical Model of Mainframe Computer Investment," manuscript, Yale University.
- [9] Das, M. (1992): "A Micro-econometric Model of Capital Utilization and Retirement: The Case of the Cement Industry," *Review of Economic Studies*, 59, 277-297.
- [10] Eckstein, Z. and I. Zilcha (1994): "The Effects of Compulsory Schooling on Growth, Income Distribution and Welfare," *Journal of Public Economics*, 54, 339-359.
- [11] Eckstein, Z. and K. Wolpin (1999): "Why Youth Drop out of High School: The Impact of Preferences, Opportunities and Abilities," *Econometrica*, 67, 1295-1339.
- [12] Ferrall, C. (1997): "Unemployment Insurance Eligibility and the Transition from School to Work in Canada and the United States," *Journal of Business and Economic Statistics*, 15, 115-129.

- [13] Hall, P. and L-S Huang (2001): "Nonparametric Kernel Regression Subject to Monotonicity Constraints," *Annals of Statistics*, 29, 624-647.
- [14] Heckman, J. and R. Robb (1985): "Alternative Methods for Evaluating the Impact of Interventions," in J. Heckman and B. Singer (eds.) *Longitudinal Analysis of Labor Market Data*. Cambridge University Press. Cambridge.
- [15] Heckman, J. and Smith, J. (1998): "Evaluating the Welfare State", Frisch Centenary, Econometric Monograph Series, Cambridge: Cambridge University Press.
- [16] Heckman, J., and E. Vytlacil (1999): "Local Instrumental Variables and Latent Variable Models for Identifying and Bounding Treatment Effects," *Proceedings of the National Academy of Sciences*, 96, 4730-4734.
- [17] Heckman, J. and S. Navarro (2005): "Dynamic Discrete Choice and Dynamic Treatment Effects," manuscript. The University of Chicago. Department of Economics.
- [18] Hotz, J., and R.A. Miller (1993): "Conditional choice probabilities and the estimation of dynamic models," *Review of Economic Studies*, 60, 497-529.
- [19] Keane, M. and R. Moffitt (1998): "A Structural Model of Multiple Welfare Program Participation and Labor Supply," *International Economic Review*, 39, 553-89.
- [20] Keane, M. and K. Wolpin (2000): "Estimating Welfare Effects Consistent with Forward-Looking Behavior. Part I: Lessons from a Simulation Exercise," PIER Working Paper 01-019. University of Pennsylvania.
- [21] Keane, M. and K. Wolpin (2000): "Estimating Welfare Effects Consistent with Forward-Looking Behavior. Part II: Empirical Results," PIER Working Paper 01-020. University of Pennsylvania.
- [22] Keane, M. and K. Wolpin (2001): "The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment," *International Economic Review*, 42, 1051-1103.
- [23] Kennet, M. (1993) "Did Deregulation Affect Aircraft Engine Maintenance? An Empirical Policy Analysis," *The RAND Journal of Economics*, 24, 542-558.
- [24] Kennet, M. (1994) "A Structural Model of Aircraft Engine Maintenance," *Journal of Applied Econometrics*, 9, 351-368.

