

Another Look at the Identification of Dynamic Discrete Decision Processes: An Application to Retirement Behavior

Victor AGUIRREGABIRIA

Department of Economics, University of Toronto, Toronto, Ontario M5S 3G7, Canada
(victor.aguirregabiria@utoronto.ca)

This article deals with the estimation of behavioral and welfare effects of counterfactual policy interventions using dynamic structural models where all the primitive functions are nonparametrically specified (i.e., preferences, technology, transition rules, and distribution of unobserved variables). It proves the nonparametric identification of agents' decision rules, before and after the policy intervention, and of the change in agents' welfare. Based on these results, I propose a nonparametric procedure to estimate the behavioral and welfare effects of a class of counterfactual policy interventions. The nonparametric estimator can be used to construct a test of the validity of a parametric specification. I illustrate this method using a simple model of labor force retirement, panel data with information on public pension wealth, and a hypothetical reform that delays by three years the eligibility ages of the public pension system in Sweden.

KEY WORDS: Counterfactual policy interventions; Dynamic discrete decision processes; Nonparametric identification; Retirement behavior.

1. INTRODUCTION

Dynamic discrete choice structural models assume that agents are forward looking and maximize expected intertemporal payoffs. The structural functions to estimate in these models represent agents' preferences and beliefs about uncertain events. These functions are estimated under the principle of revealed preference and using microdata on individuals' choices and outcomes. Dynamic discrete choice structural models have proven useful tools for the assessment of public policy interventions, hypothetical or factual, such as unemployment insurance, social security pensions, patent regulation, educational policies, programs on child poverty, and scrapping subsidies, among many others. A common feature of the econometric models in these applications is the parametric specification of the structural functions. These parametric models contrast with the emphasis on robustness and nonparametric specification in the literature on evaluation of treatment effects (see Heckman and Robb 1985; Manski 1990, and, more recently, Heckman and Smith 1998 and Heckman and Vytlacil 1999, 2005). Though robustness is an important argument in favor of the treatment effects approach, that methodology cannot be used to evaluate counterfactual policies and it has limitations in measuring welfare effects and in allowing for transitional dynamics and general equilibrium effects. It is in this kind of problem where dynamic structural models can be particularly useful. The main purpose of this article is to determine conditions under which nonparametrically specified dynamic structural models can be used to estimate the effects of counterfactual policy interventions.

Rust (1994) and Magnac and Thesmar (2002) have obtained negative results on the identification of dynamic discrete structural models. These studies show that the principle of revealed preference can identify a value function that combines both preferences and beliefs, but it cannot identify preferences separately from beliefs, even when the researcher "knows" the time

discount factor and agents' beliefs. Based on this negative result, this article takes a different look at the problem of nonparametric identification of dynamic decision models. Instead of looking at the separate identification of preferences and beliefs, I study the identification of the behavioral and welfare effects of counterfactual policy changes. I prove the identification of agents' optimal decision rules and value functions associated with a class of hypothetical policy interventions. This approach is in the spirit of Marschak (1953) who was one of the first econometricians to argue that models that are not fully identified may contain subsets of parameters that are identified and can be used for policy analysis (see also Heckman 2000 for a discussion of Jacob Marschak's approach to structural econometrics and identification).

In the Rust–Magnac–Thesmar framework, the econometrician observes agents' decisions and some state variables for a random sample of agents. In this article, I assume that the econometrician also observes an outcome variable. This outcome variable is neither a decision nor a state variable, but it is a component of the utility function. Observing this outcome variable is key for the nonparametric identification of the distribution of the unobservables. However, this additional information does not solve the Rust–Magnac–Thesmar underidentification problem. In particular, the component of the utility function that is not the outcome variable cannot be identified separately from agents' beliefs.

A first contribution of this article is to show that, for a wide class of models and counterfactual policies, agents' behaviors before and after the policy intervention, and the change in agents' welfare, are nonparametrically identified. Three assumptions are key for this identification result. First, (a) the

outcome variable enters additively in the utility function, and the unobservables affecting this variable are (conditionally) independent of the other unobservables in the utility function. Second, (b) the transition probability function and the discount factor are not affected by the policy. And third, (c) the econometrician knows the difference between the utility functions before and after the policy change but does not know either of these two utility functions. I show that, under condition (a), a recursive application of a theorem put forward by Matzkin (1992, 1994) implies the identification of the distribution of the unobservables and of two value functions: the present value of behaving optimally minus the present value of some arbitrary behavior, and the present value of a one-period deviation from the arbitrary behavior. However, knowledge of these value functions is not sufficient to identify the one-period utility function or to identify the effects of a general class of counterfactual changes in the structural functions. Instead of imposing additional restrictions to identify the utility function, the article shows that knowledge of these value functions is sufficient to identify the effects of a relevant class of counterfactual policies. Conditions (b) and (c) characterize the class of counterfactual policies whose effects can be identified without further assumptions on the primitives of the model.

A second contribution of the article is that it proposes a non-parametric method to estimate the behavioral and welfare effects of counterfactual policy interventions. When the effect of interest is conditional on agents' state variables, the estimator is subject to the standard "curse of dimensionality" in nonparametric estimation: that is, its rate of convergence is not root- n , and it declines with the number of continuous explanatory variables. However, the estimator of the unconditional average effect is root- n consistent. Therefore, it is possible to obtain precise estimates of policy effects even when the specification of the structural model contains a relatively large number of state variables. As a third contribution, I apply this method to evaluate hypothetical reforms in the rules of a public pension system using data of male blue-collar workers in Sweden. This application illustrates how the method can be used to obtain meaningful estimates of behavioral and welfare effects that do not rely on parametric assumptions on the primitives of the model.

A recent study by Heckman and Navarro (2007) deals with the identification of dynamic discrete structural models. Heckman and Navarro consider the identification of primitive structural functions; for example, the utility function and the transition probability function of state variables. Their sampling framework is very similar to the one in this article: the researcher observes agents' actions, some state variables, and an outcome variable. As mentioned above, observing an outcome variable does not solve the Rust–Magnac–Thesmar underidentification problem. In order to get around this underidentification, Heckman and Navarro incorporate several restrictions. The most substantial restriction, and the one that clearly distinguishes the Heckman and Navarro model from the approach in this article, is that they assume that the continuation value associated with one of the choice alternatives is known to the researcher; for example, it is "normalized" to zero. This restriction is also considered in Taber (2000). Though this restriction may be needed for the identification of the utility function, I show in this article that it is not needed for the identification

of the effects of a wide and relevant class of counterfactual experiments. Furthermore, if this restriction does not hold in the "true" model, then the predicted effects of some counterfactual policies based on the estimated model can be inconsistent. Bajari and Hong (2006) study the identification of a semiparametric dynamic discrete choice game where the probability distribution of the unobservables is parametric (extreme value type 1). Like Heckman and Navarro, they consider the identification of primitive structural functions. They present exclusion restrictions that can identify players' utility functions.

The rest of the article is organized as follows. In Section 2, I set up the model and the basic assumptions. Section 3 presents the identification results. Section 4 describes the estimation procedure. The empirical application is featured in Section 5. I summarize and conclude in Section 6. Proofs of propositions are in the Appendix.

2. MODEL

2.1 Framework and Basic Assumptions

Time is discrete and indexed by t . Agents have preferences defined over a sequence of states of the world between periods 0 and T , where the time horizon T can be finite or infinite. A state of the world has two components: a vector of state variables \mathbf{s}_t that is predetermined before period t ; and a discrete decision $a_t \in A = \{0, 1, \dots, J\}$ that the agent chooses at period t . The set of feasible choices at time t may depend on the state: $a_t \in A(\mathbf{s}_t)$, with $A(\mathbf{s}_t) \subseteq A$. The decision at period t affects the evolution of future values of the state variables. An agent's preferences over possible sequences of states of the world can be represented by the time-separable utility function $\sum_{j=0}^T \beta^j U_t(a_{t+j}, \mathbf{s}_{t+j})$, where $\beta \in [0, 1)$ is the discount factor and $U_t(a_t, \mathbf{s}_t)$ is the current utility function at period t . Agents have uncertainty about future values of state variables. Their beliefs about future states can be represented by a sequence of Markov transition probability functions $F_t(\mathbf{s}_{t+1} | a_t, \mathbf{s}_t)$. These beliefs are *rational* in the sense that they are the true transition probabilities of the state variables. In every period t , the agent observes the vector of state variables \mathbf{s}_t and chooses an action $a_t \in A(\mathbf{s}_t)$ to maximize the expected utility

$$E \left(\sum_{j=0}^T \beta^j U_t(a_{t+j}, \mathbf{s}_{t+j}) \mid a_t, \mathbf{s}_t \right). \quad (1)$$

Let $\alpha_t(\mathbf{s}_t)$ and $V_t(\mathbf{s}_t)$ be the optimal decision rule and the value function at period t , respectively. By Bellman's principle of optimality, the sequence of value functions can be obtained using the recursive expression:

$$V_t(\mathbf{s}_t) = \max_{a \in A(\mathbf{s}_t)} \left\{ U_t(a, \mathbf{s}_t) + \beta \int V_{t+1}(\mathbf{s}_{t+1}) dF_t(\mathbf{s}_{t+1} | a, \mathbf{s}_t) \right\}. \quad (2)$$

The optimal decision rule $\alpha_t(\mathbf{s}_t)$ is the $\arg \max_{a \in A(\mathbf{s}_t)} \{ U_t(a, \mathbf{s}_t) + \beta \int V_{t+1}(\mathbf{s}_{t+1}) dF_t(\mathbf{s}_{t+1} | a, \mathbf{s}_t) \}$.

The researcher observes a random sample of agents who behave according to this model. Agents in the sample are indexed by $i \in \{1, 2, \dots, n\}$. As is typically the case in micro panels, we observe each individual over a short period of time. I consider that each individual is observed for two periods. It is important

to note that in this article the time index t does not represent calendar time but the agent-specific period in the agent's decision problem. To emphasize this agent-specific time, I refer to t as the *agent's age*. In some parts of the article I will also emphasize this point by using the variable t_i to represent agent i 's age. For each agent in the sample, the econometrician observes his action, a_{it} , and a subvector \mathbf{x}_{it} of the vector of state variables \mathbf{s}_{it} , that is, $\mathbf{x}_{it} \subset \mathbf{s}_{it}$, and an *outcome variable* y_{it} that contains information about utility but is not one of the model's actions or state variables. For instance, in a model of firm behavior the researcher may observe firms' output or revenue, or in a model of individual behavior the econometrician may observe individual earnings. This outcome variable depends on current values of the action and the state variables. I specify this relationship as $y_{it} = Y_t(a_{it}, \mathbf{s}_{it})$, where $Y_t(\cdot)$ is the *outcome function*; for example, a production function. In summary, the researcher's dataset is

$$\text{Data} = \{a_{it}, \mathbf{x}_{it}, y_{it} : i = 1, 2, \dots, N; t = t_i, t_i + 1\}. \quad (3)$$

The vector of state variables \mathbf{s}_{it} can be decomposed into three subvectors: $\mathbf{s}_{it} = (\mathbf{x}_{it}, \omega_{it}, \boldsymbol{\varepsilon}_{it})$ where ω_{it} represents unobservables that enter into the outcome function, and $\boldsymbol{\varepsilon}_{it}$ represents unobservables that enter into the utility function but not into the outcome function. Both ω_{it} and $\boldsymbol{\varepsilon}_{it}$ are *structural state variables* in the sense that they have an economic interpretation within the model; for example, a productivity shock, or a temporary shock in health status.

Assumptions 1–6 establish the key restrictions that will be used to obtain our identification results. Assumptions 1–5 refer to the structural functions of the model: preferences, beliefs, and choice set. Assumption 6 (in Section 2.2) defines the class of policy interventions that I consider in this article.

By the law of conditional probability, the transition probability of the state variables, $F(\mathbf{s}_{i,t+1} | a_{it}, \mathbf{s}_{it})$, can be written as the product of two transitions, one for the unobservables ($\boldsymbol{\varepsilon}, \omega$) and the other for the observables \mathbf{x} : that is, $F(\mathbf{s}_{i,t+1} | a_{it}, \mathbf{s}_{it}) = \Pr(\boldsymbol{\varepsilon}_{i,t+1}, \omega_{i,t+1} | \mathbf{x}_{i,t+1}, a_{it}, \mathbf{s}_{it}) \Pr(\mathbf{x}_{i,t+1} | a_{it}, \mathbf{s}_{it})$. Assumption 1 establishes conditional independence restrictions on these transition probabilities.

Assumption 1 (Conditional independence). (A) The unobserved state variables have exogenous transition probabilities in the sense that they do not depend on the agent's current action: $\Pr(\boldsymbol{\varepsilon}_{i,t+1}, \omega_{i,t+1} | \mathbf{x}_{i,t+1}, a_{it}, \mathbf{s}_{it}) = \Pr(\boldsymbol{\varepsilon}_{i,t+1}, \omega_{i,t+1} | \mathbf{x}_{i,t+1}, \omega_{it}, \boldsymbol{\varepsilon}_{it})$; (B) conditional on the contemporaneous value of \mathbf{x} , the variables $\boldsymbol{\varepsilon}$ and ω have independent transitions, and $\boldsymbol{\varepsilon}$ is not serially correlated: $\Pr(\boldsymbol{\varepsilon}_{i,t+1}, \omega_{i,t+1} | \mathbf{x}_{i,t+1}, \omega_{it}, \boldsymbol{\varepsilon}_{it}) = F_{\boldsymbol{\varepsilon}}(\boldsymbol{\varepsilon}_{i,t+1} | \mathbf{x}_{i,t+1}) F_{\omega}(\omega_{i,t+1} | \mathbf{x}_{i,t+1}, \omega_{it})$; and (C) the evolution of \mathbf{x}_{it} may be endogenous (i.e., dependent on the agent's actions), but conditional on \mathbf{x}_{it} and a_{it} , the vector $\mathbf{x}_{i,t+1}$ does not depend on ω_{it} and $\boldsymbol{\varepsilon}_{it}$, $\Pr(\mathbf{x}_{i,t+1} | a_{it}, \mathbf{x}_{it}, \omega_{it}, \boldsymbol{\varepsilon}_{it}) = F_{\mathbf{x}}(\mathbf{x}_{i,t+1} | a_{it}, \mathbf{x}_{it})$. Under these conditions, the cumulative transition probability of the state variables factors as

$$F(\mathbf{s}_{i,t+1} | a_{it}, \mathbf{s}_{it}) = F_{\boldsymbol{\varepsilon}}(\boldsymbol{\varepsilon}_{i,t+1} | \mathbf{x}_{i,t+1}) F_{\omega}(\omega_{i,t+1} | \mathbf{x}_{i,t+1}, \omega_{it}) \times F_{\mathbf{x}}(\mathbf{x}_{i,t+1} | a_{it}, \mathbf{x}_{it}). \quad (4)$$

This assumption is similar to Rust's conditional independence assumption (see Rust 1994). It is weaker than Rust's because it allows the unobservable ω_{it} to be serially correlated.

Relative to the conditional independence assumptions in Heckman and Navarro (2007), Assumption 1 is neither more general nor more restrictive. Heckman and Navarro also assume that $\boldsymbol{\varepsilon}$ and ω are conditionally independent, but they allow for time-invariant unobserved heterogeneity in these unobservables (see their page 365). However, they assume that all the observed state variables (other than the indicator of the lagged stopping decision) follow strictly exogenous stochastic processes that do not depend on the agent's actions [see their theorem 4, condition (ii)].

