INTERNATIONAL ECONOMIC REVIEW
Vol. 56, No. 2, May 2015

ESTIMATING DYNAMIC DISCRETE CHOICE MODELS WITH HYPERBOLIC DISCOUNTING, WITH AN APPLICATION TO MAMMOGRAPHY DECISIONS*

BY HANMING FANG AND YANG WANG[1]

University of Pennsylvania, U.S.A.; Lafayette College, U.S.A.

We extend the semiparametric estimation method for dynamic discrete choice models using Hotz and Miller's (*Review of Economic Studies* 60 (1993), 497–529) conditional choice probability approach to the setting where individuals may have hyperbolic discounting time preferences and may be naive about their time inconsistency. We illustrate the proposed identification and estimation method with an empirical application of adult women's decisions to undertake mammography to evaluate the importance of present bias and naivety in the underutilization of this preventive health care. Our results show evidence for both present bias and naivety.

1. INTRODUCTION

Dynamic discrete choice models have been used to understand a wide range of economic behavior. The early dynamic discrete choice models that are empirically implemented tend to be parametric,[2] but recently, a growing list of authors have addressed the non- or semiparametric identification of dynamic discrete choice models. The earliest attempt in this regard is Hotz and Miller (1993), which pioneered the approach of using conditional choice probabilities to infer about choice-specific continuation values. Rust (1994a, 1994b) showed that the discount factor in standard dynamic discrete choice models is generically not identified; Magnac and Thesmar (2002) expanded Rust's nonidentification results and proposed exclusion restrictions that lead to the identification of the standard discount factor.

All of the above-mentioned literature models the impatience of the decision makers by assuming that agents discount future streams of utility or profits *exponentially* over time. As is now well known, exponential discounting is not just an analytically convenient assumption; without this assumption, intertemporal marginal rates of substitution will change as time passes, and preferences will be time inconsistent (see Strotz, 1956, p. 172). A recent theoretical literature has built on the work of Strotz (1956) and others to explore the consequences of relaxing the standard assumption of exponential discounting. Drawing both on experimental research and on common intuition, economists have built models of quasi-hyperbolic discounting to capture the tendency of decision makers to seize short-term rewards at the expense

*Manuscript received February 2014; revised September 2014.

[1] We are grateful to the editor and two anonymous referees for useful suggestions that improved the article. We would also like to thank Peter Arcidiacono, Patrick Bayer, Doug Bernheim, Xiaohong Chen, Stefano DellaVigna, Bo Honore, Han Hong, Joe Hotz, Yingyao Hu, Edward Kung, David Laibson, Thierry Magnac, Aprajit Mahajan, Ted O'Donoghue, Ulrike Malmandier, Andriy Norets, Dan Silverman, Frank Sloan, Xun Tang, Ken Wolpin, and seminar/conference participants at Cornell, Princeton, AEA Meetings in San Francisco (2009), Midwest Health Economics Conference in Chicago (2010), and SITE Psychology and Economics Conference (2011) for helpful comments. The first draft of the paper (January 2009) was completed when both authors were at Duke University. Fang gratefully acknowledges financial support from NSF Grant SES 0844845. We are responsible for all remaining errors. Please address correspondence to: Hanming Fang, Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 19104. E-mail: *hanming.fang@econ.upenn.edu*.

[2] The earliest formulation and estimation of parametric dynamic discrete choice models includes Wolpin (1984) for fertility choice, Miller (1984) for occupational choice, Pakes (1986) for patent renewal, and Rust (1987) for bus engine replacement.

of long-term preferences.[3] This literature studies the implications of time-inconsistent preferences and their associated problems of self-control for a variety of economic choices and environments.[4]

A small list of empirical papers that attempted to estimate dynamic models with hyperbolic discounting time preferences has followed the *parametric* approach (Laibson et al., 2007; Paserman, 2008; Fang and Silverman, 2009). Fang and Silverman (2009) empirically implement a dynamic structural model of labor supply and welfare program participation for never-married mothers with potentially time-inconsistent preferences. Using panel data on the choices of single women with children from the National Longitudinal Survey of Youth (NLSY 1979), they provide estimates of the degree of time-inconsistency and of its influence on the welfare take-up decision. For the particular population of single mothers with dependent children, they estimate the present bias and the standard discount factors to be 0.338 and 0.88, respectively, implying a one-year-ahead discount rate of 238%. Laibson et al. (2007) use the method of simulated moments (MSM) to estimate time preferences—both short- and long-run discount rates—from a structural buffer stock consumption model that includes many realistic features such as stochastic labor income, liquidity constraints, child and adult dependents, liquid and illiquid assets, revolving credit, and retirement. Under parametric assumptions, the model is identified from matching the model's predictions of retirement wealth accumulation, credit card borrowing, and consumption–income comovement with those observed in the data. Their benchmark estimates imply a 48.5% short-term annualized discount rate and a 4.3% long-term annualized discount rate. Paserman (2008) estimates the structural parameters of a job search model with hyperbolic discounting and endogenous search effort, using data on duration of unemployment spells and accepted wages from the NLSY 1979. Under parametric assumptions of the model, identification of the hyperbolic discounting parameters comes from the variation in the relative magnitude of unemployment duration and accepted wages. Indeed, he finds that the results are sensitive to the specific structure of the model and to the functional form assumption for the distribution of offered wages. For low-wage workers, he rejects the exponential discounting model and estimates a one-year discount rate of about 149%.[5] Chung et al. (2009) estimated a dynamic structural model of sales force response to a bonus-based compensation plan where the salesman might have hyperbolic discounting time preferences. Exploiting the bonus-based compensation structure, they found some evidence consistent with present bias.

None of the above papers allow for the possibility that a hyperbolic discounting decision maker may also be naive. More importantly, the identification of the present bias and standard discount factors in these papers is often based on parametric assumptions imposed on the model. To the best of our knowledge, it is not known whether dynamic discrete choice models

[3] A body of experimental research, reviewed in Ainslie (1992) and Loewenstein and Elster (1992), indicates that hyperbolic time discounting may parsimoniously explain some basic features of the intertemporal decision making that are inconsistent with simple models with exponential discounting. Specifically, it is not easy to reconcile standard decision models with exponential discounting with commonly observed preference reversals: Subjects choose the larger and later of two prizes when both are distant in time, but prefer the smaller but earlier one as both prizes draw nearer to the present. Recently, similar experimental strategies to infer about present bias based on choice reversal have been implemented in the field; see, for example, Gine et al. (2010) for a field experiment on smoking cessation and Gine et al. (2011) on savings choices. Also, see Rubinstein (2003) for an alternative explanation of preference reversals.

[4] For example, models of time-inconsistent preferences have been applied by Laibson (1997) and O'Donoghue and Rabin (1999a, 1999b) to consumption and savings, by Barro (1999) to growth, by Gruber and Koszegi (2001) to smoking decisions, by Krusell et al. (2002) to optimal tax policy, by Carrillo and Mariotti (2000) to belief formation, by Fang and Silverman (2004) to welfare program participation and labor supply of single mothers with dependent children, by Della Vigna and Paserman (2005) to job search, and by Della Vigna and Malmendier (2004, 2006) on contract designs.

[5] There are other inferential studies about discount rates that exploit specific clear-cut intertemporal trade-offs. For example, Hausman (1979) and Warner and Pleeter (2001) estimate discount rates ranging from 0% to 89% depending on the characteristics of the individual and intertemporal trade-offs at stake.

with hyperbolic discounting preferences can be semiparametrically identified using standard short-panel data that are typically used in these papers.[6,7]

In this article, we consider a standard dynamic discrete choice model where decision makers potentially exhibit hyperbolic discounting preferences in the form of a present bias factor (β), an exponential discounting factor (δ), and a potential naivety parameter ($\tilde{\beta}$) (as in O'Donoghue and Rabin, 1999a), and examine the conditions under which the primitive parameters of the model, including the three hyperbolic discounting time preference parameters, can be identified using short-panel (two periods) data. Our arguments proceed in two steps. First, we show in Proposition 1 that the preference parameters $\mathbf{u}^*$ can be uniquely identified from the observed choice probabilities *locally* if $(\beta, \tilde{\beta}, \delta)$ are at their true values $(\beta^*, \tilde{\beta}^*, \delta^*)$. Second, we show in Proposition 2 that if there exist *exclusion variables* that affect the transition probabilities of states over time but do not affect the decision makers' static payoff functions (formally stated in Assumption 5)—a condition similar to that in Magnac and Thesmar (2002) that is necessary for the identification of dynamic discrete choice models with standard exponential discounting—then we can *generically* identify the true values of the discount factors $(\beta^*, \tilde{\beta}^*, \delta^*)$. Taken together these two propositions imply that all the preference parameters $\langle \mathbf{u}^*, \beta^*, \tilde{\beta}^*, \delta^* \rangle$ are generically identified.

The intuition for why exclusion variables that affect the transition of state variables but not static payoffs might provide the source of identification for the discount factors can be described as follows. Consider two decision makers who share the same static-payoff relevant state variables but differ only in the exclusion variables. Because the exclusion variables only affect the transition of the payoff-relevant state variables, their effects on the choices in the current period will inform us about the degree to which the agents discount the future. The intuition for why β, $\tilde{\beta}$, and δ can be separately identified will be provided later in Section 3.2.