- [25] Kasahara, H. (2004): "Temporary Increases in Tariffs and Machine Replacement: The Chilean Experience,' manuscript, Queen's University, Kingston, Canada.
- [26] Magnac, T. and Thesmar, D. (2002) "Identifying dynamic discrete decision processes". *Econometrica*, 70, 801-816.
- [27] Manski, C. (1989): "Anatomy of the Selection Problem," *Journal of Human Resources*, 24, 343-360.
- [28] Manski, C. (1990): "Nonparametric Bounds on Treatment Effects," *American Economic Review, Papers and Proceedings*, 80, 319-323.
- [29] Manski, C. (1991): "Nonparametric estimation of expectations in the analysis of discrete choice under uncertainty", in *Nonparametric and Semiparametric Methods in Econometrics and Statistics*, W. Barnett, J. Powell, and G. Tauchen (eds.). Cambridge University Press.
- [30] Manski, C. (1993): "Dynamic choice in social settings: Learning from the experience of others", *Journal of Econometrics* 58, pp. 121-36.
- [31] Miller, R. (1997): "Estimating Models of Dynamic Optimization with Microeconomic Data," in H. Pesaran and P. Smidh (eds.) *Handbook of Applied Econometrics: Microeconomics (Vol 2)*. Blackwell.
- [32] Pakes, A. (1986): "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, 54, 755-84.
- [33] Pakes, A. and M. Simpsons (1989): "Patent renewal data," in Bailey, Winston (eds.), *Brookings Papers on Economic Activity*. The Brookings Institute, Washington.
- [34] Rota, P. (2004): "Estimating Labor Demand with Fixed Costs," *International Economic Review*, 45, 25-48.
- [35] Rust, J. (1986): "When Is It Optimal to Kill Off the Market for Used Durable Goods?" *Econometrica*, 54, 65-86.
- [36] Rust, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, 55, 999-1033.
- [37] Rust, J. (1994): "Structural estimation of Markov decision processes," in R. E. Engle and McFadden (eds.) *Handbook of Econometrics Volume 4*, North-Holland.

- [38] Rust, J. and G. Rothwell (1995): "Optimal Response to a Shift in Regulatory Regime: The Case of the US Nuclear Power Industry," *Journal of Applied Econometrics*, 10, S75-S118.
- [39] Rust, J. and C. Phelan (1997): "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets," *Econometrica*, 65, 781-832.
- [40] Sanders, S. and R. Miller (1997): "Human Capital Development and Welfare Participation" Carnegie-Rochester Conference Series on Public Policy, 46.
- [41] Sturm, R. (1991) "A Structural Economic Model of Operating Cycle Management in European Nuclear Power Plants," manuscript, RAND Corporation.
- [42] Taber, C. (2000): "Semiparametric Identification and Heterogeneity in Discrete Choice Dynamic Programming Models," *Journal of Econometrics*, 96, 201-229.
- [43] Wolpin, K. (1996): "Public-Policy Uses of Discrete-Choice Dynamic Programming Models," *American Economic Review*, 86, 427-32.

Table 1
Design of the Monte Carlo Experiment

Production Function	$y(K_t, \omega_t) = K_t^\alpha \exp(\omega_t),$	with $\alpha = 0.95$
Stochastic Process ω_t	$\omega_t = \rho \omega_{t-1} + \eta_t,$	with $\rho = 0.8$ $\sigma_\eta = 0.2$
Transition capital	$K_t = a_t + (1 - a_t) (1 - \lambda) K_{t-1},$	with $\lambda = 0.1$
Maintenance costs	$mc(K_t) = (1 - a) \overline{mc} (Age_t + 1),$	with $\overline{mc} = 0.01$
Replacement costs	$rc(K_{t-1}) = p_{BUY} - p_{SELL} K_{t-1},$	with $p_{BUY} = 1.6$ $p_{SELL} = 0.4$
Distribution of $\tilde{\varepsilon}_t$	$\tilde{\varepsilon}_t$ is $N(0, \sigma_\varepsilon^2),$	with $\sigma_\varepsilon = 0.2$
Time discount factor		$\beta = 0.95$
Counterfactual policy	As in equation (12)	with $\tau_0 = 0.12 p_{BUY}$ $\tau_1 = 0.04 p_{BUY}$ <u>$Age = 3, \overline{Age} = 5$</u>
Sample size	Cross section of 1,000 firms	
Monte Carlo replications	Number of replications = 1,000.	

Table 2
Statistics from Monte Carlo Experiment

	True Value	Empirical Median Monte Carlo	95% Conf. Interval Monte Carlo
Avg. Age Before Policy	2.451	2.451	(2.349 , 2.556)
Avg. Age After Policy	2.086	1.973	(1.849 , 2.178)
Policy Effect on Avg. Age	-0.365	-0.473	(-0.594 , -0.283)
Avg. Productivity Before Policy	0.797	0.797	(0.790 , 0.803)
Avg. Productivity After Policy	0.823	0.831	(0.816 , 0.840)
Policy Effect on Avg. Productivity	0.026	0.033	(0.020 , 0.042)