Assumption 2 (Additivity of the outcome in the utility function). The utility function U_t is additive in the outcome function. For any possible action $a \in A$,

$$U_t(a, \mathbf{s}_{it}) = Y_t(a, \mathbf{x}_{it}, \omega_{it}) + C_t(a, \mathbf{x}_{it}, \boldsymbol{\varepsilon}_{it}), \quad (5)$$

where $C_t(\cdot)$ is a real-valued function.

Together with the conditional independence in Assumption 1, this assumption implies that the utility of alternative a is the sum of two random variables, $Y_t(a, \mathbf{x}_{it}, \omega_{it})$ and $C_t(a, \mathbf{x}_{it}, \boldsymbol{\varepsilon}_{it})$, which are independent conditional on \mathbf{x}_{it} . This assumption is not innocuous and it restricts the structure of the utility function. For instance, in the retirement model of the application in Section 5, this assumption implies that an individual's utility is the sum of annual earnings and the utility of leisure, and that these two components are independent once we condition on age, marital status, and pension wealth.

For our analysis, it is convenient to decompose $C_t(a, \mathbf{x}_{it}, \boldsymbol{\varepsilon}_{it})$ into two additive components: the median of $C_t(a, \mathbf{x}_{it}, \boldsymbol{\varepsilon}_{it})$ conditional on \mathbf{x}_{it} , and the deviation with respect to this median. For any action $a \in A$ we have that

$$C_t(a, \mathbf{x}_{it}, \boldsymbol{\varepsilon}_{it}) = M_t(a, \mathbf{x}_{it}) + \eta_{it}(a), \quad (6)$$

where $M_t(a, \mathbf{x}_{it}) \equiv \text{Median}(C_t(a, \mathbf{x}_{it}, \boldsymbol{\varepsilon}_{it}) | \mathbf{x}_{it})$ and $\eta_{it}(a) \equiv C_t(a, \mathbf{x}_{it}, \boldsymbol{\varepsilon}_{it}) - M_t(a, \mathbf{x}_{it})$. I use $\boldsymbol{\eta}_{it}$ to represent the vector $\{\eta_{it}(0), \eta_{it}(1), \dots, \eta_{it}(J)\}$. By construction, the random variables in $\boldsymbol{\eta}_{it}$ have median zero and are median independent of \mathbf{x}_{it} . It is possible to show that $\boldsymbol{\eta}_{it}$ and $\boldsymbol{\varepsilon}_{it}$ generate the same sigma-algebra and that they share some properties. For instance, if $\boldsymbol{\varepsilon}_{it}$ satisfies the conditional independence in Assumption 1, then $\boldsymbol{\eta}_{it}$ also satisfies

$$F(\mathbf{s}_{i,t+1} | a_{it}, \mathbf{s}_{it}) = F_{\boldsymbol{\eta}}(\boldsymbol{\eta}_{i,t+1} | \mathbf{x}_{i,t+1}) F_{\omega}(\omega_{i,t+1} | \mathbf{x}_{i,t+1}, \omega_{it}) \times F_{\mathbf{x}}(\mathbf{x}_{i,t+1} | a_{it}, \mathbf{x}_{it}). \quad (7)$$

Nevertheless, it is important to note that $\boldsymbol{\varepsilon}_{it}$ and $\boldsymbol{\eta}_{it}$ are different vectors of random variables. While $\boldsymbol{\varepsilon}_{it}$ contains structural variables with a clear economic interpretation within the model, this is not generally the case for $\boldsymbol{\eta}_{it}$. Furthermore, the identification of the probability distribution of $\boldsymbol{\eta}_{it}$ does not imply the identification of the distribution of $\boldsymbol{\varepsilon}_{it}$ or of the effect of $\boldsymbol{\varepsilon}_{it}$ on the function $C_t(a, \mathbf{x}_{it}, \boldsymbol{\varepsilon}_{it})$. However, for the identification results in this article, we need to know the distribution of $\boldsymbol{\eta}_{it}$ and it is not necessary to know the distribution of $\boldsymbol{\varepsilon}_{it}$ or the function $C_t(a, \mathbf{x}_{it}, \boldsymbol{\varepsilon}_{it})$.

Assumption 3 (Distribution of unobservables). (A) $\boldsymbol{\eta}_{it}$ is a vector of continuous random variables with support the Euclidean space; (B) for any value of \mathbf{x} , $F_{\boldsymbol{\eta}}(\boldsymbol{\eta} | \mathbf{x})$ is continuously differentiable in $\boldsymbol{\eta}$; (C) ω is a continuous random variable

with support \mathbb{R} ; and (D) the transition probability function $F_\omega(\omega_{t+1}|\mathbf{x}_{t+1}, \omega_t)$ is such that for any value of \mathbf{x}_{t+1} and any strictly increasing function $q(\cdot)$, the conditional expectation $E(q(\omega_{t+1})|\mathbf{x}_{t+1}, \omega_t) = \int q(\omega_{t+1})F_\omega(d\omega_{t+1}|\mathbf{x}_{t+1}, \omega_t)$ is an increasing function of ω_t .

Assumptions 3(A), 3(B), and 3(C) are standard. Assumption 3(D) establishes that, given \mathbf{x}_{t+1} , ω_{t+1} depends positively and monotonically on ω_t . Though this is clearly a restriction on the stochastic process of the unobservables ω_t , it is also a common and plausible assumption in many applications.

Assumption 4 (Monotonicity of the outcome function). (A) The outcome function Y_t is strictly monotonic in ω_{it} , such that there is an inverse function Y_t^{-1} and $\omega_{it} = Y_t^{-1}(a_{it}, \mathbf{x}_{it}, y_{it})$; and (B) for any $\mathbf{x} \in X$ and any action $a \neq 0$, the function $\tilde{Y}_t(a, \mathbf{x}, \omega) \equiv Y_t(a, \mathbf{x}, \omega) - Y_t(0, \mathbf{x}, \omega)$ is strictly increasing in ω , and for any value $u \in \mathbb{R}$ there exists a value $\omega \in \mathbb{R}$ such that $\tilde{Y}_t(a, \mathbf{x}, \omega) = u$.

Assumption 4(A) is used to identify point-wise the sample values of the unobserved variable ω_{it} . The monotonicity of $\tilde{Y}_t(a, \mathbf{x}, \omega)$ with respect to ω in Assumption 4(B) is used to identify the distribution function $F_\eta(\eta|\mathbf{x})$ from the discrete choices. The strict monotonicity in Assumption 4 is a strong condition that rules out some interesting cases such as discrete or censored outcome variables. However, it can be relaxed to a certain extent. In particular, Assumption 4 can be replaced by a similar assumption but in terms of an observable variable in the vector \mathbf{x} .

Assumption 4'. There is a continuous observable variable $x_K \subseteq \mathbf{x}$ such that η_{it} and ω_{it} are independent of x_{Kit} and the function $\tilde{Y}_t(a, \mathbf{x}, \omega) \equiv Y_t(a, \mathbf{x}, \omega) - Y_t(0, \mathbf{x}, \omega)$ is strictly increasing in x_K , and for any value $u \in \mathbb{R}$ there exists a value $x_K \in \mathbb{R}$ such that $\tilde{Y}_t(a, \mathbf{x}, \omega) = u$.

Under Assumption 4', we do not need to identify point-wise the sample values of ω_{it} . We can exploit the monotonicity of $\tilde{Y}_t(a, \mathbf{x}, \omega)$ with respect to x_K to identify the distribution function $F_\eta(\eta|\mathbf{x})$, and we can use a similar approach as in Aakvik, Heckman, and Vytlačil (2005) to identify the distribution of ω without having to identify point-wise the sample values of ω_{it} .

Assumption 5. The set of feasible choice alternatives at time t depends only on observed state variables: $A(\mathbf{s}_{it}) = A(\mathbf{x}_{it})$.

Assumptions 1 and 5 imply that the optimal decision rule $\alpha_t(\mathbf{s}_{it})$ can be described as $\alpha_t(\mathbf{s}_{it}) = \arg \max_{a \in A(\mathbf{x}_{it})} \{v_t(a, \mathbf{x}_{it}, \omega_{it}) + \eta_{it}(a)\}$, where the functions $v_t(0, \mathbf{x}, \omega)$, $v_t(1, \mathbf{x}, \omega)$, \dots , $v_t(J, \mathbf{x}, \omega)$ are choice-specific value functions and they are defined as

$$\begin{aligned} v_t(a, \mathbf{x}_t, \omega_t) &\equiv Y_t(a, \mathbf{x}_t, \omega_t) + M_t(a, \mathbf{x}_t) \\ &+ \int \max_{a' \in A(\mathbf{x}_{t+1})} \{v_{t+1}(a', \mathbf{x}_{t+1}, \omega_{t+1}) + \eta_{t+1}(a')\} \\ &\times F(ds_{t+1}|a, \mathbf{x}_t, \omega_t). \end{aligned} \quad (8)$$

The optimal decision rule represents individuals' behavior. Individuals' welfare is given by the value function $V_t(\mathbf{s}_t) = \max_{a \in A(\mathbf{x}_t)} \{v_t(a, \mathbf{x}_t, \omega_t) + \eta_t(a)\}$. For the econometric analysis, it is convenient to define versions of these functions that are

integrated over the unobservables in η_t . The choice probability function is defined as

$$P_t(a|\mathbf{x}_t, \omega_t) \equiv \int I\{\alpha(\mathbf{x}_t, \omega_t, \eta_t) = a\} F_\eta(d\eta_t|\mathbf{x}_t). \quad (9)$$

The integrated-valued function is

$$\begin{aligned} \bar{V}_t(\mathbf{x}_t, \omega_t) &\equiv \int V_t(\mathbf{s}_t) F_\eta(d\eta_t|\mathbf{x}_t) \\ &= \int \max_{a \in A(\mathbf{x}_t)} \{v_t(a, \mathbf{x}_t, \omega_t) + \eta_t(a)\} F_\eta(d\eta_t|\mathbf{x}_t). \end{aligned} \quad (10)$$

2.2 Policy Interventions

Consider a hypothetical policy intervention that modifies the current utility function. Let U_t be the utility function in the *data generating process*, and let U_t^* be the utility function under the counterfactual policy. Assumption 6 describes the class of policy interventions that I consider in this article.

Assumption 6. The counterfactual policy can be represented as a change in the utility function from function U_t to function U_t^* such that: (A) the econometrician knows the difference between the two utility functions [i.e., the econometrician knows the function $\tau_t(a, \mathbf{s}) \equiv U_t^*(a, \mathbf{s}) - U_t(a, \mathbf{s})$]; and (B) the function $\tau_t(a, \mathbf{s})$ does not depend on η , that is, $\tau_t(a, \mathbf{s}) = \tau_t(a, \mathbf{x}, \omega)$.

The function τ_t represents the policy intervention and it is known to the researcher, though the functions U_t and U_t^* are unknown. It may depend on (a, \mathbf{x}, ω) in a completely unrestricted way, but it cannot depend on the unobservable η .

Let $\{P_t, \bar{V}_t\}$ and $\{P_t^*, \bar{V}_t^*\}$ be the choice probability functions and the integrated value functions before and after the policy intervention, respectively. I represent the behavioral effects of the policy by comparing the functions P_t^* and P_t . Similarly, the difference between the functions \bar{V}_t^* and \bar{V}_t represents the welfare effects of the policy.

2.3 An Example: A Model of Capital Replacement

The example is a simplified version of the model in Kasahara (2009), which examines the impact on firms' equipment investment of a temporary increase in import tariffs in Chile. A firm produces a good using capital (a machine) and perfectly flexible inputs. The firm has multiple plants and each plant consists of only one machine. Production at different plants is independent and there are constant returns to scale. Therefore, we can concentrate on the decision problem for an individual plant or machine. The current profit of a plant is equal to variable profits minus maintenance costs and machine replacement costs. Let x_t represent *the time since last machine replacement*, or the age of the existing machine at the beginning of month t . And let $a_t \in \{0, 1\}$ be the indicator of the decision of replacing the old machine by a new machine at the beginning of month t . Therefore, the age of the machine that is used during month t is $(1 - a_t)x_t$. The variable profit during month t depends on the age of the machine and on a productivity shock ω_t , and it is represented by the variable profit function $Y((1 - a_t)x_t, \omega_t)$. The amount of variable inputs is a deterministic function of the age of the machine and the productivity shock and therefore it

is implicit in the function $Y(\cdot)$. The productivity shock ω_t follows a Markov process, and the transition rule of the age of the machine is $x_{t+1} = (1 - a_t)(x_t + 1) + a_t$. The profit function has the following form:

$$U_t = Y((1 - a_t)x_t, \omega_t) - C^M((1 - a_t)x_t, \varepsilon_t^M) - a_t C^R(x_t, \varepsilon_t^R). \quad (11)$$

$C^M(\cdot, \cdot)$ is the maintenance cost, which depends on the age of the machine and on a random shock ε_t^M that is unobservable to the researcher. $C^R(\cdot, \cdot)$ is the replacement cost net of the scrapping value of the retired capital. The replacement cost is paid only if the machine is replaced, and it depends on the age of the old machine and on a random shock ε_t^R that is also unobservable to the researcher. In this model, Assumption 1 implies that the cost shocks ε_t^R and ε_t^M are independently distributed over time (though their distribution may depend on the age of the machine) and distributed independently of the productivity shock ω_t . The specification of the profit function imposes the additive separability of Assumption 2. Using the same decomposition as in expression (6), I can write the profit function as

$$U_t = Y((1 - a_t)x_t, \omega_t) - M_t(a_t, x_t) - \eta_t(a_t), \quad (12)$$

where $M_t(0, x_t) \equiv \text{Median}(C^M(x_t, \varepsilon_t^M)|x_t)$, $M_t(1, x_t) \equiv \text{Median}(C^M(0, \varepsilon_t^M) + C^R(x_t, \varepsilon_t^R)|x_t)$, $\eta_t(0) \equiv C^M(x_t, \varepsilon_t^M) - M_t(0, x_t)$, and $\eta_t(1) \equiv C^M(0, \varepsilon_t^M) + C^R(x_t, \varepsilon_t^R) - M_t(1, x_t)$. Assumption 1(B) implies that $\eta_t(0)$ and $\eta_t(1)$ are independent of ω_t .

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 replacement subsidy that depends on the age of the old machine. The amount of the subsidy is

$$\tau(a_t, x_t) = \begin{cases} 0 & \text{if } x_t < x_{\text{low}}^* \\ a_t(\lambda_0 - \lambda_1[x_t - x_{\text{low}}^*]) & \text{if } x_{\text{high}}^* \leq x_t \leq x_{\text{low}}^* \\ 0 & \text{if } x_t > x_{\text{high}}^* \end{cases} \quad (13)$$

where $\{\lambda_0, \lambda_1, x_{\text{low}}^*, x_{\text{high}}^*\}$ are parameters that characterize the policy, and they have the following properties: $\lambda_0 > 0$, $\lambda_1 > 0$, and $(x_{\text{high}}^* - x_{\text{low}}^*) \leq \lambda_0/\lambda_1$. The subsidy is zero if replacement takes place too early (i.e., before age x_{low}^*) or too late (i.e., after age x_{high}^*). For replacement ages within the range $[x_{\text{low}}^*, x_{\text{high}}^*]$, the subsidy is strictly positive and it decreases linearly with the age of the machine.