Our article also represents an interesting intermediate case between the literature on estimating dynamic discrete choice single-agent decision problems (see Miller, 1984, Wolpin, 1984, Pakes, 1986, Rust, 1987, and Hotz and Miller, 1993, for early contributions and Rust, 1994a, 1994b, and Aguirregabiria and Mira, 2007, for surveys) and the more recent literature on estimating dynamic games (see Pakes and McGuire, 1994, Pakes et al., 2007, Aguirregabiria and Mira, 2007a, Bajari et al., 2007, and Pesendorfer and Schmidt-Dengler, 2008, among others, and see Bajari et al., 2013, for a survey). We discuss the connections between our article and the above literature in Section 2.3.

We propose a maximum pseudolikelihood estimation method that is intimately related to our identification arguments. Monte Carlo experiments show that the proposed estimator performs well in samples with sizes comparable to what are typically available in empirical applications.

We illustrate our identification argument and estimation method with an empirical application of adult women's decisions to undertake mammography to investigate the role of time inconsistent preferences in the under-utilization of this preventive care. We consider a simple model where mammography can potentially lower the probability of death in the next two years and it may also lower the probability of bad health conditional on surviving in two years; however, undertaking mammography may involve immediate costs (most of which we would like to interpret as psychological and physical costs instead of financial costs). For the purpose of identifying the hyperbolic discounting preference parameters, we use several variables, including the indicator for whether the woman's mother is still alive and/or whether she died at age greater than 70, as the exclusion variables that do not enter the relevant instantaneous

[6] Fang and Silverman (2006) is an exception. They show that exponential discounting and hyperbolic discounting models are distinguishable, using an argument based on observed choice probabilities.

[7] Mahajan and Tarozzi (2011) discuss how data about elicited beliefs and responses to time-preference questions, together with observational data, can be used to identify and estimate hyperbolic discounting parameters in their study of the adoption of insecticide treated nets. They allow for three types of agents in terms of their time preferences—time consistent, completely naive, and completely sophisticated—but do not allow for partially naive agents. Moreover, their results apply in finite horizon settings only.

payoff function but affect the transition probability of other instantaneous payoff-relevant state variables.[8] Our estimates indicate that individuals exhibit both present bias and naivety: Our estimates of the present bias factor β range from 0.562 to 0.709 depending on specifications of the exclusion variables, and our estimate of the partial naivety parameter $\tilde{\beta}$ is 1 in all specifications. Hypothesis testing rejects both the null of no present bias and the null of no naivety with p-values less than 0.01. These results suggest that both present bias and naivety might have played an important role in the fact that nearly 25% of the women do not undertake mammography as advised by American Cancer Association, which is universally regarded as a very cost effective way for early detection of breast cancer (see Degnan et al., 1992).

The remainder of the article is structured as follows. In Section 2, we describe a general dynamic discrete choice model with hyperbolic discounting time preferences. In Section 3, we provide detailed analysis for identification. In Section 4, we propose an estimation strategy closely related to our identification arguments for the discount factors, and we also evaluate the performance of our proposed estimation method using Monte Carlo experiments. In Section 5, we provide the background information for mammography, which is the decision we examine in our empirical application; we also describe the data set used in our study and provide some basic descriptive statistics of the samples. We then provide details about the empirical specification of our model of the decision for undertaking mammography and present the main estimation results. Finally, Section 6 concludes and discusses a few important issues abstracted away in our analysis. Due to space limitations, we collect some omitted details, elaborations, proofs, and extensions in an online Appendix (Supporting Information).

2. DYNAMIC DISCRETE CHOICE MODEL WITH HYPERBOLIC DISCOUNTING TIME PREFERENCES

2.1. *Basic Model Setup.*[9] Consider a decision maker whose intertemporal utility is additively time separable. The agent's instantaneous preferences are defined over the action she chooses from a discrete set of alternatives $i \in \mathcal{I} = \{0, 1, \ldots, I\}$ and a list of state variables denoted by $h \equiv (x, \boldsymbol{\varepsilon})$, where $x \in \mathcal{X}$ are observed by the researcher, and $\boldsymbol{\varepsilon} \equiv (\varepsilon_1, \ldots, \varepsilon_I) \in R^I$ are the vector of random preference shocks for each of the I alternatives. We assume that $\mathcal{X}$ is a finite set and denote $X = |\mathcal{X}|$ to be the size of the state space. We make the following assumption about the instantaneous utility from taking action $i \in \mathcal{I}$, $u_i^*(h) \equiv u_i^*(x, \boldsymbol{\varepsilon})$.

ASSUMPTION 1 (ADDITIVE SEPARABILITY). *The instantaneous utilities are given by, for each* $i \in \mathcal{I} \setminus \{0\}$,

$$u_i^*(x, \boldsymbol{\varepsilon}) = u_i(x) + \varepsilon_i, \tag{1}$$

where $u_i(x)$ is the deterministic component of the utility from choosing i at x, and $(\varepsilon_1, \ldots, \varepsilon_I)$ has a joint distribution G, which is absolutely continuous with respect to the Lebesgue measure in R^I. Without loss of generality, we normalize $u_0(x) = 0$ for all $x \in \mathcal{X}$.

We assume that the *time horizon is infinite* with time denoted by $t = 1, 2, \ldots$[10] The decision maker's intertemporal preferences are represented by a simple and now commonly used formulation of potentially time-inconsistent preferences: (β, δ)-preferences (Phelps and Pollak, 1968; Laibson, 1997; O'Donoghue and Rabin, 1999a).

[8] See Section 5.5 and the last panel of Table 6 for more details about other exclusion variables used in our analysis.

[9] See online Appendix A (Supporting Information) for a summary of the key notations used in this article.

[10] For a discussion on how the analysis can be extended to the finite horizon case, see online Appendix H (Supporting Information).

DEFINITION 1. (β, δ)-Preferences are intertemporal preferences represented by

$$U_t(u_t, u_{t+1}, \ldots) \equiv u_t + \beta \sum_{k=t+1}^{+\infty} \delta^{k-t} u_k,$$

where $\beta \in (0, 1]$ and $\delta \in (0, 1]$.

Following the terminology of O'Donoghue and Rabin (1999a), the parameter δ is called the *standard discount factor*, which captures long-run, time-consistent discounting, and the parameter β is called the *present bias factor*, which captures short-term impatience. The standard model is nested as a special case of (β, δ)-preferences when $\beta = 1$. When $\beta \in (0, 1)$, (β, δ)-preferences capture "quasi-hyperbolic" time discounting (Laibson, 1997). We say that an agent's preferences are time consistent if $\beta = 1$ and are present biased if $\beta \in (0, 1)$.

The literature on time-inconsistent preferences distinguishes between *naive* and *sophisticated* agents (Strotz, 1956; Pollak, 1968; O'Donoghue and Rabin, 1999a, 1999b). An agent is *partially naive* if the self in every period t underestimates the present bias of her future selves, believing that her future selves' present bias is $\tilde{\beta} \in [\beta, 1]$; in the extreme, if the present self believes that her future selves are time consistent, i.e., $\tilde{\beta} = 1$, she is said to be *completely naive.* On the other hand, an agent is *sophisticated* if the self in every period t correctly knows her future selves' present bias β and anticipates their behavior when making her period-t decision, i.e., if $\tilde{\beta} = \beta$.

Following previous studies of time-inconsistent preferences, we will analyze the behavior of an agent by thinking of the individual as consisting of many autonomous *selves,* one for each period. Each period-t self chooses her current behavior to maximize her current utility $U_t(u_t, u_{t+1}, \ldots)$, while her future selves control their subsequent decisions.

More specifically, let the observable state variable in period t be $x_t \in \mathcal{X}$ and the unobservable choice-specific shocks be $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, \ldots, \varepsilon_{It}) \in R^I$. A (Markovian) *strategy profile* for all selves is $\boldsymbol{\sigma} \equiv \{\sigma_t\}_{t=1}^{\infty}$ where $\sigma_t : \mathcal{X} \times R^I \to \mathcal{I}$ for all t. It specifies for each self her action in all possible states and under all possible realizations of shock vectors. For any strategy profile $\boldsymbol{\sigma}$**,** write $\boldsymbol{\sigma}_t^+ \equiv \{\sigma_k\}_{k=t}^{\infty}$ as the *continuation strategy profile* from period t on.