Figure 1
Probability of Replacement Conditional on Age and Productivity: $P(\omega, \text{Age})$
True Model Before Policy Change




Figure 2
Steady-State Distribution of Age $H(\text{Age})$
True Model Before Policy Change

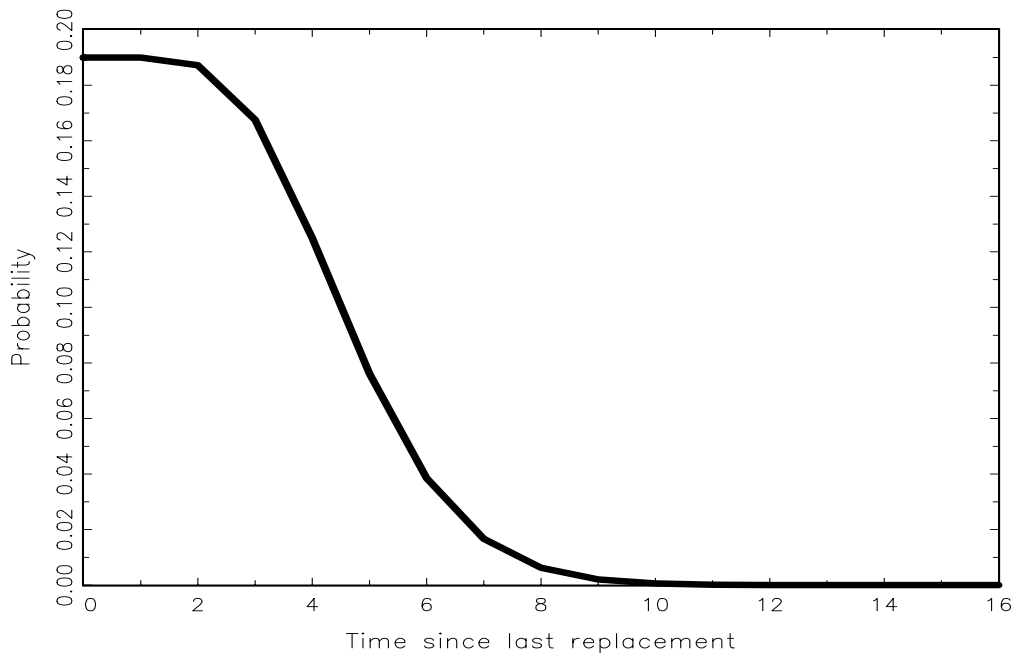


Figure 3
Probability of Replacement Conditional on Age: $\bar{P}(\text{Age})$
True Model Before and After Policy Change