3. IDENTIFICATION

Suppose that we have a random sample of individuals with information on the variables $\{a_{it}, a_{i,t+1}, t_i, \mathbf{x}_{it}, \mathbf{x}_{i,t+1}, y_{it}, y_{i,t+1}\}$. As usual, I study identification with a very large (i.e., infinite) sample of individuals. Furthermore, I assume that the sample has variability over the whole support of the observable variables, $A^2 \times [1, T] \times X^2 \times Y^2$. This assumption of full-support variation is needed to identify the reduced form of the model. For the sake of simplicity, I concentrate on binary choice models: $a \in \{0, 1\}$. For notational simplicity I use $P_t(\mathbf{x}, \omega)$ to denote $P_t(1|\mathbf{x}, \omega)$. Also, to simplify notation, I use $F_{\eta_t}(\cdot)$ to represent the distribution of η_t conditional on \mathbf{x}_t at period t (i.e., I omit \mathbf{x}_t as an argument).

3.1 Preliminaries

I assume that the outcome function $Y_t(\cdot)$ is identified without having to estimate the rest of the structural model. There are different conditions under which one can consistently estimate wage equations or production functions using instrumental variables or control function approaches that do not require the estimation of the complete structural model (see Olley and Pakes 1996 and Imbens and Newey 2006). The empirical application of retirement behavior in Section 4 provides an example for the identification of $Y_t(\cdot)$. Following most of the literature on dynamic discrete structural models, I assume that the discount factor β is known to the researcher.

Given $Y_t(\cdot)$, Assumption 4(A) implies that we can recover the sample values of unobservables $\{\omega_{it}\}$ as $\omega_{it} = Y_t^{-1}(a_{it}, \mathbf{x}_{it}, y_{it})$. That is, the sample values of ω_{it} and $\omega_{i,t+1}$ are identified pointwise, and we can treat these variables as observables. Then, Assumption 1 implies that the transition probability functions $F_{\mathbf{x}}$ and F_{ω} are nonparametrically identified. We can identify $F_{\mathbf{x}}$ on $A \times X^2$ from the transition frequencies $\Pr(\mathbf{x}_{i,t+1}|a_{it}, \mathbf{x}_{it})$ in the data. F_{ω} is also identified on $\mathbb{R} \times \mathbb{R} \times X$ from the frequencies $\Pr(\omega_{i,t+1}|\omega_{it}, \mathbf{x}_{it})$ in the data. It is also clear that we can identify the choice probability functions $P_t(\mathbf{x}, \omega)$ on $X \times \mathbb{R}$ from the frequencies $\Pr(a_{it} = 1|\mathbf{x}_{it}, \omega_{it}, t_i = t)$ in the data. However, without further restrictions, we cannot identify the structural functions $\{C_t, F_{\varepsilon t}\}$ or even the *semi-structural* functions $\{M_t, F_{\eta t}\}$. Instead, this article considers the identification of the functions P_t^* and $(\bar{V}_t^* - \bar{V}_t)$ associated with a counterfactual policy intervention.

The following Proposition, based on Matzkin (1994), will be used for some of the identification results.

Proposition 1 (Based on Matzkin 1994). Consider a binary choice model, $a = I\{\eta \leq W + \varphi(X)\}$, where a , W , and X are observable variables to the econometrician, η is unobservable, and $\varphi(\cdot)$ is a real-valued function. Consider the following definitions: $S_X \subseteq \mathbb{R}^K$ is the support of the vector X ; φ^0 is the true function φ in the population; and $F_{\eta|X}^0$ is the true CDF of η conditional on X . Assume that: (i) η and W are independent conditional on X ; (ii) for any $X \in S_X$, the CDF $F_{\eta|X}^0$ is strictly increasing in η , its support is \mathbb{R} , and $F_{\eta|X}^0(0) = 0.5$ (zero conditional median); (iii) the distribution of W conditional on X has a Lebesgue density that is everywhere positive on \mathbb{R} ; and (iv) for any $X \in S_X$, there is a value $W \in \mathbb{R}$ such that $W + \varphi^0(X) = 0$. Under these conditions, the function φ^0 is identified on S_X and the function $F_{\eta|X}^0$ is identified on \mathbb{R} .

An implication, or corollary, of Proposition 1 is the following. Consider the model and the conditions of Proposition 1. Define the counterfactual choice probability function $P^*(X, W) \equiv F_{\eta|X}^0(W + \varphi^0(X) + \pi(X, W))$, where $\pi(X, W)$ is a function from $S_X \times \mathbb{R}$ into \mathbb{R} that is known to the econometrician. Then, P^* is identified on $S_X \times \mathbb{R}$, such that

$$P^*(X, W) = F_{\eta|X}^0(W + \varphi^0(X) + \pi(X, W)) = P^0(X, W + \pi(X, W)). \quad (14)$$

McFadden (1981) defined the *social surplus function* of a random utility model as the difference between the expected utility of behaving optimally minus the expected utility of an arbitrary (suboptimal) behavior. Proposition 2 is a corollary of theorem 5.1 in McFadden (1981), and it establishes that the so-

cial surplus function depends on the optimal choice probability and the CDF of the unobservables only. In the context of a dynamic discrete choice model, Propositions 2 and 3 (below) are key to showing that Matzkin's proposition 1 can be applied recursively to obtain the identification of certain value functions.

Proposition 2. Consider the utility maximization problem $\max\{u(0) + \eta(0), u(1) + \eta(1)\}$, and the following definitions. G is the expected utility of behaving optimally minus the expected utility of choosing alternative zero (McFadden's social surplus), that is,

$$G \equiv \int \max\{u(0) + \eta(0), u(1) + \eta(1)\} dF_{\eta}(\eta) - u(0).$$

P is the probability that alternative 1 is the optimal choice: that is,

$$P \equiv \int I\{u(0) + \eta(0) < u(1) + \eta(1)\} dF_{\eta}(\eta).$$

And $F_{\tilde{\eta}}$ is the CDF of the random variable $\tilde{\eta} \equiv \eta(0) - \eta(1)$. If $F_{\tilde{\eta}}$ is a continuous and strictly increasing function, then the social surplus G depends on P and $F_{\tilde{\eta}}$ only. That is, $G = \int \max\{0, F_{\tilde{\eta}}^{-1}(P) - \tilde{\eta}\} dF_{\tilde{\eta}}(\tilde{\eta})$.

Going back to the dynamic binary choice model, Proposition 3 provides a representation of the choice-specific value functions $v_t(a, \mathbf{x}_t, \omega_t)$ that will be useful to prove the identification results. The proposition establishes that we can decompose additively the differential value functions v_t into three functions. This decomposition is not arbitrary. I show in Section 3.2 that we can identify these three components of the value function, and that these components, together with the CDF of $\tilde{\eta}_t$, can be used to construct the counterfactual choice probability function P_t^* .

Proposition 3. For any age t , the choice-specific value function v_t can be written as

$$v_t(a, \mathbf{x}, \omega) = V_t^Y(a, \mathbf{x}, \omega) + V_t^M(a, \mathbf{x}) + V_t^{OPT}(a, \mathbf{x}, \omega), \quad (15)$$

where: $V_t^Y(a, \mathbf{x}, \omega)$ is the expected, discounted value of the sum of current and future realizations of the outcome variable y if the current choice is a and then alternative 0 is chosen forever in the future; $V_t^M(a, \mathbf{x})$ is the expected, discounted value of the sum of current and future realizations of the component M of the utility function if the current choice is a and then alternative 0 is chosen forever in the future; and $V_t^{OPT}(a, \mathbf{x}, \omega)$ is the value of behaving optimally in the future minus the value of always choosing alternative 0, given that the current choice is a . These functions can be obtained recursively as follows: at last period T , $V_T^Y(a, \mathbf{x}, \omega) = Y_T(a, \mathbf{x}, \omega)$, $V_T^M(a, \mathbf{x}) = M_T(a, \mathbf{x})$, and $V_T^{OPT}(a, \mathbf{x}, \omega) = 0$; and any period $t < T$,

$$\begin{aligned} V_t^Y(a, \mathbf{x}, \omega) &= Y_t(a, \mathbf{x}, \omega) \\ &+ \beta \int V_{t+1}^Y(0, \mathbf{x}_{t+1}, \omega_{t+1}) dF(\mathbf{x}_{t+1}, \omega_{t+1} | a, \mathbf{x}, \omega), \\ V_t^M(a, \mathbf{x}) &= M_t(a, \mathbf{x}) + \beta \int V_{t+1}^M(0, \mathbf{x}_{t+1}) dF_{\mathbf{x}}(\mathbf{x}_{t+1} | a, \mathbf{x}), \end{aligned} \quad (16)$$

$$\begin{aligned} V_t^{OPT}(a, \mathbf{x}, \omega) &= \beta \int [G(P_{t+j}(\mathbf{x}_{t+1}, \omega_{t+1}), F_{\tilde{\eta}_t, t+j}) \\ &+ V_{t+1}^{OPT}(0, \mathbf{x}_{t+1}, \omega_{t+1})] dF(\mathbf{x}_{t+1}, \omega_{t+1} | a, \mathbf{x}, \omega), \end{aligned}$$

where $G(P, F)$ is the social surplus function $\int \max\{0; F_{\tilde{\eta}}^{-1}(P) - \tilde{\eta}\} dF_{\tilde{\eta}}(\tilde{\eta})$, as derived in Proposition 2.

3.2 Identification of Counterfactual Policy Effects

I present here identification results in the context of a finite horizon model, that is, $T < \infty$. However, the proof can be extended to infinite horizon models. The proof of identification in a model with infinite horizon appears in an earlier version of this article, which is available online at ideas.repec.org. The factual choice probability function is $P_t(\mathbf{x}, \omega) = F_{\tilde{\eta}_t}(v_t(1, \mathbf{x}, \omega) - v_t(0, \mathbf{x}, \omega))$, or, given Proposition 3

$$P_t(\mathbf{x}, \omega) = F_{\tilde{\eta}_t}(\tilde{V}_t^Y(\mathbf{x}, \omega) + \tilde{V}_t^M(\mathbf{x}) + \tilde{V}_t^{OPT}(\mathbf{x}, \omega)), \quad (17)$$

where $\tilde{V}_t^Y(\mathbf{x}, \omega) \equiv V_t^Y(1, \mathbf{x}, \omega) - V_t^Y(0, \mathbf{x}, \omega)$, $\tilde{V}_t^M(\mathbf{x}) \equiv V_t^M(1, \mathbf{x}) - V_t^M(0, \mathbf{x})$, and $\tilde{V}_t^{OPT}(\mathbf{x}, \omega) \equiv V_t^{OPT}(1, \mathbf{x}, \omega) - V_t^{OPT}(0, \mathbf{x}, \omega)$. The value function V_t^{OPT} depends on the agent's future optimal behavior and therefore on the optimal choice probability functions from period $t + 1$ until period T . Proposition 4 establishes the identification of the functions $F_{\tilde{\eta}_t}$, \tilde{V}_t^Y , \tilde{V}_t^M , and \tilde{V}_t^{OPT} .

Proposition 4. Suppose that Assumptions 1–5 hold and that the discount factor β is known. Then, for any period $t \leq T$, the functions $F_{\tilde{\eta}_t}$, \tilde{V}_t^Y , \tilde{V}_t^M , and \tilde{V}_t^{OPT} are identified.

The value $\tilde{V}_t^M(\mathbf{x})$ depends on $\tilde{M}_t(\mathbf{x})$, the future utilities $\{M_{t+j}(0, \mathbf{x}_{t+j}) : j > 0\}$, the transition probability function of \mathbf{x} , and the time discount factor β . Without further restrictions, the identification of the functions $\{\tilde{V}_t^M\}$ does not provide identification of the one-period utility function $M_t(a, \mathbf{x})$, or even of the differential utility function $\tilde{M}_t(\mathbf{x}) \equiv M_t(1, \mathbf{x}) - M_t(0, \mathbf{x})$. The one-period utility function is not identified. As in the case of static discrete choice models, there is a "normalization" restriction that implies the identification of utility function $M_t(a, \mathbf{x})$. If we assume that $M_t(0, \mathbf{x}) = 0$ for any period t and any value of \mathbf{x} , then, using the definition of \tilde{V}_t^M in Equation (16), it is simple to verify that $V_t^M(0, \mathbf{x}) = 0$ and $\tilde{V}_t^M(\mathbf{x}) = V_t^M(1, \mathbf{x}) = M_t(1, \mathbf{x})$. However, this "normalization" is not innocuous. If the normalization restriction is not true (and we expect that to be the case in most applications), the estimated model will imply inconsistent predictions of the effects of certain policy interventions, such as those that modify transition probabilities of the state variables or the time discount factor.

Next, we show that, despite the fact that we cannot identify individuals' utility functions, we can identify the behavioral and welfare effects of counterfactual experiments in the class defined in Assumption 6. Define the counterfactual outcome function $Y_t^*(a, \mathbf{x}, \omega) \equiv Y_t(a, \mathbf{x}, \omega) + \tau_t(a, \mathbf{x}, \omega)$, where $\tau_t(\cdot)$ is the function defined in Assumption 6. Therefore, the counterfactual choice probability function is

$$P_t^*(\mathbf{x}, \omega) = F_{\tilde{\eta}_t}(\tilde{V}_t^{Y^*}(\mathbf{x}, \omega) + \tilde{V}_t^{M^*}(\mathbf{x}) + \tilde{V}_t^{OPT^*}(\mathbf{x}, \omega)), \quad (18)$$

where \tilde{V}_t^{Y*} , \tilde{V}_t^{M*} , and \tilde{V}_t^{OPT*} are the counterfactual versions of \tilde{V}_t^Y , \tilde{V}_t^M , and \tilde{V}_t^{OPT} , respectively. It is clear that V_t^{Y*} is known, and that $V_t^{M*} = V_t^M$. Therefore, by Proposition 4, the functions \tilde{V}_t^{Y*} and \tilde{V}_t^{M*} are identified. Also, note that all the primitives that enter into the value function \tilde{V}_t^{OPT*} (i.e., transition probabilities, distribution of $\tilde{\eta}_t$, and discount factor) are policy invariant: they are the same before and after the policy change. Therefore, the only way in which the value function \tilde{V}_t^{OPT*} is affected by the policy change is through the change in the optimal behavior between periods $t + 1$ and T , that is, through the change in the optimal choice probabilities between $t + 1$ and T . To emphasize this aspect, I represent the functions $V_t^{OPT}(a, \mathbf{x}, \omega)$ and $V_t^{OPT*}(a, \mathbf{x}, \omega)$ as $\tilde{V}_t^{OPT}(\mathbf{x}, \omega; \{P_{t+j}: j > 0\})$ and $\tilde{V}_t^{OPT*}(\mathbf{x}, \omega; \{P_{t+j}^*: j > 0\})$, respectively. Taking into account these considerations, we can write the counterfactual choice probability functions as follows:

$$P_t^*(\mathbf{x}, \omega) = F_{\tilde{\eta}_t}(\tilde{V}_t^{Y*}(\mathbf{x}, \omega) + \tilde{V}_t^M(\mathbf{x}) + \tilde{V}_t^{OPT*}(\mathbf{x}, \omega; \{P_{t+j}^*: j > 0\})). \quad (19)$$

Proposition 5 establishes the identification of the counterfactual choice probabilities.