To define and characterize the equilibrium of the intrapersonal game of an agent with potentially time-inconsistent preferences, we first write $V_t\left(x_t, \boldsymbol{\varepsilon}_t; \boldsymbol{\sigma}_t^+\right)$ as the agent's period-t expected continuation utility when the state variable is x_t and the shock vector is $\boldsymbol{\varepsilon}_t$ under her *long-run* time preference for a given continuation strategy profile $\boldsymbol{\sigma}_t^+$. We can think of $V_t\left(x_t, \boldsymbol{\varepsilon}_t; \boldsymbol{\sigma}_t^+\right)$ as representing, hypothetically, her intertemporal preferences from some prior perspective when her own present bias is irrelevant. Specifically, $V_t\left(x_t, \boldsymbol{\varepsilon}_t; \boldsymbol{\sigma}_t^+\right)$ must satisfy

$$V_t\left(x_t, \boldsymbol{\varepsilon}_t; \boldsymbol{\sigma}_t^+\right) = u^*_{\sigma_t(x_t, \boldsymbol{\varepsilon}_t)}\left(x_t, \varepsilon_{\sigma_t(x_t, \boldsymbol{\varepsilon}_t)t}\right) + \delta \mathrm{E}\left[V_{t+1}\left(x_{t+1}, \boldsymbol{\varepsilon}_{t+1}; \boldsymbol{\sigma}_{t+1}^+\right) | x_t, \sigma_t\left(x_t, \boldsymbol{\varepsilon}_t\right)\right], \tag{2}$$

where $\sigma_t\left(x_t, \boldsymbol{\varepsilon}_t\right) \in \mathcal{I}$ is the choice specified by strategy σ_t, and the expectation is taken over both the future state x_{t+1} and $\boldsymbol{\varepsilon}_{t+1}$.

We will define the equilibrium for a partially naive agent whose period-t self believes that, beginning next period, her future selves will behave optimally with a present bias factor of $\tilde{\beta} \in [\beta, 1]$. Following O'Donoghue and Rabin (1999b, 2001), we first define the concept of an agent's *perceived continuation strategy profile* by her future selves.

DEFINITION 2. The perceived continuation strategy profile for a partially naive agent is a strategy profile $\tilde{\boldsymbol{\sigma}} \equiv \{\tilde{\sigma}_t\}_{t=1}^{\infty}$ such that for all $t = 1, 2, \ldots$, all $x_t \in \mathcal{X}$, and all $\boldsymbol{\varepsilon}_t \in R^I$,

$$\tilde{\sigma}_t\left(x_t, \boldsymbol{\varepsilon}_t\right) = \arg\max_{i \in \mathcal{I}} \left\{u_i^*\left(x_t, \epsilon_{it}\right) + \tilde{\beta}\delta \mathrm{E}\left[V_{t+1}\left(x_{t+1}, \boldsymbol{\varepsilon}_{t+1}; \tilde{\boldsymbol{\sigma}}_{t+1}^+\right) | x_t, i\right]\right\}.$$

That is, if an agent is partially naive with perceived present bias by future selves at $\tilde{\beta}$, then her period-t self will anticipate that her future selves will follow strategies $\tilde{\boldsymbol{\sigma}}_{t+1}^{+} \equiv \{\tilde{\sigma}_k\}_{k=t+1}^{\infty}$. Note, importantly, that what the strategy profile $\tilde{\boldsymbol{\sigma}} \equiv \{\tilde{\sigma}_t\}_{t=1}^{\infty}$ describes is the *perception of the partially naive agent regarding what her future selves will play.* It is *not* what will generate the *actual play* that we observe in the data. What we actually observe is generated from the perception-perfect strategy profile that we now define.

DEFINITION 3. A perception-perfect strategy profile for a partially naive agent is a strategy profile $\boldsymbol{\sigma}^* \equiv \{\sigma_t^*\}_{t=1}^{\infty}$ such that, for all $t = 1, 2, \ldots$, all $x_t \in \mathcal{X}$, and all $\boldsymbol{\varepsilon}_t \in R^I$,

$$\sigma_t^*(x_t, \boldsymbol{\varepsilon}_t) = \arg\max_{i \in \mathcal{I}} \left\{ u_i^*(x_t, \varepsilon_{it}) + \beta\delta \mathrm{E}\left[V_{t+1}\left(x_{t+1}, \boldsymbol{\varepsilon}_{t+1}; \tilde{\boldsymbol{\sigma}}_{t+1}^{+}\right) | x_t, i \right] \right\}.$$

That is, $\boldsymbol{\sigma}^*$ is the best response of the current self with (β, δ)-preference against $\tilde{\boldsymbol{\sigma}}$, the perceived continuation strategy profile of her future selves. It is key to note the *difference* and *connection* between $\tilde{\boldsymbol{\sigma}}$ and $\boldsymbol{\sigma}^*$. $\tilde{\boldsymbol{\sigma}}$ is the unobserved *perception* of the partially naive agent regarding what her future selves will do, under the partial naivety assumption that her future selves do not suffer from the present bias as described by the parameter β, but instead is governed by present bias parameter $\tilde{\beta}$ that may differ from β. $\boldsymbol{\sigma}^*$ is what the self in each period will optimally choose to do, and the actions generated from $\boldsymbol{\sigma}^*$ are what will be observed in the data. Note also that when β and $\tilde{\beta}$ coincide, i.e., when the agent is sophisticated, we have $\boldsymbol{\sigma}^* = \tilde{\boldsymbol{\sigma}}$.

The existence of the perception-perfect strategy profile is shown in Peeters (2004) for the same class of stochastic games.[11] For the main text of the article, we make the following stationarity assumption on the data generating process.[12]

ASSUMPTION 2 (STATIONARITY). *We assume that the observed choices are generated under the stationary perception-perfect strategy profile of the infinite horizon dynamic game played among different selves of the decision makers.*

We also make two additional but standard assumptions about the transition of the state variables and the distribution of the shocks (see, e.g., Rust, 1994b).

ASSUMPTION 3 (CONDITIONAL INDEPENDENCE).

$$\pi(x_{t+1}, \boldsymbol{\varepsilon}_{t+1} | x_t, \varepsilon_t, d_t) = q(\boldsymbol{\varepsilon}_{t+1} | x_{t+1}) \pi(x_{t+1} | x_t, d_t)$$
$$q(\boldsymbol{\varepsilon}_{t+1} | x_{t+1}) = q(\boldsymbol{\varepsilon}).$$

ASSUMPTION 4 (EXTREME VALUE DISTRIBUTION). *$\boldsymbol{\varepsilon}_t$ is i.i.d extreme value distributed.*

It is well known that the distribution of the choice-specific shocks to payoffs in discrete choice models is not nonparametrically identified (see, e.g., Magnac and Thesmar, 2002). Thus, one has to make an assumption about the distribution of ε. The extreme value distribution assumption can in principle be replaced by any other *known* distribution G, at least in empirical applications. In what follows, we will discuss how the analysis will be modified if a known distribution G is assumed instead of the extreme-value distribution.[13] However, it should be noted that our formal identification arguments as summarized in Propositions 1 and 2 do invoke Assumption

[11] It is well known that this class of stochastic games may have multiple equilibria (see, e.g., Krusell and Smith, 2003).

[12] The importance of the stationarity assumption is discussed in Remark 6.

[13] Assumption 4 is invoked in three places in our subsequent analysis. See Remark 5 and the text surrounding (19).

4. In that sense, we would like to acknowledge that our notion of "semi-parametric" identification is weaker than what is typically implied in the literature.[14] However, it is important to emphasize that Assumption 4 is invoked in *all* of the structural empirical papers on hyperbolic discounting we referenced in the introduction, and it is also a maintained assumption (Assumption CLOGIT) in the dynamic discrete choice literature reviewed in Aguirregabiria and Mira (2010).

2.2. *Characterizing $\tilde{\boldsymbol{\sigma}}$ and $\boldsymbol{\sigma}^*$.* Now we describe the decision process of the decision maker. First, define the deterministic component of the *current choice-specific value function*, $W_i(x)$, as follows:

$$W_i(x) = u_i(x) + \beta\delta \sum_{x' \in \mathcal{X}} V(x')\pi(x'|x, i), \tag{3}$$

where $\pi(x'|x, i)$ denotes the transition probabilities for state variables x when action i is taken and $V(\cdot)$ is the *perceived long-run* value function defined as

$$V(x) \equiv \mathrm{E}_{\boldsymbol{\varepsilon}} V(x, \varepsilon; \tilde{\boldsymbol{\sigma}}), \tag{4}$$

where $V(x, \varepsilon; \tilde{\boldsymbol{\sigma}})$ is the *stationary* value function defined according to (2) under the perceived continuation strategy profile $\tilde{\boldsymbol{\sigma}}$ for a partially naive agent as defined in Definition 2.

Using $V(\cdot)$ as defined in (4), we can also define the *choice-specific value function of the next-period self as perceived by the current self*, $Z_i(x)$, as follows:

$$Z_i(x) = u_i(x) + \tilde{\beta}\delta \sum_{x' \in \mathcal{X}} V(x')\pi(x'|x, i). \tag{5}$$

There are two key difference between $W_i(x)$ and $Z_i(x)$. The first difference is in how they discount the future streams of payoffs: In $W_i(x)$ the payoff t-periods removed from the current period is discounted by $\beta\delta^t$, whereas in $Z_i(x)$ the payoff t-periods removed from now is discounted by $\tilde{\beta}\delta^t$. The second difference is interpretational: $W_i(x)$ represents how the current period self evaluates the deterministic component of the payoff from choosing alternative i, whereas $Z_i(x)$ is how the current-period self perceives how her next-period self would evaluate the deterministic component of the payoff from choosing alternative i. It is obvious but important to note that $W_i(x)$ will regulate the current self's optimal choice, but $Z_i(x)$ will regulate the perception of the current self regarding the choices of her future selves.