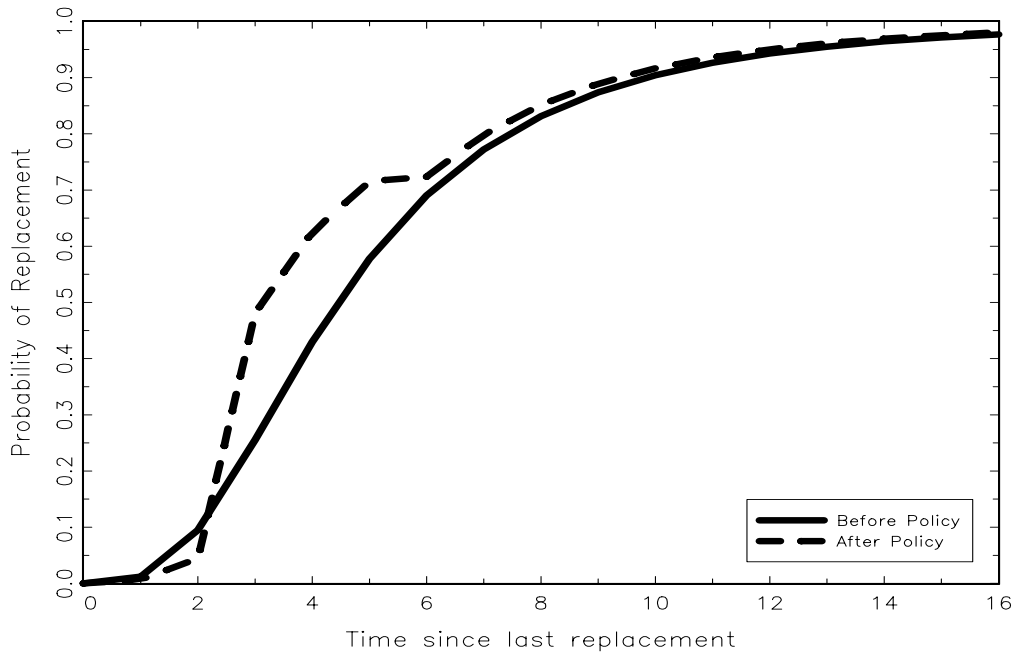


Figure 4
Steady-State Distribution of Age $H(\text{Age})$
True Model Before and After Policy Change




Figure 5
Frequency Estimate of Age Distribution
with 95% Confidence Bands (based on asymptotic distribution)




Figure 6
Frequency Estimate of Probability of Replacement Conditional on Age
with 95% Confidence Bands (based on asymptotic distribution)

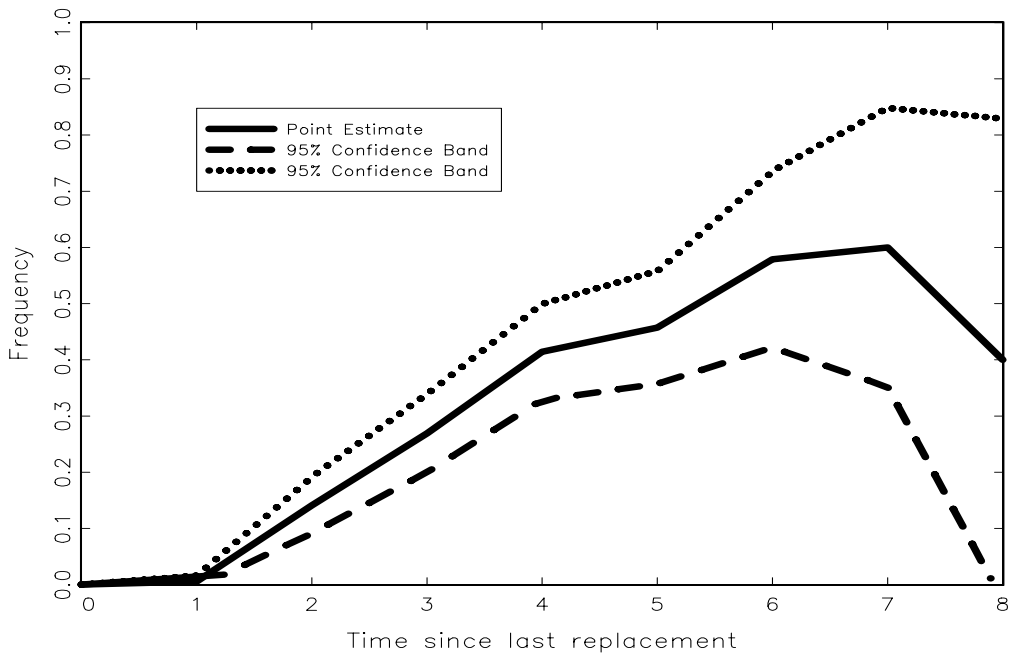


Figure 7
Function $\tilde{Y}(\text{Age}, \omega) = Y(1, \text{Age}, \omega) - Y(0, \text{Age}, \omega)$