Proposition 5. Suppose that Assumptions 1–6 hold and that the discount factor β is known. Then, the counterfactual choice probability functions $\{P_t^*\}$ are identified over $S_X \times \mathbb{R}$. These probability functions can be obtained using the following method. Starting at the last period, $t = T$, the probability $P_t^*(\mathbf{x}, \omega)$ can be obtained recursively as

$$P_t^*(\mathbf{x}, \omega) = P_t(\mathbf{x}, \omega_t^*(\mathbf{x}, \omega)), \quad (20)$$

where $\omega_t^*(\mathbf{x}, \omega)$ is a function from $S_X \times \mathbb{R}$ into \mathbb{R} that is implicitly defined as the value $\omega_t^* \in \mathbb{R}$ that solves the equation

$$\begin{aligned} & \tilde{V}_t^Y(\mathbf{x}, \omega_t^*) + \tilde{V}_t^{OPT}(\mathbf{x}, \omega_t^*; \{P_{t+j}: j > 0\}) \\ &= \tilde{V}_t^{Y*}(\mathbf{x}, \omega) + \tilde{V}_t^{OPT*}(\mathbf{x}, \omega; \{P_{t+j}^*: j > 0\}). \end{aligned} \quad (21)$$

Proposition 6 establishes the identification of the welfare effect function $\bar{V}_t^* - \bar{V}_t$.

Proposition 6. Under the conditions in Proposition 5, the welfare effect function $\bar{V}_t^* - \bar{V}_t$ is identified. We can obtain this function as:

$$\begin{aligned} & \bar{V}_t^*(\mathbf{x}, \omega) - \bar{V}_t(\mathbf{x}, \omega) \\ &= [V_t^{Y*}(0, \mathbf{x}, \omega) - V_t^Y(0, \mathbf{x}, \omega)] \\ &+ [G(P_t^*(\mathbf{x}, \omega), F_{\tilde{\eta}_t}) - G(P_t(\mathbf{x}, \omega), F_{\tilde{\eta}_t})] \\ &+ [\tilde{V}_t^{OPT}(0, \mathbf{x}, \omega; \{P_{t+j}: j > 0\}) \\ &- \tilde{V}_t^{OPT*}(0, \mathbf{x}, \omega; \{P_{t+j}^*: j > 0\})]. \end{aligned} \quad (22)$$

4. ESTIMATION METHOD

Suppose that we observe $\{a_{it_i}, a_{i,t_i+1}, \mathbf{x}_{it_i}, \mathbf{x}_{i,t_i+1}, y_{it_i}, y_{i,t_i+1}\}$ for a random sample of N individuals. For any $u \in \mathbb{R}$, define $w_t^u(\mathbf{x})$ as the value of ω that solves the equation $\tilde{V}_t^Y(\mathbf{x}, \omega) + \tilde{V}_t^M(\mathbf{x}) + \tilde{V}_t^{OPT}(\mathbf{x}, \omega) = u$. By the zero median of $\tilde{\eta}_t$, $w_t^0(\mathbf{x})$ is

the value of ω that solves the equation $P_t(\mathbf{x}, \omega) = 0.5$. The estimation method proceeds in two steps.

Step 1. Estimation of the outcome function Y , transition density functions f_ω and f_x , factual choice probability function P_t , and threshold function w_t^0 .

First, we estimate the outcome function. Given this estimated function, we can get the residuals $\hat{\omega}_{it_i} = \hat{Y}_{it_i}^{-1}(y_{it_i}, a_{it_i}, \mathbf{x}_{it_i})$ at every observation i . For the derivation of asymptotic properties, I assume that the outcome function is either parametrically or semiparametrically specified such that the residuals $\{\hat{\omega}_{it_i}\}$ are root- n consistent estimates of the true ω 's. This implies that the rate of convergence of $\{\hat{\omega}_{it_i}\}$ to $\{\omega_{it_i}\}$ is faster than the convergence of the kernel estimators that I define below. Therefore, we can ignore the estimation error in $\{\hat{\omega}_{it_i}\}$ in the derivation of the asymptotic properties of the estimator of policy effects.

Given $\{\hat{\omega}_{it_i}\}$ we can construct procedures to estimate non-parametrically transition densities $f_\omega(\omega'|\omega)$ and $f_x(\mathbf{x}'|a, \mathbf{x})$, choice probability $P_t(\mathbf{x}, \omega)$, and threshold values $w_t^0(\mathbf{x})$ for arbitrary values of $(t, \mathbf{x}, \mathbf{x}', \omega, \omega', a)$. These procedures are called by the estimators in Step 2. For instance, we can use a kernel method to estimate the transition densities $f_\omega(\omega'|\omega)$ and $f_x(\mathbf{x}'|a, \mathbf{x})$. Similarly, a kernel estimator of the choice probability $P_t(\mathbf{x}, \omega)$ is

$$\begin{aligned} \hat{P}_t(\mathbf{x}, \omega) &= \sum_{i=1}^N a_{it_i} I\{t_i = t\} K\left(\frac{\mathbf{x} - \mathbf{x}_{it_i}}{b_n^x}\right) K\left(\frac{\omega - \hat{\omega}_{it_i}}{b_n^\omega}\right) \\ &/ \left[\sum_{i=1}^N I\{t_i = t\} K\left(\frac{\mathbf{x} - \mathbf{x}_{it_i}}{b_n^x}\right) K\left(\frac{\omega - \hat{\omega}_{it_i}}{b_n^\omega}\right) \right]. \end{aligned} \quad (23)$$

An estimator of $w_t^0(\mathbf{x})$ can be obtained by solving in ω the equation $\hat{P}_t(\mathbf{x}, \omega) - 0.5 = 0$. The solution to this equation can be easily obtained using Newton's method. That is, starting at some initial ω_0 , we obtain a sequence of values $\{\omega_k: k \geq 0\}$ using the recursive formula $\omega_{k+1} = \omega_k - \hat{P}_t(\mathbf{x}, \omega_k) / (\partial \hat{P}_t(\mathbf{x}, \omega_k) / \partial \omega)$. We iterate in this formula until $|\omega_{k+1} - \omega_k|$ is smaller than a pre-specified small constant. Since $\hat{P}_t(\mathbf{x}, \omega)$ is strictly increasing in ω , Newton's method always converges to the unique solution of the equation. In some applications, the kernel estimate of the function P may not be strictly monotonic in ω . Monotonicity of \hat{P} is necessary for the subsequent estimation of the effects of the counterfactual policy. We can impose monotonicity using the isotonic smooth (IS) kernel estimator proposed by Mukerjee (1988) and Mammen (1991).

Step 2. Recursive (backward induction) estimation of the functions $\{P_t^*\}$, $\{\tilde{V}_t^M\}$, and $\{F_{\tilde{\eta}_t}\}$.

(At period T) We estimate, in this order: (a) $\omega_T^*(\mathbf{x}, \omega)$; (b) $P_T^*(\mathbf{x}, \omega)$; (c) $\tilde{V}_T^M(\mathbf{x})$; (d) $w_T^u(\mathbf{x})$; and (e) $\hat{F}_{\tilde{\eta}_T}(u)$.

(a) A consistent estimator of $\omega_T^*(\mathbf{x}, \omega)$ is the value ω_T^* that solves the equation $\hat{Y}_T(\mathbf{x}, \omega_T^*) = \hat{Y}_T^*(\mathbf{x}, \omega)$. For instance, suppose that $\hat{Y}_T(\mathbf{x}, \omega) = g_T(\mathbf{x}) \exp(\omega)$, where $g_T(\cdot)$ is a real-valued function of \mathbf{x} . Then there is a closed form expression for the estimator of $\omega_T^*(\mathbf{x}, \omega)$:

$$\hat{\omega}_T^*(\mathbf{x}, \omega) = \omega + \log\left(1 + \frac{\tilde{\tau}_T(\mathbf{x}, \omega)}{g_T(\mathbf{x}) \exp\{\omega\}}\right). \quad (24)$$

(b) Based on Proposition 5, a consistent estimator of $P_T^*(\mathbf{x}, \omega)$ is $\hat{P}_T(\mathbf{x}, \hat{\omega}_T^*(\mathbf{x}, \omega))$, that is:

$$\hat{P}_T^*(\mathbf{x}, \omega) = \sum_{i=1}^N a_{i_i} I\{t_i = T\} K\left(\frac{\mathbf{x} - \mathbf{x}_{i_i}}{b_n^{\mathbf{x}}}\right) K\left(\frac{\hat{\omega}_T^*(\mathbf{x}, \omega) - \hat{\omega}_{i_i}}{b_n^{\omega}}\right) / \left[\sum_{i=1}^N I\{t_i = T\} K\left(\frac{\mathbf{x} - \mathbf{x}_{i_i}}{b_n^{\mathbf{x}}}\right) K\left(\frac{\hat{\omega}_T^*(\mathbf{x}, \omega) - \hat{\omega}_{i_i}}{b_n^{\omega}}\right) \right]. \quad (25)$$

(c) By Proposition 4, a consistent estimator of $\tilde{V}_T^M(\mathbf{x})$ is $\hat{V}_T^M(\mathbf{x}) = -\hat{Y}_T(\mathbf{x}, \hat{w}_T^0(\mathbf{x})) = -g_T(\mathbf{x}) \exp(\hat{w}_T^0(\mathbf{x}))$.

(d) By definition, $w_T^u(\mathbf{x})$ is the value of ω that solves the equation $\tilde{Y}_T(\mathbf{x}, \omega) + \tilde{V}_T^M(\mathbf{x}) = u$. Therefore, a consistent estimator of $w_T^u(\mathbf{x})$ is $\hat{w}_T^u(\mathbf{x}) = \hat{w}_T^0(\mathbf{x}) + \log(1 + \frac{u}{g_T(\mathbf{x}) \exp(\hat{w}_T^0(\mathbf{x}))})$.

(e) By Proposition 4, for $u \in \mathbb{R}$ a consistent estimator of $F_{\tilde{\eta}T}(u)$ is $\hat{P}_T(\mathbf{x}, \hat{w}_T^u(\mathbf{x}))$, that is:

$$\hat{F}_{\tilde{\eta}T}(u) = \sum_{i=1}^N a_{i_i} I\{t_i = T\} K\left(\frac{\mathbf{x} - \mathbf{x}_{i_i}}{b_n^{\mathbf{x}}}\right) K\left(\frac{\hat{w}_T^u(\mathbf{x}) - \hat{\omega}_{i_i}}{b_n^{\omega}}\right) / \left[\sum_{i=1}^N I\{t_i = T\} K\left(\frac{\mathbf{x} - \mathbf{x}_{i_i}}{b_n^{\mathbf{x}}}\right) K\left(\frac{\hat{w}_T^u(\mathbf{x}) - \hat{\omega}_{i_i}}{b_n^{\omega}}\right) \right]. \quad (26)$$

(At period $T - 1$) We estimate, in this order: (a) $\tilde{V}_{T-1}^{OPT}(\mathbf{x}, \omega)$ and $\tilde{V}_{T-1}^{OPT*}(\mathbf{x}, \omega)$; (b) $\omega_{T-1}^*(\mathbf{x}, \omega)$; (c) $P_{T-1}^*(\mathbf{x}, \omega)$; (d) $\tilde{V}_{T-1}^M(\mathbf{x})$; (e) $w_{T-1}^u(\mathbf{x})$; and (f) $\hat{F}_{\tilde{\eta}T-1}(u)$.

(a) First, we obtain an estimator of $\tilde{V}_{T-1}^{OPT}(\mathbf{x}, \omega)$:

$$\hat{V}_{T-1}^{OPT}(\mathbf{x}, \omega) = \beta \int \max\{\hat{F}_{\tilde{\eta}T}^{-1}(\hat{P}_T(\mathbf{x}', \omega')) - u; 0\} \hat{f}_{\tilde{\eta}T}(u) \hat{f}_{\omega}(\omega' | \omega) \times [\hat{f}_{\mathbf{x}}(\mathbf{x}' | 1, \mathbf{x}) - \hat{f}_{\mathbf{x}}(\mathbf{x}' | 0, \mathbf{x})] du d\mathbf{x}' d\omega'. \quad (27)$$

If (\mathbf{x}, ω) includes several continuous state variables, the multiple integral may be approximated using Monte Carlo integration in order to reduce computation time. Similarly, we can obtain a consistent estimator of $\tilde{V}_{T-1}^{OPT*}(\mathbf{x}, \omega)$.

(b) A consistent estimator of $\omega_{T-1}^*(\mathbf{x}, \omega)$ is the value ω_{T-1}^* that solves the equation $\hat{V}_{T-1}^Y(\mathbf{x}, \omega_{T-1}^*) + \hat{V}_{T-1}^{OPT}(\mathbf{x}, \omega_{T-1}^*) = \hat{V}_{T-1}^{Y*}(\mathbf{x}, \omega) + \hat{V}_{T-1}^{OPT*}(\mathbf{x}, \omega)$. In contrast to the case at period T , there is not a closed-form analytical expression for the solution to this equation. However, given that $\hat{V}_{T-1}^Y(\mathbf{x}, \omega) + \hat{V}_{T-1}^{OPT}(\mathbf{x}, \omega)$ is strictly increasing in ω , the solution to this equation can be easily obtained using Newton's method.

(c) A consistent estimator of $P_{T-1}^*(\mathbf{x}, \omega)$ is $\hat{P}_{T-1}^*(\mathbf{x}, \omega) = \hat{P}_{T-1}(\mathbf{x}, \hat{\omega}_{T-1}^*(\mathbf{x}, \omega))$.

(d) By Proposition 4, a consistent estimator of $\tilde{V}_{T-1}^M(\mathbf{x})$ is $\hat{V}_{T-1}^M(\mathbf{x}) = -\hat{V}_{T-1}^Y(\mathbf{x}, \hat{w}_{T-1}^0(\mathbf{x})) - \hat{V}_{T-1}^{OPT}(\mathbf{x}, \hat{w}_{T-1}^0(\mathbf{x}))$.

(e) A consistent estimator of $w_{T-1}^u(\mathbf{x})$ is the value of ω that solves the equation $\hat{V}_{T-1}^Y(\mathbf{x}, \omega) + \hat{V}_{T-1}^M(\mathbf{x}) + \hat{V}_{T-1}^{OPT}(\mathbf{x}, \omega) = u$. Again, the unique solution to this equation can be obtained using Newton's method.

(f) A consistent estimator of $F_{\tilde{\eta}, T-1}(u)$ is $\hat{F}_{\tilde{\eta}, T-1}(u) = \hat{P}_{T-1}(\mathbf{x}, \hat{w}_{T-1}^u(\mathbf{x}))$.

(At period $t < T$) We proceed in the same way as for $t = T - 1$.

Remark 1. The computational cost of estimating the sequence of choice probability functions $\{P_t^*\}$ is equivalent to the cost of solving the dynamic programming problem once.