Given $Z_i(x)$, we know that the current self's perception of her future self's choice, i.e., $\tilde{\sigma}$ as defined in Definition 2, is simply given as

$$\begin{aligned} \tilde{\sigma}(x, \boldsymbol{\varepsilon}) &= \max_{i \in \mathcal{I}} \left[u_i(x) + \varepsilon_i + \tilde{\beta}\delta \sum_{x' \in \mathcal{X}} V(x')\pi(x'|x, i) \right] \\ &= \max_{i \in \mathcal{I}} \left[Z_i(x) + \varepsilon_i \right]. \end{aligned} \tag{6}$$

Let us define the probability of choosing alternative j by the next-period self as perceived by the current-period self when the next-period's state is x, $\tilde{P}_j(x)$, as

$$\begin{aligned} \tilde{P}_j(x) &= \Pr\left[\tilde{\sigma}(x, \boldsymbol{\varepsilon}) = j\right] \\ &= \Pr\left[Z_j(x) + \varepsilon_j \geq Z_{j'}(x) + \varepsilon_{j'} \text{ for all } j' \neq j\right]. \end{aligned} \tag{7}$$

[14] Although we conjecture that the identification results can be generalized to general known distribution G, the formal proof is beyond the scope of this article and left for future research.

With the characterization of $\tilde{\sigma}(x, \boldsymbol{\varepsilon})$, we can now provide a characterization of $V(\cdot)$. For this purpose, further denote the *perceived choice-specific long-run value function,* $V_i(x)$, as follows:

$$V_i(x) = u_i(x) + \delta \sum_{x' \in \mathcal{X}} V(x')\pi(x'|x, i). \tag{8}$$

According to the definition of $V(\cdot)$ as given by (4), $V(x)$ is simply the expected value of $V_i(x) + \varepsilon_i$, where i is the chosen alternative according to $\tilde{\sigma}(x, \boldsymbol{\varepsilon})$; that is, we have the following relationship:

$$V(x) = \mathrm{E}_{\boldsymbol{\varepsilon}} \left[V_{\tilde{\sigma}(x,\boldsymbol{\varepsilon})}(x) + \varepsilon_{\tilde{\sigma}(x,\boldsymbol{\varepsilon})} \right]. \tag{9}$$

Now note from (5) and (8), we have

$$V_i(x) = Z_i(x) + (1 - \tilde{\beta})\delta \sum_{x' \in \mathcal{X}} V(x')\pi(x'|x, i). \tag{10}$$

Relationship (10) is crucial as it allows us to rewrite (9) as[15]

$$V(x) = \mathrm{E}_{\boldsymbol{\varepsilon}} \max_{i \in \mathcal{I}} [Z_i(x) + \varepsilon_i] + (1 - \tilde{\beta})\delta \sum_{j \in \mathcal{I}} \left[\tilde{P}_j(x) \sum_{x' \in \mathcal{X}} V(x')\pi(x'|x, j) \right]. \tag{11}$$

With the above preliminaries, now we can describe how the agent will make the choices when the state variables are given by $(x, \boldsymbol{\varepsilon})$. From Definition 3 for perception-perfect strategy profile and Equation (3), we know that the current period decision maker will choose i if and only if

$$i \in \arg\max_{i \in \mathcal{I}} \{W_i(x) + \varepsilon_i\}.$$

That is, the perception-perfect strategy profile $\sigma^*(x, \boldsymbol{\varepsilon})$ is

$$\sigma^*(x, \boldsymbol{\varepsilon}) = \arg\max_{i \in \mathcal{I}} \{W_i(x) + \varepsilon_i\}.$$

Under Assumption 4, the probability of *observing* action i being chosen at a given state variable x is

$$P_i(x) = \Pr\left[W_i(x) + \varepsilon_i > \max_{j \in \mathcal{I} \setminus \{i\}} \{W_j(x) + \varepsilon_j\} \right] = \frac{\exp[W_i(x)]}{\sum_{j=0}^{I} \exp[W_j(x)]}. \tag{12}$$

$P_i(x)$ is the current-period self's equilibrium choice probabilities and will be observed in the data.

Now we derive some important relationships that will be used in our identification exercise below. First, note that by combining (3) and (5), we have that

$$Z_i(x) - u_i(x) = \frac{\tilde{\beta}}{\beta} [W_i(x) - u_i(x)]. \tag{13}$$

[15] See online Appendix B (Supporting Information) for details of the derivation of (11).

Since both $Z_i(\cdot)$ and $W_i(\cdot)$ depend on $V(\cdot)$, we would like to use (11) to derive a characterization of $V(\cdot)$. Note that under Assumption 4, we have (see Rust, 1994b)

$$\mathrm{E}_\varepsilon \max_{i\in\mathcal{I}}\{Z_i(x)+\varepsilon_i\} = \ln\left\{\sum_{i\in\mathcal{I}} \exp\left[Z_i(x)\right]\right\}. \tag{14}$$

Moreover, from (7), we have that

$$\tilde{P}_j(x) = \frac{\exp\left[Z_j(x)\right]}{\sum_{i=0}^{I}\exp\left[Z_i(x)\right]}. \tag{15}$$

Using (14) and (15), we can rewrite (11) as

$$V(x) = \ln\left\{\sum_{i\in\mathcal{I}} \exp\left[Z_i(x)\right]\right\} + (1-\tilde{\beta})\delta \sum_{j\in\mathcal{I}} \frac{\exp\left[Z_j(x)\right]}{\sum_{i=0}^{I}\exp\left[Z_i(x)\right]} \sum_{x'\in\mathcal{X}} V(x')\pi(x'|x,j). \tag{16}$$

The three sets of equations (5), (13), and (16) will form the basis of our identification argument below. Let us first make a few useful remarks.

REMARK 1. We have three value functions $\{W_i(x), Z_i(x), V_i(x) : x \in \mathcal{X}\}$ as defined, respectively, in (3), (5), and (8). Both $W_i(\cdot)$ and $Z_i(\cdot)$ are related to $V_i(\cdot)$. It is worth emphasizing that $W_i(x)$ will regulate the current self's choice behavior as demonstrated by (12), and $Z_i(x)$ will regulate the current self's perception of future selves' choices as demonstrated by (15). $V_i(x)$ is an auxiliary value function that simply uses the long-run discount factor δ to evaluate the payoffs from the choices that the current self perceives that will be made by her future selves.

REMARK 2. If $\tilde{\beta} = 1$, i.e., if the decision maker is completely naive, we can see from (10) that $V_i(x) = Z_i(x)$ for all x. This makes sense because when $\tilde{\beta} = 1$, the current self perceives her future selves to be time consistent. Thus, the current self is already perceiving her future selves to be behaving according to the long-run discount factor δ only.

REMARK 3. If $\tilde{\beta} = \beta$, i.e., when an agent is sophisticated, then Equations (3) and (5) tell us that $W_i(x) = Z_i(x)$. That is, if the decision maker is sophisticated, then the current self's own choice rule will be identical to what she perceives to be her future self's choice rule.

REMARK 4. When the decision maker is partially naive, there are two distinct value functions $W_i(x)$ and $Z_i(x)$ that separately regulate the choice of the current self and the perceived choice of her future selves. Equation (13) clarifies that it is the fact that we allow for potential naivety in the hyperbolic discounting model that is creating the wedge between $W_i(x)$ and $Z_i(x)$: If $\tilde{\beta} = \beta$, i.e., if agents are sophisticated (even when they suffer from present bias), it would be true that $W_i(x) = Z_i(x)$. This is an important point because, as we see in (12), the observed choice probabilities (i.e., data) would provide direct information about $W_i(x)$ without needing any information about the discount factors. Thus, when $\tilde{\beta} = \beta$, the observed choice probabilities also provide direct information about $Z_i(x)$, but when $\tilde{\beta}$ and β are potentially not equal, we can no longer learn about $Z_i(x)$ directly from the observed choice probabilities.

Finally, we would like to remark about the role played by Assumption 4 in the above derivation.

REMARK 5. Under a general distribution G for ε, we will lose the analytical expression (14) for $E_\varepsilon \max_{i\in\mathcal{I}}\{Z_i(x)+\varepsilon_i\}$, but nonetheless $\mathrm{E}_\varepsilon \max_{i\in\mathcal{I}}\{Z_i(x)+\varepsilon_i\}$ can be expressed as a function of

$\mathbf{z}(x) \equiv \{Z_i(x), i \in \mathcal{I}\}$, the form of which only depends on $G(\cdot)$. Let us denote this, as in Magnac and Thesmar (2002), by

$$R(\mathbf{z}(x); G) \equiv E_\varepsilon \max_{i \in \mathcal{I}} \{Z_i(x) + \varepsilon_i\}.$$

Similarly, $\tilde{P}_j(x)$ will not take the convenient logit form (15), but again it can be expressed as a function of $\mathbf{z}(x)$, the form of which only depends on $G(\cdot)$:

$$\tilde{P}_j(\mathbf{z}(x); G) \equiv \int_{\{Z_j(x)+\varepsilon_j \geq Z_{j'}(x)+\varepsilon_{j'} \text{ for all } j' \neq j\}} dG(\boldsymbol{\varepsilon}).$$

For a known distribution G, both $R(\mathbf{z}(x); G)$ and $\tilde{P}_j(\mathbf{z}(x); G)$ will be known functions that can be substituted in place of their corresponding expressions in (14) and (15). Also, both can be easily simulated in empirical applications.