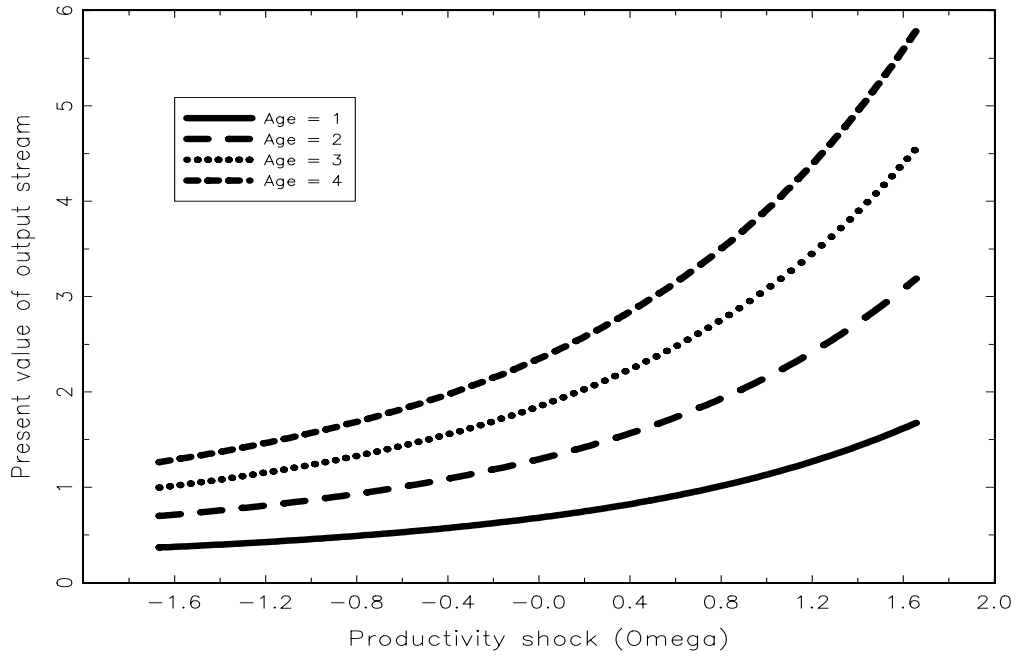




Figure 8

Nadaraya-Watson Estimates of Replacement Probability $\hat{P}(\omega, \text{Age})$
with Bootstrapped 95% Confidence Bands

$\omega = -0.514$ (Percentile 0.5%)



$\omega = 0.0$ (Percentile 50%)



$\omega = 0.514$ (Percentile 99.5%)

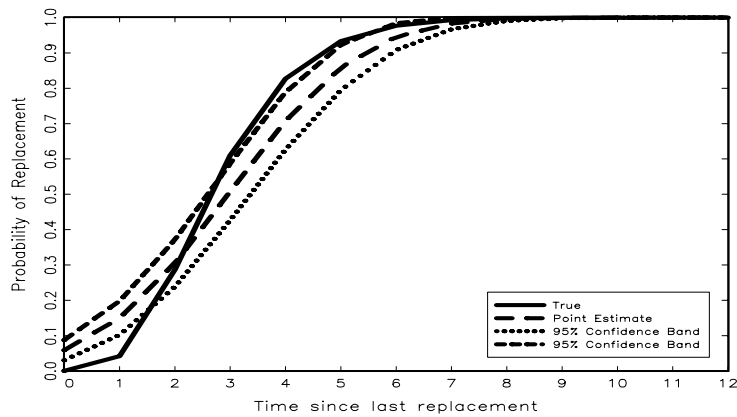


Figure 9
Estimated Steady-State Distribution of Age
with Bootstrapped 95% Confidence Bands