Remark 2. Under standard regularity conditions, our kernel estimator of $\{P_t^*\}$ is consistent and asymptotically normal, and its rate of convergence is $\sqrt{Nb_N^k}$, where b_N is the bandwidth and k is the number of continuous variables in the vector of state variables (\mathbf{x}, ω) . The speed of convergence declines with the number of continuous regressors. In applications with more than two or three continuous state variables, the estimator of the function P_t^* can be quite imprecise unless the number of observations in our dataset is very large. However, there are situations in which we may be interested in *partial means* of the policy effects. A partial mean of the policy effect is an average of $P_t^*(\mathbf{x}_{i_i}, \omega_{i_i}) - P_t(\mathbf{x}_{i_i}, \omega_{i_i})$ over some state variables holding other state variables constant. In the literature on evaluation of treatment effects, these partial means are called (conditional) average treatment effects (ATE). Let z_i be a vector included in $(\mathbf{x}_{i_i}, \omega_{i_i})$, and let $\mathbf{x}_{(-z)_i}$ represent the vector $(\mathbf{x}_{i_i}, \omega_{i_i})$ excluding z_i , such that $(\mathbf{x}_{i_i}, \omega_{i_i}) = (z_i, \mathbf{x}_{(-z)_i})$. Then, the ATE conditional on z_i taking a constant value z is:

$$ATE(z) = E[P_{t_i}^*(z, \mathbf{x}_{(-z)_i}) - P_{t_i}(z, \mathbf{x}_{(-z)_i})]. \quad (28)$$

When z_i does not include any variable, we have the unconditional ATE, that is, $ATE = E(P_{t_i}^*(\mathbf{x}_{i_i}, \omega_{i_i}) - P_{t_i}(\mathbf{x}_{i_i}, \omega_{i_i}))$. A consistent estimator of the (conditional or unconditional) ATE is simply the sample mean of the kernel estimator $\hat{P}_{t_i}^*(z, \mathbf{x}_{(-z)_i}) - \hat{P}_{t_i}(z, \mathbf{x}_{(-z)_i})$, that is,

$$\widehat{ATE}(z) = \frac{1}{N} \sum_{i=1}^N \hat{P}_{t_i}^*(z, \mathbf{x}_{(-z)_i}) - \hat{P}_{t_i}(z, \mathbf{x}_{(-z)_i}). \quad (29)$$

Newey (1994) studies the asymptotic properties of a partial mean of kernel estimates. He shows that averaging out some continuous conditional variables implies an improvement in the rate of convergence of the estimator. Furthermore, averaging over all the (continuous) state variables to obtain the unconditional ATE provides an estimator that is root- n consistent. Therefore, the nonparametric approach in this article may provide precise and meaningful estimates of average policy effects even when the model has a relatively large number of continuous state variables and the dataset does not contain many observations. The empirical application in Section 5 is an example of this.

Remark 3. Based on the results presented in Remark 2, it is possible to use our estimator of ATE, based on our nonparametric model, to define the following test of a parametric specification. Let \widehat{ATE}_{np} and \widehat{ATE}_p be the estimators of ATE using our nonparametric model and using a parametric model, respectively. Under the null hypothesis that the parametric model is *correctly specified*, we have that

$$\frac{\widehat{ATE}_{np} - \widehat{ATE}_p}{\sqrt{\text{Var}(\widehat{ATE}_{np} - \widehat{ATE}_p)}} \sim_a N(0, 1), \quad (30)$$

where \sim_a means “is asymptotically distributed as.” The asymptotic variance $\text{Var}(\widehat{ATE}_{np} - \widehat{ATE}_p)$ can be approximated using a bootstrap method (see Remark 4 below). An interesting feature of the test of specification is that it tests the validity of a parametric model only in the context of estimating a specific policy effect. That is, some specifications may be invalid for some policy experiments but perfectly valid for other policy experiments.

Remark 4. In the empirical application in Section 5, I use a bootstrap method to approximate the standard errors of estimated average policy effects (unconditional ATEs). The method is a nonparametric bootstrap that resamples the data and replicates the whole estimation procedure for each bootstrap sample. Define a *block* as an individual’s history of observable choices and state variables in the sample. The sample can be described as N blocks, which are independent random draws from a certain population. The bootstrap sampling method that I use takes random draws (independent and with replacement) of blocks from the sample. I take $B = 500$ bootstrap samples of size N ; that is, each bootstrap sample has the same number of blocks/individuals as the original sample. The bootstrap standard error of the estimate \widehat{ATE} is just the square root of $(B - 1)^{-1} \sum_{b=1}^B (\widehat{ATE}_b - \overline{\widehat{ATE}})^2$, where \widehat{ATE}_b is the estimate of the ATE using the b th bootstrap sample and $\overline{\widehat{ATE}}$ is $B^{-1} \sum_{b=1}^B \widehat{ATE}_b$.

5. AN APPLICATION

This section presents an application of this methodology to evaluate the effects of a hypothetical reform in the social security pension system in Sweden. The main purpose of this application is to illustrate the implementation of the method and to show that it can provide meaningful results. Our model of retirement behavior follows Rust and Phelan (1997) and Karlstrom, Palme, and Svensson (2004). The hypothetical reform consists of a delay of three years in the eligibility ages of the public pension system. The minimum age to claim a public pension increases from 60 to 63 years and the *normal* retirement age increases from 65 to 68. This type of reform has been and is still being considered in many Organization for Economic Cooperation and Development (OECD) countries.

Additivity of the outcome variable in the utility function (Assumption 2) and conditional independence between the unobservables in the outcome function and in the rest of the utility function (Assumption 1) are key conditions to implement our method. In order to have a model of retirement that satisfies these assumptions, I impose the following restrictions in the model below: (a) the utility function is additively separable in consumption and leisure; (b) the utility of consumption belongs to the CRRA family and the relative risk aversion parameter is known to the researcher; (c) in every period, consumption is equal to disposable income; and (d) conditional on age, marital status, and public pension wealth, there is independence between the unobservable in labor earnings and the unobservables in the utility of leisure. These are strong assumptions. In particular, assumption (c) rules out consumption smoothing

and accumulation of private wealth, which is an obvious substitute for public pension wealth. Therefore, in our policy analysis we assume that when the public pension system becomes less generous (through a delay in eligibility ages), individuals do not decide to reduce their consumption before retirement in order to accumulate wealth and to retire not much later than they planned before the policy change. This type of response is ignored by our model. Interestingly, this assumption is also present in most of the literature on structural models of retirement behavior, for example, Berkovec and Stern (1991), Rust and Phelan (1997), Karlstrom, Palme, and Svensson (2004), or, more recently, Bound, Stinebrickner, and Waidmann (2009). Some important exceptions are the recent papers by French (2005), Blau (2008), and van der Klaauw and Wolpin (2008).

5.1 A Model of Retirement Behavior

Every year, individuals decide whether to continue working ($a_t = 1$) or to retire and claim social security pension benefits ($a_t = 0$). This decision is irreversible. Individuals have a utility function that is additively separable in consumption (C_t) and leisure (L_t). More specifically,

$$\tilde{U}_t = U_C(C_t) + U_L(L_t, t, m_t, \boldsymbol{\varepsilon}_t), \quad (31)$$

where t represents age, m_t is marital status, and $\boldsymbol{\varepsilon}_t$ is an individual idiosyncratic shock in the utility of leisure that is unobservable to the econometrician (e.g., unobserved health status). These variables capture individual heterogeneity in the utility of leisure. If the individual works, hours of leisure are equal to \bar{L}_1 and annual earnings are equal to labor earnings W_t . If the individual decides to retire, then hours of leisure are \bar{L}_0 and earnings are equal to retirement benefits B_t . Thus, we can write leisure as $L_t = a_t \bar{L}_1 + (1 - a_t) \bar{L}_0$ and annual earnings as $Y_t = a_t W_t + (1 - a_t) B_t$. Labor earnings depend on age, marital status, and a wage shock:

$$W_t = \exp\{h_W(m_t, t) + \omega_{t+1}\}. \quad (32)$$

$h_W(\cdot)$ is a function and ω_{t+1} is a wage shock. This shock follows a Markov process with transition rule $\omega_{t+1} = \rho(\omega_t) + \xi_{t+1}$, where $\rho(\cdot)$ is a function and ξ_{t+1} is the innovation of the process. The individual knows ω_t when he decides whether to retire at age t , but he does not know the innovation ξ_{t+1} . This assumption is similar to the one in Rust and Phelan (1997), and it is important for the identification of the wage function (see below). Retirement benefits depend on current age (t), retirement age (ra_t), and *pension points* (pp_t): $B_t = B(t, m_t, ra_t, pp_t)$. I describe this function in Section 5.2 below.

The state variables of the model are $\boldsymbol{\varepsilon}_t$, ω_t , and $\mathbf{x}_t = (t, m_t, ra_t, pp_t)$. Note that retirement status (i.e., a_{t-1}) is implicitly given by retirement age. Since the individual has uncertainty about current labor earnings, the relevant current utility is the expected utility $U_t \equiv E_t(\tilde{U}_t)$, where the information set at period t is $(a_t, \mathbf{x}_t, \omega_t, \boldsymbol{\varepsilon}_t)$:

$$U_t \equiv E_t(\tilde{U}_t) = E_t(U_C(C_t)) + U_L(a_t \bar{L}_1 + (1 - a_t) \bar{L}_0, t, m_t, \boldsymbol{\varepsilon}_t). \quad (33)$$

In order to have a utility function U_t with the form postulated in Assumptions 1 and 2, we need the value of $E_t(U_C(C_t))$ to be an observable outcome variable y_t for the researcher. If we

could observe individuals' consumption C_t , we would need relatively weak assumptions to have that structure. However, we do not observe individual consumption. Instead, we observe labor earnings W_t for those individuals who are working and retirement benefits B_t (either potential or actual benefits) for every individual, working or not. I assume that consumption is equal to earnings, $C_t = Y_t$, and that the utility of consumption is a CRRA function; that is, $U_C(C_t) = C_t^\alpha$, where α is the parameter of relative risk aversion. Furthermore, I also assume that the risk aversion parameter is known to the researcher. Under these conditions, we have that

$$\begin{aligned} E_t(U_C(C_t)) & \\ & \equiv Y(a_t, \mathbf{x}_t, \omega_t) \\ & = \begin{cases} \exp\{h_W(m_t, t) + \rho(\omega_t) + \mu\}^\alpha & \text{if } a_t = 1 \\ B(\mathbf{x}_t)^\alpha & \text{if } a_t = 0, \end{cases} \quad (34) \end{aligned}$$

where μ is a parameter that represents $\ln(E(\exp\{\xi_{t+1}\}))$. Define the functions $M(0, \mathbf{x}_t)$ and $M(1, \mathbf{x}_t)$ and the variables $\eta_t(0)$ and $\eta_t(1)$ such that, for $a = 0, 1$: $M(a, \mathbf{x}_t) \equiv \text{Median}\{U_L(a\bar{L}_1 + (1-a)\bar{L}_0, t, m_t, \boldsymbol{\varepsilon}_t) | \mathbf{x}_t\}$ and $\eta_t(a) \equiv U_L(a\bar{L}_1 + (1-a)\bar{L}_0, t, m_t, \boldsymbol{\varepsilon}_t) - M(a, \mathbf{x}_t)$. Using these definitions, we can rewrite the utility function as $U_t = Y(a_t, \mathbf{x}_t, \omega_t) + M(a_t, \mathbf{x}_t) + \eta_t(a_t)$.

I conclude this subsection by presenting sufficient conditions for the identification of the outcome function $Y(\cdot)$. Since the pension benefits function $B(\mathbf{x}_t)$ is known and the risk aversion parameter α is assumed to be known to the econometrician, we have to identify only the labor earnings function. A least-squares estimation of the log-wage equation $\log(W_t) = h_W(m_t, t) + \omega_{t+1}$ (using the sample of observations with $a_t = 1$) suffers from selection bias because ω_{t+1} is not mean independent of the individual's choice at period t . However, the Markov structure of the wage shock and our assumption on the arrival of information provide moment conditions that can be used to nonparametrically estimate the functions $h_W(\cdot)$ and $\rho(\cdot)$. Consider the subsample of individuals working at periods t and $t-1$, that is, $a_t = a_{t-1} = 1$. For these individuals, we have that: (a) $\ln W_t = h_W(\mathbf{x}_t) + \rho(\omega_t) + \xi_{t+1}$; and (b) $\omega_t = \ln W_{t-1} - h_W(\mathbf{x}_{t-1})$. Solving Equation (b) into Equation (a), we have that

$$\begin{aligned} \ln W_t = h_W(\mathbf{x}_t) + \rho(\ln W_{t-1} - h_W(\mathbf{x}_{t-1})) + \xi_{t+1} \\ \text{if } a_t = a_{t-1} = 1. \quad (35) \end{aligned}$$

The innovation ξ_{t+1} is iid over time and unknown to the individual when he makes his working decision at period t , or at period $t-1$. This implies that ξ_{t+1} is independent of (a_t, a_{t-1}) , and it is also independent of \mathbf{x}_t and W_{t-1} . The orthogonality conditions $E(\xi_{t+1} | a_t = 1, a_{t-1} = 1, \mathbf{x}_t, W_{t-1}) = 0$ provide moment conditions that can be used to estimate the functions $h_W(\cdot)$ and $\rho(\cdot)$.

5.2 Social Security Pensions and the Counterfactual Reform

I present here a brief description of Sweden's public pension system during the sample period 1983–1997. For a more detailed explanation, see section 2 in Karlstrom, Palme, and

Svensson (2004, hereafter KPS). I discuss first the pension benefits function, $B(\mathbf{x})$, and then the rules for accumulation of pension points.

In Sweden, the measurement unit for the amount of pension benefits is the so-called *basic amount* (*BA*). This measure is indexed by the consumer price index (CPI), and it was equal to \$4,765 U.S. dollars in 2003. Notice that to translate Swedish kronas into U.S. dollars (of the year 2003), I have used an exchange rate of 8.1 Swedish kronas per U.S. dollar. An individual younger than 60 years old is not entitled to receive public pension benefits. Therefore, $B(\mathbf{x}) = 0$ for $t < 60$. Of course, individuals can retire before age 60, but they cannot claim and receive benefits until they become 60 years old. Let $\bar{B}_{65}(m, pp)$ represent the amount of annual pension benefits that the individual is entitled to if he/she retires at age 65 (see below). Then, for ages $t \geq 60$, we have that

$$B(\mathbf{x}) = \begin{cases} \bar{B}_{65}(m, pp)(1 - \kappa_1(65 - 60)) & \text{if } ra < 60 \\ \bar{B}_{65}(m, pp)(1 - \kappa_1(65 - ra)) & \text{if } 60 \leq ra < 65 \\ \bar{B}_{65}(m, pp)(1 + \kappa_2(ra - 65)) & \text{if } 65 \leq ra < 70 \\ \bar{B}_{65}(m, pp)(1 + \kappa_2(70 - 65)) & \text{if } ra \geq 70. \end{cases} \quad (36)$$

κ_1 is a permanent actuarial reduction in benefits per year of early retirement. κ_2 is a permanent actuarial increase in benefits per year of delayed retirement. For our sample period, 1983–1997, the values of these parameters were $\kappa_1 = 6.0\%$ per year (0.5% per month) and $\kappa_2 = 8.4\%$ per year (0.7% per month). The amount $\bar{B}_{65}(m, pp)$ is the combination of three pension programs: basic pension (*BP*), special supplementary pension (*SSP*), and normal supplementary pension (*NSP*). More specifically,

$$\bar{B}_{65}(m, pp) = BP(m) + \min\{SSP, NSP(pp)\}. \quad (37)$$

The basic pension $BP(m)$ depends on marital status but not on the individual's pension points. It is equal to 96.0% of the *BA* (i.e., \$4,574 per year) if the individual is single and 78.5% of the *BA* (i.e., \$3,740 per year) if married. The *SSP* applies only if the individual's normal supplementary pension is not large enough, that is, if $NSP(pp) < SSP$. It is the same for every individual and is equal to 55.5% of the *BA*. The normal supplementary pension $NSP(pp)$ is the most important part of pension benefits for most individuals. It is proportional to the individual's pension points: $NSP(pp) = 0.6 * pp * BA$. As I described below, pension points can take values between 0 and 6.5. Therefore, the range of variation of the annual pension benefits upon retirement at age 65 (i.e., the range of variation of $\bar{B}_{65}(\mathbf{x})$) is $[1.515 * BA, 7.46 * BA]$ (i.e., between \$7,219 and \$35,547) for a single person, and $[1.34 * BA, 7.285 * BA]$ (i.e., between \$6,385 and \$34,713) for a married person. The range of variation of actual pension benefits is larger because individuals retire at different ages, before and after age 65.