2.3. *Relationship with the Dynamic Games Literature.* We analyze the observed outcome of the dynamic discrete choice problem of a hyperbolic discounting decision process as the equilibrium outcome of an *intrapersonal game* with the "players" being *the selves at different periods.* Thus, our article represents an interesting intermediate case between the classic literature on estimating single-agent dynamic discrete choice decision problems and the more recent literature on estimating dynamic games. It is worth pointing out that there are two crucial differences between the intrapersonal games we analyze for agents with time-inconsistent time preferences and those in the existing dynamic games literature.

The first key difference is that in our case, we *do not observe* the actions of all the "players." More specifically, the outcomes—choices and the evolutions of the state variables—we observe in the data are affected only by the current selves, even though the current selves' choices are impacted by their perception of future selves' actions. The current self's perception of how her future selves will play has to be inferred by the researcher using the equilibrium restriction imposed by the theory. As can be seen from the above discussion, $\tilde{P}_j(x)$, as defined in (15), captures the current self's perception of how her future selves will play, which is crucial for us to understand the current self's actual choices. However, as researchers, we do not observe $\tilde{P}_j$, but only observe P_j, the choice probabilities by the current self. In the standard dynamic games literature, it is always assumed that the actions of all the players are observed.

Second, in our intrapersonal game, the future selves know the state variables and the realizations of payoff shocks of the current and past selves. This differs from the existing incomplete-information dynamic games literature where players only know the distributions of opponents' payoff shocks. However, the fact that future selves have information about current selves does not matter in the analysis of the equilibrium because the future self, when called upon to play, would be best responding to the subsequent future selves, not the prior selves.

Third, the dynamic games literature (e.g., Bajari et al., 2007) may allow for players to have different contemporaneous payoff functions; in our setting, however, the payoffs for the players—the current self and the future selves—differ only in their time preferences. Moreover, under hyperbolic discounting, we are assuming a rather restricted form of difference in time preferences between the selves in different periods.

3. IDENTIFICATION

3.1. *Data and Preliminaries.* Before we describe our results on identification, let us assume that we have access to a data set that provides us with the following information.

DATA.

- Conditional choice probabilities: For all $x \in \mathcal{X}$, we observe the choice probabilities $P_i(x)$ for all $i \in \mathcal{I}$;
- Transitional probabilities for observable state variables: For all $(x, x') \in \mathcal{X}^2$, all $i \in \mathcal{I}$, we observe the transition probabilities $\pi(x'|x, i)$; we denote

$$\boldsymbol{\pi} \equiv \{\pi(x'|x, i) : (x, x') \in \mathcal{X}^2, i \in \mathcal{I}\};$$

- Short panels: We have access to at least two periods of the above data, even though the data result from a stationary infinite horizon model.

REMARK 6. The experimental literature cited in footnote 3 relies on observing preference reversal to infer about hyperbolic discounting. To observe preference reversals in these experiments requires that we observe the choices of the decision maker at a minimum of three points in time. Thus, it is important to explain how present biased preferences may be identified from panel data as short as two periods. First, note that preference reversals in the experimental literature rely on eliciting in an earlier date hypothetical choices that would be made in a later date. Our identification result applies to observational data, which does not have information about the hypothetical choices. Second, even though we assume that we have access to panel data as short as two periods, the stationarity assumption (Assumption 2) ensures that we are able to learn from the short-panel data about the choices of the decision maker under any state variables. The reason is that a decision maker's choice probabilities in a future period depend on the value of her state variables then. But in a stationary environment, we can estimate these choice probabilities from the individuals' choices in the short panel who have these same states.

Following Magnac and Thesmar (2002), we assume that the *structure* of the model, denoted by $\boldsymbol{\theta}$, is defined by the following parameters:[16]

$$\boldsymbol{\theta} = \{(\beta, \tilde{\beta}, \delta), G, \langle\{u_i(x), Z_i(x'), V_i(x') : i \in \mathcal{I}, x \in \mathcal{X}, x' \in \mathcal{X}\}\rangle\}, \tag{17}$$

where $G(\cdot)$, the distribution of the choice-specific payoff shocks ε_{it}, is assumed to have a Type-I extreme value distribution by Assumption 4. Note that the elements in $\boldsymbol{\theta}$ in our setting differ from those in Magnac and Thesmar (2002) in that we have two additional parameters β and $\tilde{\beta}$ that measure present bias and naivety, respectively; moreover, the interpretation of $V_i(x')$ in our article differs from theirs. In their paper, $V_i(x')$ directly informs about the *actual* choice probabilities of the decision maker in the second period because $\Pr(i|x') = \Pr(V_i(x') + \varepsilon_i > V_j(x') + \varepsilon_j \; \forall j \neq i)$. In our article, $Z_i(x')$ captures the current self's *perception* of the choice probability of the next period's self, which is *never actually* observed in the data, whereas $V_i(x')$ is just an auxiliary value function to account for the exponentially discounted payoff streams from the perceived choices made according to $\tilde{\sigma}(x, \boldsymbol{\varepsilon})$. Another difference is that in Magnac and Thesmar (2002), the vector $\{V_i(x') : x' \in \mathcal{X}\}$ are completely free parameters; in our setting, however, neither $Z_i(x')$ nor $V_i(x')$ are completely free parameters, as they are subject to the restriction that they have to satisfy (5), (13), and (16).

We denote by $\boldsymbol{\Theta}$ the set of all permissible structures. The set $\boldsymbol{\Theta}$ requires that the structure satisfies the assumptions we adopted in the model as well as the restrictions (5), (13), and (16).

[16] We could have included the vector $\{W_i(x): i \in \mathcal{I}, x \in \mathcal{X}\}$ as part of the model parameters as well. But as seen from Equation (19), $W_i(x)$ can be straightforwardly inferred from the data of choice probabilities $P_i(x)$, subject to a normalization. For this reason, we exclude them from our list of model parameters.

Given any structure $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, the model predicts the probability that an agent will choose alternative $i \in \mathcal{I}$ in state $x \in \mathcal{X}$, which we denote by $\hat{P}_i(x;\boldsymbol{\theta})$ and is given by

$$\hat{P}_i(x;\boldsymbol{\theta}) = \Pr\left\{u_i(x) + \varepsilon_i + \beta\delta \sum_{x' \in \mathcal{X}} V(x')\pi(x'|x,i) \right.$$

$$(18) \qquad \left. = \max_{j \in \mathcal{I}} \left[u_j(x) + \varepsilon_j + \beta\delta \sum_{x' \in \mathcal{X}} V(x')\pi(x'|x,j)\right] \middle| x, \boldsymbol{\theta} \right\}.$$

As is standard in the identification literature, we call the predicted choice probabilities $\hat{P}_i(x;\boldsymbol{\theta})$ as the *reduced form* of structure $\boldsymbol{\theta} \in \boldsymbol{\Theta}$. We say that two structures $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \boldsymbol{\Theta}$ are *observationally equivalent* if

$$\hat{P}_i(x;\boldsymbol{\theta}) = \hat{P}_i(x;\boldsymbol{\theta}')\, \forall i \in \mathcal{I} \text{ and } x \in \mathcal{X}.$$

A model is said to be *identified* if and only if for any $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \boldsymbol{\Theta}$, $\boldsymbol{\theta} = \boldsymbol{\theta}'$ if they are observationally equivalent.

3.2. *Identification Results.* We first describe, for a given $\langle \beta, \tilde{\beta}, \delta \rangle$, the identification of $\langle \{u_i(x), Z_i(x'), V_i(x') : i \in \mathcal{I}, x \in \mathcal{X}, x' \in \mathcal{X}\} \rangle$ based on the three set of equations (5), (13), and (16); then, we provide conditions pertinent to the identification of $\langle \beta, \tilde{\beta}, \delta \rangle$.