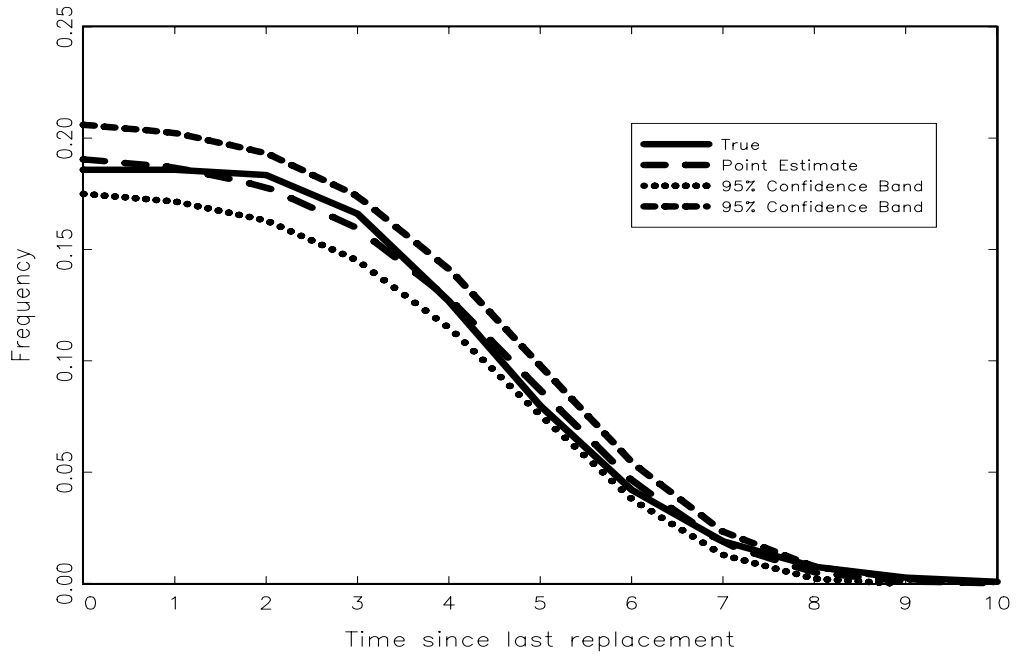


Figure 10
True and Estimated m Function
with Bootstrapped 95% Confidence Bands




Figure 11
 True and Estimated \tilde{C} Function
 Quantiles 2.5%, 50%, and 97.5% from Monte Carlo Distribution




Figure 12
 Monte Carlo Distribution of the Estimated
 Standard Deviation of $\tilde{\epsilon}$




Figure 13
Effect of Subsidy on Age Distribution
True Effect and Median Estimated Effect from Monte Carlo Distribution

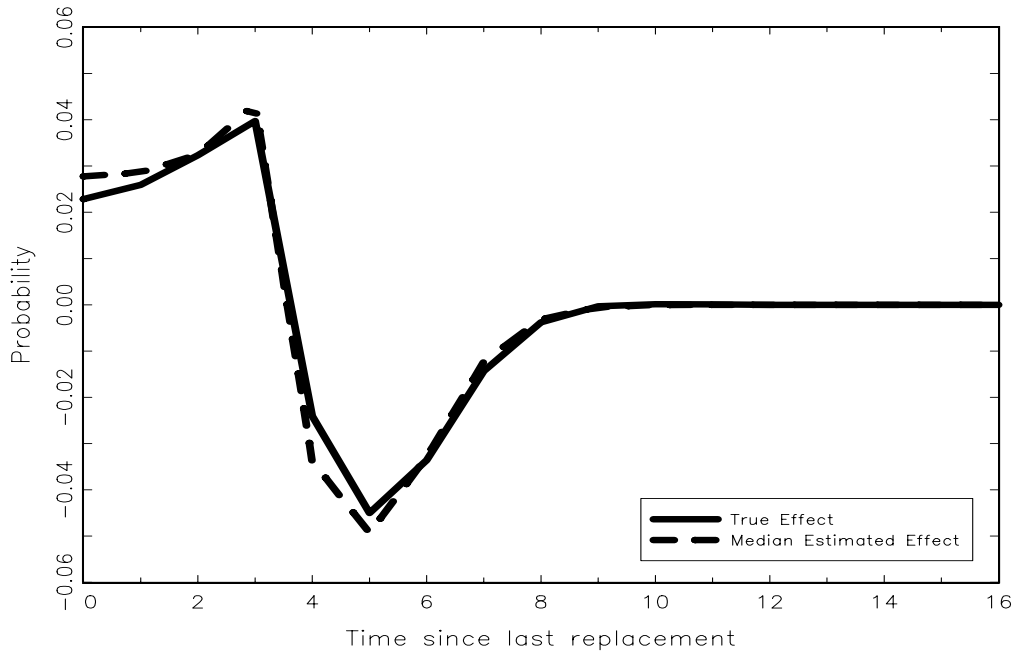


Figure 14
Effect of Policy on Age Distribution
Median, Quantile 2.5% and Quantile 97.5% from Monte Carlo Distribution