Pension points are a deterministic function of an individual's whole history of earnings. Every year, an individual with annual labor earnings W greater than *BA* obtains an amount of social security points between 0 and 6.5, according to the formula $\max\{6.5; \frac{W}{BA} - 1\}$. The pension points are equal to the average social security points over an individual's best 15 years of earnings, adjusted by a factor that depends on the total number of years with earnings greater than the *BA*. In order to provide an explicit formula, let $\{W^{(1)}, W^{(2)}, \dots, W^{(15)}\}$ be the annual labor

earnings over an individual’s best 15 years. Let n_W be the number of years in which the individual obtained annual earnings greater than the BA . Then,

$$pp = \min \left\{ 1, \frac{n_W}{30} \right\} * \left(\frac{1}{15} \sum_{j=1}^{15} \max \left\{ 6.5; \frac{W^{(j)}}{BA} - 1 \right\} \right). \quad (38)$$

KPS show that the evolution of an individual’s pension points over time, though not Markovian, can be very closely approximated by a first-order Markov process. This is also the case for public pension wealth in the U.S. Social Security System, as shown by Rust and Phelan (1997). Following these previous studies, I consider a deterministic Markov rule for pension points; that is, $\ln(pp_{t+1}) = h_{pp}(t, pp_t)$, where $h_{pp}(\cdot, \cdot)$ is a function that I estimate from the data. I consider a specification with pension point in logs because this variable always takes positive values. As shown below, the R -squared coefficient from this regression is very close to one (i.e., 0.998), which justifies treating the transition of pension points as deterministic.

The counterfactual policy that I evaluate is a three-year delay in the eligibility ages. The new eligibility ages are 63 (the “minimum”), 68 (the “normal”), and 73 (the “maximum”). Given that the discount for early retirement is $\kappa_1 = 6.0\%$ per year, an individual who would have retired at age 65 before the policy change will have an 18% reduction in annual pension if he/she still retires at the same age. The function $\tau(a, \mathbf{x})$ represents the change in the one-period utility induced by the policy. The new policy does not affect the utility function when working, and therefore $\tau(1, \mathbf{x}) = 0$. The effect of the utility when retired is $\tau(0, \mathbf{x}) = B^*(\mathbf{x}) - B(\mathbf{x})$, where B^* and B are the benefits functions after and before the reform, respectively.

5.3 Data

The data come from the Swedish Longitudinal Individual Panel (LINDA). This dataset has been used before by KPS to estimate a (parametric) dynamic structural model of retirement. Following KPS, I apply the following filters for the selection of the working sample: (a) sample period 1983–1997; (b) men; (c) blue-collar workers; (d) born between 1927 and

1940; (e) exclude individuals who exit from the labor force through the disability insurance program; and (f) exclude individual year observations in which the individual is younger than 50 years old. The working sample contains 3,129 individuals and 34,593 observations during the sample period 1983–1997.

The dataset also reports annual labor earnings for the year in which an individual retires. However, the average value of annual labor earnings for the year of retirement is clearly lower than before retirement. This is because the individual works only part of the retirement year. Therefore, I exclude observations on labor earnings for the year of retirement, both for the summary statistics in Table 1 and for the estimation of the earnings equation. That is, I consider that the individual is retired during the whole year in which he decides to retire. For the same reason, for the estimation of the transition rule of pension points, I exclude observations of the year of retirement.

Table 1 presents summary statistics for retirement age, labor earnings, and pension wealth as measured by entitled benefits upon retirement at age 65. The dispersion of labor earnings in this sample is quite small: at age 59, more than 80% of individuals have annual earnings between \$18,600 and \$31,000. As expected, the distribution of public pension wealth is even more concentrated: at age 59, more than 80% of individuals have entitled benefits (upon retirement at age 65) between \$14,000 and \$20,400. Interestingly, from age 59 to age 65, median labor earnings and median pension wealth decline with age. Given that it is very unlikely that, for a given individual, pension wealth declines with age, this effect is most likely due to selection; that is, individuals with higher earnings and pension wealth are more likely to retire between ages 60 and 65. This trend reverses after age 65, and median earnings and median pension wealth increase between ages 65 and 68. Therefore, it seems that the selection effect is weaker after age 65. The quantiles of the distribution of retirement age are based on a Kaplan–Meier estimator that takes into account the existence of right-censored spells. Figure 1 presents the Kaplan–Meier estimator of the CDF of retirement age. The median retirement age is between 64 and 65.

The dataset includes information on aggregate mortality rates for men in these cohorts. I have used this information to construct an age variant discount factor that accounts for mortality

Table 1. Summary statistics: 3,129 male blue-collar workers; cohorts 1927–1940; years 1983–1997

Variable	Median	Quantile 10%	Quantile 90%	No. of observation
<i>Annual labor earnings</i> ⁽¹⁾ :				
Age 59	24.0	18.6	31.0	2323
Age 62	23.3	18.3	30.4	1286
Age 65	22.8	16.7	30.7	176
Age 68	24.0	19.5	35.8	23
<i>Entitled benefits upon retirement at 65</i> ⁽¹⁾ :				
Age 59	16.8	14.0	20.4	2342
Age 62	16.7	14.0	20.5	1332
Age 65	16.4	13.7	20.2	557
Age 68	17.2	14.4	21.4	30
<i>Retirement age</i> ⁽²⁾	64.5	62.0	69.0	3,129

⁽¹⁾In thousands of U.S. dollars of 2003. Exchange rate: U.S.\$ 1 \approx 8.1 Swedish kronas.

⁽²⁾Based on Kaplan–Meier estimator of the CDF of retirement age.




Figure 1. Kaplan–Meier estimator of CDF of retirement age. Sample of 3,129 male blue-collar workers.




Figure 2. Estimated labor earnings function. Married men with $\omega = 0$.

risk. The discount factor at age t is $\beta_t \equiv \bar{\beta} \lambda_t$, where λ_t is the mortality rate at age t and $\bar{\beta}$ is equal to 0.95.

5.4 Wage Function, Choice Probabilities, and Transitions

(a) *Wage equation and stochastic process of the wage shock.* From Section 5.1, we have that

$$\ln(W_{it}) = h_W(\mathbf{x}_{it}) + \rho(\ln W_{i,t-1} - h_W(\mathbf{x}_{i,t-1})) + \xi_{i,t+1}$$

$$\text{if } a_{it} = a_{i,t-1} = 1. \quad (39)$$

I consider a polynomial series approximation to the functions h_W and ρ . I use Akaike’s information criterion (AIC) to choose the orders of these polynomials. Given the polynomial series approximation, Equation (39) is a regression model with nonlinear restrictions between the parameters. I estimate this model using nonlinear least squares. Table 2 presents the estimated functions, and Figure 2 shows the age profile of annual labor earnings for the case of married men with median value of ω . The functions $h_W(\cdot)$ and $\rho(\cdot)$ are quadratic. Labor earnings

reach their life-cycle maximum at age 62, both for married and for unmarried men. There is very strong persistence in the evolution of the shock ω . The significance of the regressor “square ω_{t-1} ” in the estimated function $\rho(\cdot)$ illustrates that this persistence is larger at high values of ω . Given the residuals of this regression, $\{\hat{\xi}_{it}\}$, I estimate the density function of the innovation ξ using a kernel method with Gaussian kernel and cross-validation for the choice of bandwidth. The estimated density, not reported here, has strong kurtosis at zero and a thick left tail. Innovations in labor earnings are typically very close to zero, but there is a nonnegligible probability of having a very negative shock (e.g., a period of unemployment).

(b) *Transition rule of pension points.* Table 3 presents the estimation of different specifications of the transition function $h_{pp}(t, pp_t)$. Following AIC, I select the specification in column (3).

(c) *Transition of marital status.* Not surprisingly (especially for males, given that their mortality rate is higher than for females), there is very high persistence in marital status at these ages. The estimated probabilities $\Pr(m_{t+1} = 1 | m_t = 1, t)$ and $\Pr(m_{t+1} = 0 | m_t = 0, t)$ are very close to one for every age t .

Table 2. Wage equation and stochastic process of wage shock

Wage equation		Process of wage shock	
$lm(W) = h_{W,0} + h_{W,1}married + h_{W,2}age + h_{W,3}age^2$		$\omega_t = \rho_1 \omega_{t-1} + \rho_2 \omega_{t-1}^2 + \xi_t$	
Parameter	Estimate (Std. error)	Parameter	Estimate (Std. error)
$h_{W,0}$ (constant)	2.5709 (0.2064)	ρ_1	0.9144 (0.0038)
$h_{W,1}$ (married)	0.0816 (0.0020)	ρ_2	0.0952 (0.0023)
$h_{W,2}$ (age)	0.0170 (0.0072)	Std. dev. ξ_t	0.163
$h_{W,3}$ (age ²)	−0.00014 (0.00006)		
R-square	0.050	R-square	0.701
No. of observations	30,630	No. of observations	30,630

Table 3. Transition rule of pension points

Parameter	(1)	(2)	(3)	(4)	(5)
	Estimate (Std. error)	Estimate (Std. error)	Estimate (Std. error)	Estimate (Std. error)	Estimate (Std. error)
<i>constant</i>	0.0326 (0.0005)	0.0535 (0.0011)	0.0782 (0.0012)	0.0532 (0.0148)	0.0279 (0.0160)
$\ln(pp)$	0.9837 (0.0003)	0.9840 (0.0003)	0.9437 (0.0009)	0.9438 (0.0009)	0.9634 (0.0049)
<i>age</i>		-0.0004 (0.00002)	-0.0004 (0.00002)	0.0005 (0.0005)	0.0009 (0.0005)
$\ln(pp)^2$			0.0161 (0.0004)	0.0161 (0.0004)	0.0161 (0.0004)
age^2				-7.7×10^{-6} (4.5×10^{-6})	-6.64×10^{-6} (4.53×10^{-6})
$\ln(pp) \times age$					-0.0003 (0.0001)
Std. dev. η	0.012	0.012	0.012	0.012	0.012
R-square	0.9969	0.9970	0.9971	0.9971	0.9971
No. of observations	30,630	30,630	30,630	30,630	30,630

(d) *Choice probability function P.* For the choice probability function, I have considered a logit model with a polynomial specification in terms of the state variables (\mathbf{x}, ω) . Note that this specification of $P(\cdot)$ does not imply that we impose a logistic probability distribution for the unobservable $\tilde{\eta}_t$. The polynomial specification of the function P with respect to ω implies a flexible specification of the CDF $F_{\tilde{\eta}}$. In fact, as shown below, our estimate of P generates an estimate of the CDF of $\tilde{\eta}_t$ that is significantly different from the CDF of a logistic random variable. Table 4 presents the estimated logit model, and Figure 3 shows the age profile of the probability of working. The selected specification is the one that minimizes AIC within the class of polynomial functions of $(t, m, \ln(pp), \omega)$. It is important to note that including age dummies does improve very significantly the goodness of fit and even AIC. However, a specification with age dummies generates a very implausible prediction, that is, that the probability of working declines very significantly at age 65 and then increases very abruptly at age 66. Given the age profile of pension benefits and of the estimated

labor earnings function, the only way in which the model can generate this nonmonotonicity is through the utility of leisure. The difference between the utility of leisure if working minus the utility of leisure if retired should decline (abruptly) at age 65 and then increase very significantly at age 66. The increase at age 66 is very unrealistic. A more plausible interpretation of this evidence is that the nonmonotonicity is due to persistent unobserved heterogeneity in health status that we do not account for in this model. Individuals who continue working after age 65 may have significantly better health status than those who retire at age 65. Since we do not observe health status, and the model does not allow for persistent unobserved heterogeneity in the utility of leisure, it seems as if the utility of leisure if working increases at age 66. To avoid this spurious interpretation, I omit age dummies in the estimation of the choice probability function.

Table 4. Reduced form probability of working logit model

Variable	Estimate	(SE)
<i>constant</i>	97.46	(21.57)
<i>married</i>	0.1038	(0.0888)
ω	12.42	(3.441)
ω^2	0.3250	(0.1455)
$\ln(pp)$	-10.76	(5.352)
$\ln(pp)^2$	-0.3325	(0.2472)
<i>age</i>	-2.119	(0.6496)
age^2	0.0096	(0.0051)
$age \times \ln(pp)$	0.1768	(0.0856)
$age \times \omega$	-0.1666	(0.0509)
$\ln(pp) \times \omega$	-0.5352	(0.3526)
log-likelihood	-2359.75	
Akaike AIC	4741.49	
No. of observations	31,464	

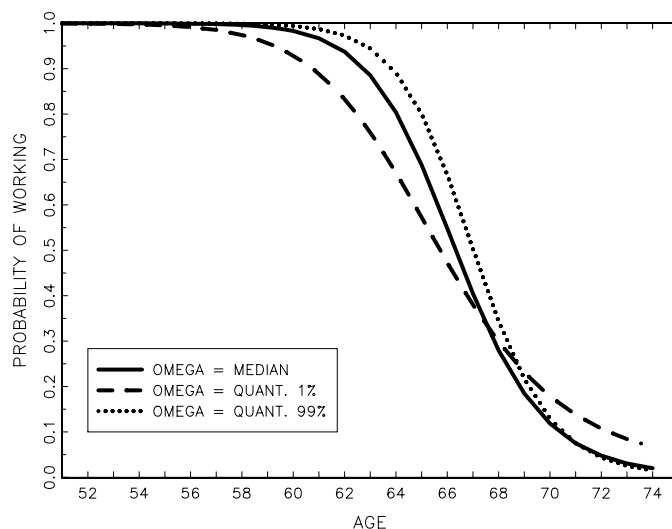


Figure 3. Estimated probability of working $P(t, m, pp, \omega)$. At median values of (m, pp) and at quantiles 1%, 50%, and 99% of ω .