3.2.1. *Identification of* $\langle \{u_i(x), Z_i(x'), V_i(x') : i \in \mathcal{I}, x \in \mathcal{X}, x' \in \mathcal{X}\} \rangle$ *for given* $\langle \beta, \tilde{\beta}, \delta \rangle$. For any given joint distribution G of $\tilde{\varepsilon} \equiv (\varepsilon_1, \ldots, \varepsilon_I)$, the choice probability vector $\mathbf{P}(x) = (P_1(x), \ldots, P_I(x))$ is a mapping Q of $\mathbf{W}(x) = (W_0(x), W_1(x), \ldots, W_I(x))$. Hotz and Miller (1993) showed that the mapping Q can be inverted if for each $i \in \mathcal{I}$, one of the $W_i(x)$ is normalized. That is, one can find

$$D_i(x) \equiv W_i(x) - W_0(x) = Q_i(\mathbf{P}(x); G),$$

where Q_i is the ith component of the inverse Q. Under our Assumption 4 (that ε_i has i.i.d type-I extreme value distribution), the mapping Q_i is especially simple (following from (12)):

$$(19) \qquad D_i(x) = W_i(x) - W_0(x) = \ln \frac{P_i(x)}{P_0(x)}.$$

Since we observe $P_i(x)$ and $P_0(x)$ from the data, we immediately learn about $D_i(x)$. We thus proceed as if $D_i(x)$ is observable.

From Equation (13), we have that, for all $i \in \mathcal{I}$,

$$(20) \qquad Z_i(x) = \frac{\tilde{\beta}}{\beta} W_i(x) + \left(1 - \frac{\tilde{\beta}}{\beta}\right) u_i(x).$$

Together with (19), we have, for all $i \in \mathcal{I} \backslash \{0\}$ and $x \in \mathcal{X}$,

$$(21) \qquad Z_i(x) - Z_0(x) = \frac{\tilde{\beta}}{\beta} D_i(x) + \left(1 - \frac{\tilde{\beta}}{\beta}\right) [u_i(x) - u_0(x)].$$

Now let us examine the system of equations defined by Equation (16), which can be rewritten, using (15), as, for all $x \in \mathcal{X}$,

$$(22)\quad V(x) = Z_0(x) + \ln\left\{\sum_{i\in\mathcal{I}} \exp\left[Z_i(x) - Z_0(x)\right]\right\} + \left(1-\tilde{\beta}\right)\delta \sum_{j\in\mathcal{I}} \tilde{P}_j(x) \sum_{x'\in\mathcal{X}} V(x')\pi(x'|x,j).$$

Note that if we take $\{\beta, \tilde{\beta}, \delta, \langle Z_i(x) : x \in \mathcal{X}\rangle\}$ as given, Equation (22) is just a system of X linear equation in X unknowns, namely, $\mathbf{V} \equiv [V(1), \ldots, V(x)]^T$.[17]

Specifically, if we stack $[Z_0(x) + \ln\{\sum_{i\in\mathcal{I}} \exp[Z_i(x) - Z_0(x)]\}]_{x\in\mathcal{X}}$ into an $X \times 1$ column vector $\mathbf{A}$, denote the $X \times [(I+1)X]$ matrix of choice probabilities by $\tilde{\mathbf{P}} \equiv [\tilde{\mathbf{P}}_0 \cdots \tilde{\mathbf{P}}_I]$,[18] and properly stack up the transition matrices $\pi(x'|x,j)$ into an $[(I+1)X] \times X$ matrix as:[19]

$$\mathbf{\Pi} \equiv \begin{bmatrix} \mathbf{\Pi}_0 \\ \vdots \\ \mathbf{\Pi}_I \end{bmatrix},$$

we can write (22) as

$$\mathbf{V} = \mathbf{A} + (1-\tilde{\beta})\delta\tilde{\mathbf{P}}\mathbf{\Pi}\mathbf{V},$$

which yields the solution of $\mathbf{V}$ as a function of $\{\beta, \tilde{\beta}, \delta, \langle Z_i(x) : x \in \mathcal{X}\rangle\}$ to be

$$(23)\qquad \mathbf{V} = [\mathbf{I} - (1-\tilde{\beta})\delta\tilde{\mathbf{P}}\mathbf{\Pi}]^{-1}\mathbf{A}.$$

Thus, for fixed values of $\langle\beta, \tilde{\beta}, \delta\rangle$, we can plug (23) into (5) and obtain, for all $x \in \mathcal{X}$ and all $i \in \mathcal{I}$,

$$(24)\qquad Z_i(x) = u_i(x) + \tilde{\beta}\delta\mathbf{\Pi}_i(x)[\mathbf{I} - (1-\tilde{\beta})\delta\tilde{\mathbf{P}}\mathbf{\Pi}]^{-1}\mathbf{A},$$

where $\mathbf{\Pi}_i(x) = [\pi(1|x,i), \ldots, \pi(X|x,i)]$ is an $X \times 1$ vector as defined in footnote 19.

Now consider the system of equations given by (21) and (24). We know from the standard theories of discrete choice that we have to normalize the utility for the reference alternative 0, for which without loss of generality we set $u_0(x) = 0$ for all $x \in \mathcal{X}$.[20] Thus, the unknowns contained in the equation system of (21) and (24) include the $(I+1) \times X$ values for $\{Z_i(x) : i \in \mathcal{I}, x \in \mathcal{X}\}$ and $I \times X$ values for $\{u_i(x) : i \in \mathcal{I}/\{0\}, x \in \mathcal{X}\}$; that is, the total number of unknowns is $(2I+1) \times X$. It is also easy to see that the total number of equations in the system is also equal to $(2I+1) \times X$:

[17] Note that all the terms in $\tilde{P}_j(x)$ depend just on $Z_i(\cdot)$.

[18] For each $j \in \mathcal{I}$, $\tilde{\mathbf{P}}_j$ is the $X \times X$ matrix organized as

$$\tilde{\mathbf{P}}_j = \begin{bmatrix} \tilde{P}_j(0) & 0 & \cdots & 0 \\ 0 & \tilde{P}_j(1) & \cdots & 0 \\ \mathbf{0} & \mathbf{0} & \ddots & \mathbf{0} \\ 0 & 0 & \cdots & \tilde{P}_j(x) \end{bmatrix}.$$

[19] Here, $\mathbf{\Pi}_j$ is an $X \times X$ matrix defined by

$$\mathbf{\Pi}_j = \begin{bmatrix} \mathbf{\Pi}_j(1) \\ \mathbf{\Pi}_j(2) \\ \vdots \\ \mathbf{\Pi}_j(x) \end{bmatrix} = \begin{bmatrix} \pi(1|1,j) & \ldots & \pi(X|1,j) \\ \pi(1|2,j) & \ldots & \pi(X|2,j) \\ & \vdots & \\ \pi(1|X,j) & \ldots & \pi(X|X,j) \end{bmatrix}.$$

[20] See Norets and Tang (2014) for interesting discussions on the effects of such normalizations.

$I \times X$ equations in (21) and $(I+1) \times X$ equations in (24). Because the main variables of interest that we would like to solve in the equation system (21) and (24) are the $I \times X$ values for $\{u_i(x) : i \in \mathcal{I}/\{0\}, x \in \mathcal{X}\}$, and the values for $\{Z_i(x)\}$ are only auxiliary, we can further eliminate $\{Z_i(x)\}$ from the equation system (21) and (24) and only solve for $\mathbf{u} \equiv \{u_i(x) : i \in \mathcal{I}/\{0\}, x \in \mathcal{X}\}$, where $\mathbf{u}$ is used as the short hand for the $I \times X$-vector; below we will write

$$\mathcal{G}(\mathbf{u};(\beta, \tilde{\beta}, \delta)) = 0 \tag{25}$$

to denote this system of $I \times X$ equations in $I \times X$ unknowns, namely, $\{u_i(x) : i \in \mathcal{I}/\{0\}, x \in \mathcal{X}\}$.

Of course, having the same number of equations as the number of unknowns does *not* guarantee the existence or the uniqueness of the solutions. However, suppose that the data are generated by primitive $\boldsymbol{\theta}^*$ where the vector of objects $\boldsymbol{\theta}$ is described in (17), that is, $\mathcal{G}(\mathbf{u}^*; (\beta^*, \tilde{\beta}^*, \delta^*)) = 0$.[21] It can be shown that the Jacobian matrix of $\mathcal{G}(\mathbf{u}; (\beta^*, \tilde{\beta}^*, \delta^*))$ with respect to $\mathbf{u}$ evaluated at $\mathbf{u}^*$, which we denote by $\partial_{\mathbf{u}}\mathcal{G}(\mathbf{u}^*;(\beta^*, \tilde{\beta}^*, \delta^*))$, has full rank. Thus, by implicit function theorem, we have the following local uniqueness result.[22]

Proposition 1. *Suppose that data are generated by the assumed data generating process with primitives* $(\mathbf{u}^*, \beta^*, \tilde{\beta}^*, \delta^*)$. *Then, the solution to the system of nonlinear equations* (25) *is locally unique when* $(\beta, \tilde{\beta}, \delta) = (\beta^*, \tilde{\beta}^*, \delta^*)$.

Note that Proposition 1 only says that the solution to equation system (25) is locally unique when the discount factors are at their true values. When the discount factors are not at the true values, it is still possible that the nonlinear equation system (25) has no solution or multiple solutions. Nonetheless, under additional exclusion restrictions described in the next subsection, Proposition 2 below will establish that the true values of the discount factors $(\beta^*, \tilde{\beta}^*, \delta^*)$ will be generically identified.[23]

We should also note that in the equation system of (21) and (24), the three discounting parameters $(\beta, \tilde{\beta}, \delta)$ appear in three different combinations: First, in Equation (24), both $\tilde{\beta}\delta$ and $(1-\tilde{\beta})\delta$ appear; second, $\tilde{\beta}/\beta$ appears in Equation (21). The fact that $\tilde{\beta}\delta$, $(1-\tilde{\beta})\delta$ and $\tilde{\beta}/\beta$ appear in the equation system is crucial for the separate estimation of the $(\beta, \tilde{\beta}, \delta)$ parameters.