5.5 Distribution of the Unobserved Utility of Leisure

I briefly discuss here the estimates of the distribution of unobservables $F_{\tilde{\eta}}$ and of the value function \tilde{V}_t^M associated with the utility of leisure. Figure 4 presents the estimated CDF of the unobservable $\tilde{\eta}_t$ at ages 58, 60, and 62 for a married man with median pension points. For any age, the distribution is very asymmetrical around the median $\tilde{\eta} = 0$. For values lower than zero, the distribution is concentrated around zero. However, for values greater than zero, the distribution has a very thick tail. Furthermore, when the individual gets older, the right tail of the distribution becomes thicker. This effect is even stronger for ages older than 64. The reason is that in order to explain why some individuals are still working at ages older than 64, we need a nonnegligible probability that the unobservable $\tilde{\eta}$ (the unobserved utility of leisure when working minus the unobserved utility of leisure when not working) takes very high values; that is, we need a nonnegligible probability of a very high preference for working.

Figure 5 presents estimates of the value $\tilde{V}_t^M(\mathbf{x})$ at ages between 56 and 68 and for the median values of the state variables (m, pp) . In the context of this retirement model, this value function represents the value of leisure if retirement is delayed one more year (relative to the value of leisure if retired today). This value remains quite constant until age 65, and then it declines very rapidly after that age.

5.6 Policy Effects on the Distribution of Retirement Age

Let ra_i and ra_i^* be individual i 's retirement age under the factual and counterfactual social security policies, respectively. Let $(m_{i51}, pp_{i51}, \omega_{i51})$ be the values of marital status, public pension points, and wage shock at age 51 for individual i in the sample. Conditional on this initial value of the state variables, the retirement ages ra_i^* and ra_i are random variables. The expected effect of the policy change on individual i 's retirement age is $TE_i \equiv E(ra_i^* - ra_i | m_{i51}, pp_{i51}, \omega_{i51})$. The (unconditional) average treatment effect is $ATE \equiv E(ra_i^* - ra_i)$. To estimate the ATE, first I estimate, for each individual i in the sample, the




Figure 4. CDF of the unobservable $\tilde{\eta}$.




Figure 5. Value $\tilde{V}_t(m, pp)$ at median values of (m, pp) .

probability distribution of the random variables ra_i^* and ra_i conditional on the individual's initial condition $(m_{i51}, pp_{i51}, \omega_{i51})$. Then, I use these distributions to obtain an estimate of TE_i , and finally I estimate the (unconditional) average treatment effect using the sample mean $N^{-1} \sum_{i=1}^N \widehat{TE}_i$.

I describe here some details of the estimation of the distributions of the retirement ages ra_i^* and ra_i conditional on the individual's state at age 51, $(m_{i51}, pp_{i51}, \omega_{i51})$. Retirement age depends on the realization of the state variables $(m_{it}, pp_{it}, \omega_{it})$ between ages 51 and 75. For each individual in the sample, I use the estimated transition probabilities of the state variables and the initial condition $(m_{i51}, pp_{i51}, \omega_{i51})$ to simulate 200 complete histories of the state variables. Given a simulated history, say $\{\tilde{m}_{it}, \tilde{pp}_{it}, \tilde{\omega}_{it} : t = 51, 52, \dots, 75\}$, and using the backward induction procedure described in Section 4, I obtain estimates of the factual and counterfactual choice probabilities $P_t(\tilde{m}_{it}, \tilde{pp}_{it}, \tilde{\omega}_{it})$ and $P_t^*(\tilde{m}_{it}, \tilde{pp}_{it}, \tilde{\omega}_{it})$ at every age $t \in [51, 75]$. These hazard rates provide the distribution of retirement age of individual i conditional on the simulated history $\{\tilde{m}_{it}, \tilde{pp}_{it}, \tilde{\omega}_{it} : t = 51, 52, \dots, 75\}$:

$$\begin{aligned} \tilde{\Pr}(ra_i = t) &= \left[\prod_{j=51}^{t-1} P_t(\tilde{m}_{ij}, \tilde{pp}_{ij}, \tilde{\omega}_{ij}) \right] \\ &\quad \times (1 - P_t(\tilde{m}_{it}, \tilde{pp}_{it}, \tilde{\omega}_{it})), \\ \tilde{\Pr}(ra_i^* = t) &= \left[\prod_{j=51}^{t-1} P_t^*(\tilde{m}_{ij}, \tilde{pp}_{ij}, \tilde{\omega}_{ij}) \right] \\ &\quad \times (1 - P_t^*(\tilde{m}_{it}, \tilde{pp}_{it}, \tilde{\omega}_{it})). \end{aligned} \tag{40}$$

To obtain the distribution of retirement age ra_i (and ra_i^*) conditional only on the state at age 51 (and not on the whole history), I average the distributions $\tilde{\Pr}(ra_i = t)$ over the 200 simulated histories. Though 200 seems a small number of simulations, notice that the transitions of marital status and pension points are almost deterministic, and ω is also a very persistent state variable. In fact, most of the simulated histories of an individual are very similar to each other such that the distribution of retirement ages using 200 simulations is practically equal to the

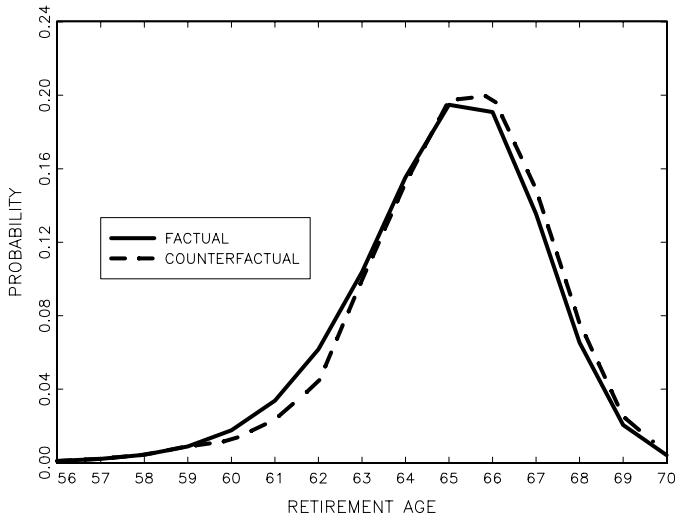


Figure 6. Factual and counterfactual distribution of retirement age.

one using only 50 simulations. Based on these distributions, we can obtain an estimate of the expected retirement age for individual i under the factual and counterfactual policies, that is, $\sum_{t=51}^{75} \Pr(ra_i = t) * t$ and $\sum_{t=51}^{75} \Pr(ra_i^* = t) * t$, respectively. Therefore, an estimate of TE_i is $\sum_{t=51}^{75} [\Pr(ra_i^* = t) - \Pr(ra_i = t)] * t$.

Figure 6 presents the estimated PDFs of retirement age under the factual and counterfactual policies for an individual with median values of the initial conditions $(m_{i51}, pp_{i51}, \omega_{i51})$. The figure shows that the effect of the policy change on the retirement behavior of the median individual in this population is very minor. Figure 7 presents the empirical distribution (PDF) of the individual-specific treatment effects TE_i for the 3,129 individuals in the sample. As expected, every individual in the sample increases his (expected) retirement age. However, this increase is very small for most individuals, with almost all the values between 0.1 and 0.4 years. Table 5 reports different characteristics of the distribution of the treatment effect as




Figure 7. Empirical distribution of individual-specific expected treatment effects: $TE_i \equiv E(ra_i^* - ra_i | m_{i51}, pp_{i51}, \omega_{i51})$.

well as bootstrap standard errors. I also report the estimated effect in KPS (table III, column S1a). KPS find that the effect of the policy change on the average retirement age is 0.78 years (65.58 – 64.80). They do not report standard errors. Instead, our estimate of the ATE is 0.22 with a standard error of 0.0743. The test of specification that I described in Section 4 clearly rejects the parametric specification for the evaluation of this policy change. Given that the most important difference between the KPS parametric model and our model is in the distribution of the unobservable $\tilde{\eta}$, it seems that the age invariant logistic specification in KPS has a significant incidence in the ATE and that we can reject that feature of their model. To verify this supposition, I have estimated a semiparametric version of the model using the same approach as before but imposing a logistic distribution for $\tilde{\eta}$. The last row of Table 5 reports the ATE for this semiparametric model. It shows that the estimated effect is very similar to the one reported by KPS.

This empirical application, though useful to illustrate the implementation of the method, has some limitations. First, a more complete analysis should include estimates of the model under different values of the relative risk aversion parameter within the class of a constant relative risk aversion (CRRA) utility function. That analysis would include risk neutrality (i.e., relative risk aversion equal to zero) and log-utility function (i.e., relative risk aversion equal to one) as extreme cases. I have presented here results for one of these extreme cases. Second, as discussed in Section 5.4(d), the lack of information on health status can introduce significant biases in our estimates of choice probabilities and policy effects. Third, the model ignores labor market institutions such as mandatory retirement at age 65, which may be generating important incentives to retire at age 65 in the population under study. I would expect that, in reality, companies’ mandatory retirement rules would change together with the social security policy; that is, mandatory retirement might move from age 65 to age 68. In this sense, the estimated ATE measures the ceteris paribus of the change in social security pension benefits keeping other (related) social norms and labor market institutions constant.

6. SUMMARY AND CONCLUSIONS

This article presents a nonparametric approach to evaluating the behavioral and welfare effects of counterfactual policies using a dynamic structural model. The nonparametric structural model retains all of the economic assumptions of a structural model (e.g., exogeneity/endogeneity assumptions, equilibrium concept, rational expectations, transition rules, independence assumptions, agents’ information), but it is more robust than a parametric model because it relaxes parametric assumptions. There are situations in which a parametric model can be preferred, either because of its parsimony or to obtain more precise estimates. In these cases, our nonparametric approach can be used to test for the validity of a parametric specification and to search for a valid parametric model.

APPENDIX

Proof of Proposition 1

Let $P^0(X, W)$ be the choice probability function in the population, that is, $P^0(X, W) \equiv \Pr(a = 1 | X, W)$. It is clear that P^0

Table 5. Distribution of treatment effects on retirement age

Model	ATE	Std. error of ATE ⁽¹⁾	Median TE	Test ⁽²⁾ (p-value)
Nonparametric	0.237	0.0743	0.234	
Parametric (KPS) ⁽³⁾	0.78			5.864 (0.0000)
Semiparametric (logit $\tilde{\eta}$)	0.612	0.0621	0.607	4.050 (0.0000)

NOTE: (1) Bootstrap standard errors based on 500 bootstrap samples with the same number of individuals as the original sample. (2) Test statistic = $|ATE_{parametric} - ATE_{nonparametric}| / SE(ATE_{parametric} - ATE_{nonparametric})$. (3) KPS do not report a standard error for the estimated average effect. For the test, I use the same value of $SE(ATE_{parametric} - ATE_{nonparametric})$ as I obtained for the semiparametric model.

is identified on its support $S_X \times \mathbb{R}$. By definition, $P^0(X, W) = F_{\eta|X}^0(W + \varphi^0(X))$. By the conditions of Proposition 1, the function $P^0(X, W)$ is strictly increasing in W and there is a unique value W that solves the equation $P^0(X, W) = 0.5$. Let $w_0(X)$ be the value of W that solves the equation $P^0(X, W) = 0.5$. By the zero conditional median of η , we have that $w_0(X) + \varphi^0(X) = 0$. Therefore, for any $X \in S_X$, we have that $\varphi^0(X) = -w_0(X)$, and $\varphi^0(X)$ is identified. Then, for any real u :

$$\begin{aligned} F_{\eta|X}^0(u) &= \Pr(\eta \leq u|X, u) \\ &= \Pr(\eta \leq W + \varphi^0(X)|X, u, W = u - \varphi^0(X)) \\ &= \Pr(a = 1|X, u, W = u - \varphi^0(X)) = P^0(X, u - \varphi^0(X)). \end{aligned}$$

Therefore, $F_{\eta|X}^0$ is identified on \mathbb{R} .

Proof of Proposition 2

By definition, the optimal choice probability P is equal to $F_{\tilde{\eta}}(u(1) - u(0))$. Since $F_{\tilde{\eta}}$ is continuous and strictly increasing, we have that $u(1) - u(0) = F_{\tilde{\eta}}^{-1}(P)$, where $F_{\tilde{\eta}}^{-1}$ is the inverse function of $F_{\tilde{\eta}}$. Then, by the definition of social surplus function, we have that

$$\begin{aligned} G &= \int \max\{0, u(1) - u(0) - \tilde{\eta}\} dF_{\tilde{\eta}}(\tilde{\eta}) \\ &= \int \max\{0, F_{\tilde{\eta}}^{-1}(P) - \tilde{\eta}\} dF_{\tilde{\eta}}(\tilde{\eta}). \end{aligned}$$

Proof of Proposition 3

For notational simplicity I omit ω as an argument in the different functions. By definition, the choice-specific value function can be written as,

$$v_t(a, \mathbf{x}) = Y_t(a, \mathbf{x}) + M_t(a, \mathbf{x}) + \beta \int \tilde{V}_{t+1}(\mathbf{x}_{t+1}) dF(\mathbf{x}_{t+1}|a, \mathbf{x}),$$

where \tilde{V}_t is the integrated value function as defined in Equation (10). Given the definitions of the integrated value function and of the surplus function $G(P, F)$, in Proposition 2, we have that: $\tilde{V}_t(\mathbf{x}_t) = v_t(0, \mathbf{x}_t) + G(P_t(\mathbf{x}_t), F_{\tilde{\eta}, t})$. Solving this expression into the previous formula of the choice-specific value function, we get

$$\begin{aligned} v_t(a, \mathbf{x}) &= Y_t(a, \mathbf{x}) + M_t(a, \mathbf{x}) + \beta \int [v_{t+1}(0, \mathbf{x}_{t+1}) \\ &\quad + G(P_{t+1}(\mathbf{x}_{t+1}), F_{\tilde{\eta}, t+1})] dF(\mathbf{x}_{t+1}|a, \mathbf{x}, \omega). \end{aligned}$$

Therefore, at the last period T , $v_T(a, \mathbf{x}) = V_T^Y(a, \mathbf{x}) + V_T^M(a, \mathbf{x})$, with $V_T^Y(a, \mathbf{x}) = Y_T(a, \mathbf{x})$, $V_T^M(a, \mathbf{x}) = M_T(a, \mathbf{x})$, and $V_T^{OPT}(a, \mathbf{x}) = 0$. Going backward, at period $T - 1$, we have that

$$\begin{aligned} v_{T-1}(a, \mathbf{x}) &= Y_{T-1}(a, \mathbf{x}) + M_{T-1}(a, \mathbf{x}) \\ &\quad + \beta \int [V_T^Y(0, \mathbf{x}_T) + V_T^M(0, \mathbf{x}_T) \\ &\quad + G(P_T(\mathbf{x}_T), F_{\tilde{\eta}, T})] dF(\mathbf{x}_T|a, \mathbf{x}) \\ &= V_{T-1}^Y(a, \mathbf{x}) + V_{T-1}^M(a, \mathbf{x}) + V_{T-1}^{OPT}(a, \mathbf{x}), \end{aligned}$$

where $V_{T-1}^Y(a, \mathbf{x}) = Y_T(a, \mathbf{x}) + \beta \int V_T^Y(0, \mathbf{x}_T) dF(\mathbf{x}_T|a, \mathbf{x}, \omega)$, $V_{T-1}^M(a, \mathbf{x}) = M_T(a, \mathbf{x}) + \beta \int V_T^M(0, \mathbf{x}_T) dF(\mathbf{x}_T|a, \mathbf{x}, \omega)$, and where $V_{T-1}^{OPT}(a, \mathbf{x}) = \beta \int [V_T^{OPT}(0, \mathbf{x}_T) + G(P_T(\mathbf{x}_T), F_{\tilde{\eta}, T})] dF(\mathbf{x}_T|a, \mathbf{x}, \omega)$. Applying backward induction, we obtain that for any t , $v_t(a, \mathbf{x}) = V_t^Y(a, \mathbf{x}) + V_t^M(a, \mathbf{x}) + V_t^{OPT}(a, \mathbf{x})$.