3.2.2. *Exclusion restriction to identify* $\langle\beta, \tilde{\beta}, \delta\rangle$. Now we discuss conditions for identification related to $\langle\beta, \tilde{\beta}, \delta\rangle$. This discussion is closely related to that in Magnac and Thesmar (2002). We impose the following *exclusion restriction assumption.*

Assumption 5 (Exclusion Restriction). *There exist state variables* $x_1 \in \mathcal{X}$ *and* $x_2 \in \mathcal{X}$ *with* $x_1 \neq x_2$, *such that*

(1) *for* all $i \in \mathcal{I}$, $u_i(x_1) = u_i(x_2)$;
(2) *for some* $i \in \mathcal{I}$, $\pi(x'|x_1, i) \neq \pi(x'|x_2, i)$.

More specifically, to satisfy the exclusion restriction assumption, there must exist at least one variable that does not directly affect the *contemporaneous utility function* $u_i(\cdot)$ for all $i \in \mathcal{I}$, but the variable may matter for choices because it affects the transition of state variables. The

[21] Even though the primitives of the model listed in (17) also include $\{Z_i(x), V_i(x)\}$, they are both functions of $(\beta, \tilde{\beta}, \delta)$ and $\mathbf{u}$.

[22] The implicit function theorem also implies that in the neighborhood around $(\beta^*, \tilde{\beta}^*, \delta^*)$, the solution to the equation system (25) is continuous and differentiable with respect to $(\beta, \tilde{\beta}, \delta)$. This result is used in the proof of Proposition 3 in online Appendix F (Supporting Information).

[23] In the estimation, we may need to worry about the number of solutions to equation system (25) when the values of the discount factors are not at the truth. We will describe in Section 4 how the number of solutions affects our estimation procedure.

extent to which individuals' choice probabilities differ at states x_1 and x_2 reveals information about the discount factors. This is the key intuition from Magnac and Thesmar's (2002) result where they are interested in identifying a single long-term discount factor δ. In their setting, if $\delta = 0$, i.e., if individuals are completely myopic, then the choice probabilities would have been the same under x_1 and x_2; to the extent that choice probabilities differ at x_1 and x_2, it reveals information about the degree of time discounting. Their intuition, however, can be extended to the hyperbolic discounting case, as we will exploit in the proposed estimation strategy below.

For notational simplicity, we will divide the state variables into two groups (x_r, x_e) where $x_r \in \mathcal{X}_r$ refers to the state variables that directly enter the instantaneous payoff function $u_i(x_r)$ and $x_e \in \mathcal{X}_e$ refers to the state variables that satisfy the exclusion restriction (5). The key idea that the existence of exclusion variables can provide the source of identification of the discount factors is as follows: For any given values of $\langle \beta, \tilde{\beta}, \delta \rangle$, Section 3.2.1 tells us that we can identify, from the observed data, values of $\langle \{u_i(x) : i \in \mathcal{I} \setminus \{0\}, x \in \mathcal{X}\} \rangle$, to the extent that $\langle \beta, \tilde{\beta}, \delta \rangle$ might be consistent with the data at all. However, consider two state vectors $x = (x_r, x_e)$ and $x' = (x_r, x_e')$ that only differ in the exclusion restriction components. If the postulated values of $\langle \beta, \tilde{\beta}, \delta \rangle$ are inconsistent with the true values, there is no guarantee that the identified values of $u_i(x)$ and $u_i(x')$ are equal, as required by the exclusion restriction. The exclusion restriction will allow us to select the values of $\langle \beta, \tilde{\beta}, \delta \rangle$ such that the values of $u_i(x)$ and $u_i(x')$ identified from the data only differ if x and x' differ in components of x_r. Note that for x_e to be a state *variable* that satisfies the exclusion restriction, it must be the case that $|\mathcal{X}_e| \geq 2$ (otherwise, x_e is just a constant).

PROPOSITION 2. *Consider the space of data sets that can be generated by the assumed data generating process for some primitives* $(\mathbf{u}^*, \beta^*, \tilde{\beta}^*, \delta^*)$. *Suppose that there exist state variables that satisfy Assumption 5. Then, all the model parameters are generically identified if* $I \times |\mathcal{X}_e| \times |\mathcal{X}_r| \geq 4$.

The formal proof of Proposition 2 involves verifying that the conditions for the transversality theorem (see Proposition 8.3.1 in Mas-Colell, 1985) are satisfied. The details are available in the online Appendix C (Supporting Information). Note that, as long as there exist state variables x_e that satisfy the exclusion restriction assumption, the condition that $I \times |\mathcal{X}_e| \times |\mathcal{X}_r| \geq 4$ is always satisfied because $|\mathcal{X}_e|$ and $|\mathcal{X}_r|$ both are at least two for any nontrivial model. Heuristically, Proposition 2 states that for *almost all* data sets generated by the assumed hyperbolic discounting model, the observed choice probabilities $\tilde{\mathbf{P}}$ and transition probabilities $\mathbf{\Pi}$ will reveal the true discount factors as long as the model has state variables that satisfy the exclusion restrictions as described by Assumption 5.[24]

REMARK 7. We should emphasize that, since we normalize $u_0(x) = 0$ for all $x \in \mathcal{X}$, all variables x_e to affect u_i in the same way as they affect u_0 would ensure that $u_i(x_r, x_e) - u_0(x_r, x_e)$ only depends on x_r, and thus satisfy condition 1 in Assumption 5. This observation is used in our empirical application in our choice of exclusion variables (see Section 5.5 for details).

REMARK 8. There is one crucial difference between the role played by the exclusive variables in our setting and that of instrument variables (IV) in regression models with endogenous regressors (see Chapters 5 and 6, Wooldridge, 2002). In the IV setting, if there are more available instrumental variables than the number of endogenous regressors, the researcher can use different subsets of the IVs, which will all yield consistent parameter estimates. Different subsets of the IVs give the researcher different sources of variations in the endogenous regressors;

[24] See online Appendix D (Supporting Information) for some intuition as to how $\langle \beta, \tilde{\beta}, \delta \rangle$ come to differentially affect the observed choice behavior of the current self depending on the values of the exclusion variables. Online Appendix E (Supporting Information) shows that the existing results in the literature on the exponential discounting case can be replicated as a special case of our setup.

however, the underlying model being estimated does not depend on the subset of included IVs. In our setting, however, the exclusion variables are an explicit part of the model—they are state variables that satisfy Assumption 5; thus, when we use a different subset of exclusion variables, we are in fact postulating different models. As a result, the parameter estimates when we use different exclusion variables are those of different models and are therefore not directly comparable.

REMARK 9. When $I \times |\mathcal{X}_e| \times |\mathcal{X}_r| > 4$, the model is overidentified, and overidentification tests can be conducted. The idea is very simple. Given $I \times |\mathcal{X}_e| \times |\mathcal{X}_r| > 4$, we can identify the model using only a subset of the additional equations of the form $u_i(x_r, x_e) = u_i(x_r)$ for each $i \in \mathcal{I} \backslash \{0\}$, each $x_e \in \mathcal{X}_e$, and each $x_r \in \mathcal{X}_r$. There are many ways the over identification test can be conducted. As an illustration, consider the simplest case in which $|\mathcal{X}_e| > 4$; then suppose that we use a subset $\mathcal{X}_e' \subset \mathcal{X}_e$ with $|\mathcal{X}_e'| \geq 4$ in our estimation of the model parameters, that is, we only restrict that the payoff function $u_1(x)$ does not vary with respect to elements in $\mathcal{X}_e'$, but allow $u_1(x)$ to potentially vary with elements in $\mathcal{X}_e'' \equiv \mathcal{X}_e \backslash \mathcal{X}_e'$. If the model is correctly specified, then we should find that the estimated values of $u_1(x)$ should not vary with elements in $\mathcal{X}_e''$ either. This can be used as the basis of the overidentification test (see footnote 26 for more details).