Proof of Proposition 4

The proof proceeds in three steps, (a), (b), and (c).

(a) The value function \tilde{V}_t^Y only depends on the sequence of outcome functions and transition probabilities $\{Y_{t+j}(\cdot), F_{\mathbf{x}, t+j}(\cdot), F_{\omega, t+j}(\cdot) : j > 0\}$. Therefore, \tilde{V}_t^Y is identified at every period t .

(b) At any period t , define the variable $W_t \equiv \tilde{V}_t^Y(\mathbf{x}_t, \omega_t) + \tilde{V}_t^{OPT}(\mathbf{x}_t, \omega_t)$. Given this definition, the discrete choice model at period t can be represented as $a_t = I\{\tilde{\eta}_t \leq W_t + \tilde{V}_t^M(\mathbf{x}_t)\}$. For the moment, suppose that $\tilde{V}_t^{OPT}(\mathbf{x}_t, \omega_t)$ (and therefore W_t) is observable to the econometrician. Now, we show that, under Assumptions 1–6, this discrete choice model satisfies conditions (i)–(iv) of Proposition 1: (i) $\tilde{\eta}_t$ and W_t are independent conditional on \mathbf{x}_t [by Assumption 1]; (ii) for any $\mathbf{x}_t \in S_X$, the CDF $F_{\tilde{\eta}_t}$ is strictly increasing in $\tilde{\eta}_t$, has support \mathbb{R} , and $F_{\tilde{\eta}_t}(0) = 0.5$ [by Assumption 3]; (iii) the distribution of W_t conditional on \mathbf{x}_t has a Lebesgue density that is everywhere positive on \mathbb{R} [by Assumptions 3(C), 3(D), and 4]; and (iv) for any $\mathbf{x}_t \in S_X$, there is a value $W_t \in \mathbb{R}$ such that $W_t + \tilde{V}_t^M(\mathbf{x}_t) = 0$ [by Assumptions 3(C), 3(D), and 4]. Therefore, if W_t is observable, then by Proposition 1 the function $\tilde{V}_t^M(\cdot)$ is identified on S_X , and the function $F_{\tilde{\eta}_t}$ is identified on \mathbb{R} .

(c) It remains to prove that, at any period t , $\tilde{V}_t^{OPT}(\mathbf{x}_t, \omega_t)$ is identified (and therefore W_t is observable). We apply backward induction. At the last period, it is clear that $\tilde{V}_T^{OPT}(\mathbf{x}_T, \omega_T) = 0$. Therefore, W_T is observable and, applying point (b), the functions \tilde{V}_T^M and $F_{\tilde{\eta}, T}$ are identified. The identification of $P_T(\cdot)$ and

$F_{\tilde{\eta}T}(\cdot)$ implies the identification of the value function \tilde{V}_{T-1}^{OPT} . To see this, note that $\tilde{V}_{T-1}^{OPT}(\mathbf{x}, \omega) = V_{T-1}^{OPT}(1, \mathbf{x}, \omega) - V_{T-1}^{OPT}(0, \mathbf{x}, \omega)$, and

$$\begin{aligned} \tilde{V}_{T-1}^{OPT}(\mathbf{x}, \omega) &= \beta \int G(P_T(\mathbf{x}', \omega'), F_{\tilde{\eta}T}) d\tilde{F}_{T-1}(\mathbf{x}', \omega' | \mathbf{x}, \omega) \\ &= \beta \int \int \max\{F_{\tilde{\eta}T}^{-1}(P_T(\mathbf{x}', \omega')) - u, \\ &\quad 0\} dF_{\tilde{\eta}T}(u) d\tilde{F}_{T-1}(\mathbf{x}', \omega' | \mathbf{x}, \omega), \end{aligned}$$

where $\tilde{F}_{T-1}(\mathbf{x}', \omega' | \mathbf{x}, \omega) \equiv [F_{\mathbf{x}, T-1}(\mathbf{x}' | 1, \mathbf{x}) - F_{\mathbf{x}, T-1}(\mathbf{x}' | 0, \mathbf{x})] F_{\omega, T-1}(\omega' | \omega)$. Then, $W_{T-1} = \tilde{V}_{T-1}^Y(\mathbf{x}_{T-1}, \omega_{T-1}) + \tilde{V}_{T-1}^{OPT}(\mathbf{x}_{T-1}, \omega_{T-1})$ is observable. Since W_{T-1} is observable, point (b) above implies that the functions \tilde{V}_{T-1}^M and $F_{\tilde{\eta}, T-1}$ are identified. And, in turn, this implies that the function \tilde{V}_{T-2}^{OPT} is identified. Applying this argument backward we can show the identification of the function \tilde{V}_t^{OPT} at any period t .

Proof of Proposition 5

The proof proceeds in three steps, (a), (b), and (c).

(a) By Proposition 4, the functions P_t , \tilde{V}_t^Y , \tilde{V}_t^M , \tilde{V}_t^{OPT} , and $F_{\tilde{\eta}t}$ are identified.

(b) By Assumptions 3 and 4, for any value of $\mathbf{x} \in S_X$, the function $\tilde{V}_t^Y(\mathbf{x}, \omega) + \tilde{V}_t^{OPT}(\mathbf{x}, \omega)$ is strictly monotonically increasing in ω and its range of variation \mathbb{R} . Therefore, for any $(\mathbf{x}, \omega) \in S_X \times \mathbb{R}$, the threshold value $\omega_t^*(\mathbf{x}, \omega)$ exists and is unique, that is, $\omega_t^*(\mathbf{x}, \omega)$ is a well-defined function from $S_X \times \mathbb{R}$ into \mathbb{R} . Using the definition of the function $\omega_t^*(\mathbf{x}, \omega)$, we have that $P_t^*(\mathbf{x}, \omega) = P_t(\mathbf{x}, \omega_t^*(\mathbf{x}, \omega))$.

(c) At period T , there is no future and therefore the decision problem is static. $\omega_T^*(\mathbf{x}, \omega)$ is implicitly defined as the value $\omega_T^* \in \mathbb{R}$ that solves the equation $\tilde{Y}_T(\mathbf{x}, \omega_T^*) = \tilde{Y}_T^*(\mathbf{x}, \omega)$. It is clear that the function ω_T^* is identified over $S_X \times \mathbb{R}$. Therefore, $P_T^*(\mathbf{x}, \omega) = P_T(\mathbf{x}, \omega_T^*(\mathbf{x}, \omega))$ is also identified over $S_X \times \mathbb{R}$. Given P_T^* and $F_{\tilde{\eta}|x, T}$, the function $\tilde{V}_{T-1}^{OPT}(\mathbf{x}, \omega; \{P_T^*\})$ is known. Now, consider the counterfactual choice probability at any period $t < T$ given that the function $\tilde{V}_t^{OPT}(\mathbf{x}, \omega; \{P_{t+j}^* : j > 0\})$ is known. The function $\omega_t^*(\mathbf{x}, \omega)$, that is implicitly defined as the solution in ω_t^* to the equation $\tilde{V}_t^Y(\mathbf{x}, \omega_t^*) + \tilde{V}_t^{OPT}(\mathbf{x}, \omega_t^*; \{P_{t+j}^* : j > 0\}) = \tilde{V}_t^{Y*}(\mathbf{x}, \omega) + \tilde{V}_t^{OPT}(\mathbf{x}, \omega; \{P_{t+j}^* : j > 0\})$, is identified over $S_X \times \mathbb{R}$. This implies that $P_t^*(\mathbf{x}, \omega) = P_t(\mathbf{x}, \omega_t^*(\mathbf{x}, \omega))$ is identified.

Proof of Proposition 6

By the definition of the value functions \bar{V}_t^* and \bar{V}_t , we have that

$$\begin{aligned} \bar{V}_t^*(\mathbf{x}, \omega) - \bar{V}_t(\mathbf{x}, \omega) &= [v_t^*(0, \mathbf{x}, \omega) - v_t(0, \mathbf{x}, \omega)] \\ &\quad + G(P_t^*(\mathbf{x}, \omega), F_{\tilde{\eta}t}) - G(P_t(\mathbf{x}, \omega), F_{\tilde{\eta}t}). \end{aligned}$$

We know from Proposition 3 that $v(0, \mathbf{x}, \omega) = V_t^Y(0, \mathbf{x}, \omega) + V_t^M(0, \mathbf{x}, \omega) + V_t^{OPT}(0, \mathbf{x}, \omega)$ and $v^*(0, \mathbf{x}, \omega) = V_t^{Y*}(0, \mathbf{x}, \omega) + V_t^M(0, \mathbf{x}, \omega) + V_t^{OPT*}(0, \mathbf{x}, \omega)$. Solving these expressions into the equation for $\bar{V}_t^*(\mathbf{x}, \omega) - \bar{V}_t(\mathbf{x}, \omega)$, we can get Equation (22).

Propositions 4 and 5 establish the identification of all the functions in the right-hand side of this equation.

ACKNOWLEDGMENTS

I would like to thank the coeditor, three anonymous referees, Russell Davidson, Jim Heckman, Chun-Yu Ho, Rasmus Lentz, Thierry Magnac, Rosa Matzkin, Pedro Mira, Salvador Navarro, Whitney Newey, John Rust, Steven Stern, Tom Stoker, Chris Taber, Gustavo Vicentini, and audiences at the North American winter meeting of the Econometric Society (Boston, 2006), the Canadian Econometric Study Group (Vancouver, 2005), the Society of Economic Dynamics conference (Vancouver, 2006), University of Chicago, Duke, Georgia, Indiana (CAEPR), MIT, Northwestern, Pittsburgh, Princeton, Washington University–St. Louis, and Western Ontario.

[Received March 2007. Revised October 2008.]

REFERENCES

- Aakvik, A., Heckman, J., and Vytlačil, E. (2005), “Estimating Treatment Effects for Discrete Outcomes When Responses to Treatment Vary: An Application to Norwegian Vocational Rehabilitation Programs,” *Journal of Econometrics*, 125, 15–51. [204]
- Bajari, P., and Hong, H. (2006), “Semiparametric Estimation of a Dynamic Game of Incomplete Information,” Technical Working Paper 320, NBER. [202]
- Berkovec, J., and Stern, S. (1991), “Job Exit Behavior of Older Men,” *Econometrica*, 59, 189–210. [209]
- Blau, D. (2008), “Retirement and Consumption in a Life Cycle Model,” *Journal of Labor Economics*, 26, 35–71. [209]
- Bound, J., Stinebrickner, T., and Waidmann, T. (2009), “Health, Economic Resources and the Work Decisions of Older Men,” *Journal of Econometrics*, to appear. [209]
- French, E. (2005), “The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior,” *Review of Economic Studies*, 72, 395–427. [209]
- Heckman, J. (2000), “Causal Parameters and Policy Analysis in Economics: A Twentieth Century Retrospective,” *The Quarterly Journal of Economics*, 115, 45–97. [201]
- Heckman, J., and Navarro, S. (2007), “Dynamic Discrete Choice and Dynamic Treatment Effects,” *Journal of Econometrics*, 136, 341–396. [202,203]
- Heckman, J., and Robb, R. (1985), “Alternative Methods for Evaluating the Impact of Interventions,” in *Longitudinal Analysis of Labor Market Data*, eds. J. Heckman and B. Singer, Cambridge: Cambridge University Press. [201]
- Heckman, J., and Smith, J. (1998), “Evaluating the Welfare State,” in *Frisch Centenary. Econometric Monograph Series*, Cambridge: Cambridge University Press. [201]
- Heckman, J., and Vytlačil, E. (1999), “Local Instrumental Variables and Latent Variable Models for Identifying and Bounding Treatment Effects,” *Proceedings of the National Academy of Sciences*, 96, 4730–4734. [201]
- (2005), “Structural Equations, Treatment Effects, and Econometric Policy Evaluation,” *Econometrica*, 73, 669–738. [201]
- Imbens, G., and Newey, W. (2006), “Identification and Estimation of Triangular Simultaneous Equations Models Without Additivity,” manuscript, MIT, Department of Economics. [205]
- Karlstrom, A., Palme, M., and Svensson, I. (2004), “A Dynamic Programming Approach to Model the Retirement Behaviour of Blue-Collar Workers in Sweden,” *Journal of Applied Econometrics*, 19, 795–807. [209,210]
- Kasahara, H. (2009), “Temporary Increases in Tariffs and Machine Replacement: The Chilean Experience,” *Journal of Business & Economic Statistics*, 27, 113–127. [204]
- Magnac, T., and Thesmar, D. (2002), “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, 70, 801–816. [201]
- Mammen E. (1991), “Estimating a Smooth Monotone Regression Function,” *The Annals of Statistics*, 19, 724–740. [207]
- Manski, C. (1990), “Nonparametric Bounds on Treatment Effects,” *American Economic Review, Papers and Proceedings*, 80, 319–323. [201]
- Marschack, J. (1953), “Economic Measurements for Policy and Prediction,” in *Studies in Econometric Method*, eds. W. Hood and T. Koopmans, New York: Wiley. [201]

- Matzkin, R. (1992), "Nonparametric and Distribution-Free Estimation of the Binary Choice and the Threshold Crossing Models," *Econometrica*, 60, 239–270. [202]
- (1994), "Restrictions of Economic Theory in Nonparametric Methods," in *Handbook of Econometrics*, Vol. 4, ed. C. F. Engel and D. L. McFadden, Amsterdam, The Netherlands: Elsevier. [202,205]
- McFadden, D. (1981), "Structural Discrete Probability Models Derived From Theories of Choice," in *Structural Analysis of Discrete Data and Econometric Applications*, eds. C. Manski and D. McFadden, Cambridge, MA: MIT Press. [205]
- Mukerjee, H. (1988), "Monotone Nonparametric Regression," *The Annals of Statistics*, 16, 741–750. [207]
- Newey, W. (1994), "Kernel Estimation of Partial Means and a General Variance Estimator," *Econometric Theory*, 10, 233–253. [208]
- Olley, S., and Pakes, A. (1996), "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64, 1263–1297. [205]
- Rust, J. (1994), "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics*, Vol. 4, eds. R. E. Engle and McFadden, Amsterdam, The Netherlands: North-Holland. [201,203]
- Rust, J., and Phelan, C. (1997), "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets," *Econometrica*, 65, 781–832. [209,211]
- Taber, C. (2000), "Semiparametric Identification and Heterogeneity in Discrete Choice Dynamic Programming Models," *Journal of Econometrics*, 96, 201–229. [202]
- van der Klaauw, W., and Wolpin, K. (2008), "Social Security and the Retirement and Savings Behavior of Low Income Households," *Journal of Econometrics*, 145, 21–42. [209]