4. AN ESTIMATION STRATEGY

In this section, we describe a two-step estimation strategy based on our discussion about identification in the previous section. In the first step, we estimate from the data the choice probabilities $P_i(x)$ for all $i \in \mathcal{I}$ and all $x \in \mathcal{X}$ as well as the state transition probabilities $\pi(x'|x, i)$ for all $i \in \mathcal{I}$ and all $(x', x) \in \mathcal{X}^2$.[25] In the second step, we maximize a pseudolikelihood function described below to estimate the discount factors $\langle \beta, \tilde{\beta}, \delta \rangle$. We first solve equation system (21) and (24) for $u_i(x) = u_i(x_r, x_e)$ for a given triple of values for $\langle \beta, \tilde{\beta}, \delta \rangle$ for all $i \in \mathcal{I} \backslash \{0\}$ and all $x \in \mathcal{X}$. Recall that condition 1 in Assumption 5 requires that $u_i(x_r, x_e)$ should not depend on x_e under the true model. We directly impose this restriction in our estimation:

$$\hat{u}_i(x_r) = \frac{1}{\left|\{(x_r, \tilde{x}_e) : \tilde{x}_e \in \mathcal{X}_e\}\right|} \sum_{\{(x_r, \tilde{x}_e) : \tilde{x}_e \in \mathcal{X}_e\}} \hat{u}_i(x_r, \tilde{x}_e), \tag{26}$$

where we recall that $\mathcal{X}_e$ is the set of possible values for the payoff-irrelevant state variables x_e we discussed in Assumption 5.[26] Note that, if our model is correctly specified and $\langle \beta, \tilde{\beta}, \delta \rangle$ are at the true values $\langle \beta^*, \tilde{\beta}^*, \delta^* \rangle$, $\hat{u}_i(x_r, x_e)$ will indeed be independent of x_e and thus (26) would not have been a restriction.

Given $\hat{u}_i(x_r)$ as defined by (26), we can then use the model to predict the choice probabilities $\hat{P}_i(x; \langle \beta, \tilde{\beta}, \delta \rangle)$ as described by (18) and formulate the pseudolikelihood of the observed data as[27]

$$\mathcal{L}(\text{data}; \boldsymbol{\psi} \equiv (\beta, \tilde{\beta}, \delta)) = \Pi_{j=1}^{n} \Pi_{i=1}^{I} \Pi_{x \in \mathcal{X}} \hat{P}_i(x; \boldsymbol{\psi})^{D_{i,j}(x)}, \tag{27}$$

[25] It is useful to note that, so far, our discussion has focused on short-panel (two period) data sets under stationarity assumption. Having two-period data allows one to nonparametrically estimate the transition probabilities $\pi(x'|x, i)$; stationarity ensures that looking at a two-period slice of a potentially long panel is sufficient. Fang and Silverman (2006) considered a case without stationarity (specifically, a finite horizon model) and showed that β and δ could be potentially distinguished without exclusion restriction if the researcher has access to at least three-period panel data. Also, see Section 6 for discussions related to nonstationarity and longer panels.

[26] If $|\mathcal{X}_e| > 4$ and if we would like to implement the overidentification test as described in Remark 9, we should replace the set $\mathcal{X}_e$ in Equation (26) by $\mathcal{X}_e'$, the subset of $\mathcal{X}_e$ that is used in identifying the discount factors, and test whether the estimated $\hat{u}_i(x_r, x_e'')$ depends on x_e'' where $x_e'' \in \mathcal{X}_e'' \equiv \mathcal{X}_e \backslash \mathcal{X}_e'$.

[27] Similar maximum pseudolikelihood estimators have been proposed by Aguirregabiria and Mira (2002) and Aguirregabia (2004).

where j stands for an individual and n is the sample size, and $D_{i,j}(x)$ is an indicator with value 1 if individual j chooses alternative i when state variable is x and 0 otherwise. Also note that we partial out the contribution from the state transitions $\pi(x'|x, i)$ in the likelihood since it is already estimated in step one. We then maximize the pseudolikelihood function to estimate $\boldsymbol{\psi} \equiv (\beta, \tilde{\beta}, \delta)$. Denote

$$\hat{\boldsymbol{\psi}} = \arg\max_{\boldsymbol{\psi}} \mathcal{L}(\text{data}; \boldsymbol{\psi}). \tag{28}$$

The following proposition, proved in online Appendix F (Supporting Information), establishes the consistency and asymptotically normality of our proposed estimator.

PROPOSITION 3. *Suppose that $\boldsymbol{\psi}^* \equiv (\beta^*, \tilde{\beta}^*, \delta^*)$ is the true parameter for the data-generating process. Under the stated identifying assumptions, $\hat{\boldsymbol{\psi}}$ is a $\sqrt{n}$-consistent estimator of $\boldsymbol{\psi}^*$. Moreover, if $\boldsymbol{\psi}^*$ lies in the interior of the parameter space, then*

$$\sqrt{n}\,(\hat{\boldsymbol{\psi}} - \boldsymbol{\psi}^*) \xrightarrow{d} \mathcal{N}(0, H^{-1}),$$

where H is the Fisher information matrix defined in online Appendix F (Supporting Information).

REMARK 10. Note that Proposition 3 applies only when the true $\boldsymbol{\psi}$ is in the interior of the parameter space. However, in our empirical application, the partial naivety parameter $\tilde{\beta}^*$ is estimated to be close to the boundary. We correct the standard error calculation when the parameter lies on the boundary using methods suggested by Moran (1971) and Andrews (1999, 2001). The details of the standard error correction are provided in online Appendix F (Supporting Information).

It is also worth discussing several practical issues that may arise in the implementation of our estimator. First, even though the equation system (25) has the same number of equations and unknowns, there is no guarantee that the system will have a unique solution for *any* arbitrary values of $\langle \beta, \tilde{\beta}, \delta \rangle$. In practice, we deal with the potential no solution or multiple solution problem as follows. If for a particular combination of $\langle \beta, \tilde{\beta}, \delta \rangle$, we are unable to find a solution to the equation system (21) and (24), we will simply set the likelihood in (27) to $-\infty$ (i.e., a sufficiently small value). On the other hand, if for some particular combination of $\langle \beta, \tilde{\beta}, \delta \rangle$, we find multiple solutions to the equation system, we will pick the highest likelihood in (27). Note that these procedures do not affect the asymptotic properties of our maximum pseudolikelihood estimator because Proposition 2 ensures that the parameters are generically identified under the stated assumptions.

Finally, even though we described how $\langle \{u_i(x) : i \in \mathcal{I} \setminus \{0\}, x \in \mathcal{X}\} \rangle$ can be nonparametrically solved from equation system (25) for given values of $\langle \beta, \tilde{\beta}, \delta \rangle$, in practice, we may often parameterize it as a function of observed state variables x, including the exclusion variables x_e. Our estimation method, however, goes through unmodified (see our Monte Carlo experiments below for an illustration).

4.1. *Monte Carlo Experiments.* In this section, we provide Monte Carlo evidence for the identification of discount factors in a dynamic discrete choice model using the maximum pseudolikelihood estimation method. In this simple Monte Carlo exercise, we consider a binary choice decision problem, $i \in \{0, 1\}$, facing an agent with infinite horizon, stationary state transition, and linear utility functions. There are two state variables x_r and x_e. The state variable $x_r \in \mathcal{X}_r = \{0, 1, 2, 3, 4, 5\}$ affects both instantaneous utility and state transition, whereas state

TABLE 1
MONTE CARLO RESULTS

Parameters	δ	β	$\tilde{\beta}$	α_0	α_1
True Values	0.8	0.6	0.7	−0.1	0.5
Sample size: 9, 600					
Mean	0.8059	0.6102	0.7035	−0.1012	0.5059
Std. dev.	0.0171	0.019	0.0168	0.0039	0.009
Sample size: 7, 200					
Mean	0.8062	0.6095	0.7028	−0.1011	0.5059
Std. dev.	0.0185	0.016	0.0159	0.0028	0.01
Sample size: 4, 800					
Mean	0.8048	0.6082	0.703	−0.1013	0.5062
Std. dev.	0.0209	0.0134	0.0168	0.003	0.0112

NOTES: For each sample size, we generate 1,000 random simulation samples. The mean and standard deviations of the estimated parameters are with respect to the 1,000 samples.

variable $x_e \in \mathcal{X}_e = \{0, 1, 2, 3\}$ affects only the state transition.[28] We parameterize the instantaneous payoff functions $u_i(x_r)$ as follows: $u_1(x_r) = \alpha_0 + \alpha_1 x_r$; and normalize $u_0(x_r) = 0$ for all x_r. The true parameters are set at $\alpha_0 = -0.1$, $\alpha_1 = 0.5$, $\delta = 0.8$, $\beta = 0.6$, and $\tilde{\beta} = 0.7$. For each sample size 9,600, 7,200, and 4,800, we randomly generate 1,000 simulation samples and estimate the discount factors and corresponding utility parameters to obtain their respective means and standard deviations.

Table 1 presents the Monte Carlo results. It shows that our proposed estimation method does an excellent job in precisely recovering the true parameter values in sample sizes similar to what one may have in actual empirical applications.[29] It also shows that estimation results improve with increases in sample sizes.

5. AN EMPIRICAL APPLICATION: MAMMOGRAPHY DECISIONS

In this section, we present an illustrative application of the identification and estimation method for $\langle \beta, \tilde{\beta}, \delta \rangle$ described in Sections 3.2 and 4. We provide estimates of these key time preference parameters and then examine the role of present bias and naivety in women's decisions to undertake mammography. We should emphasize that the main purpose of this empirical application is to illustrate the workings of the proposed identification strategy to estimate discount factors. To the best of our knowledge, this is the first paper that uses exclusive restrictions, first proposed in Magnac and Thesmar (2002), as the source of identification to estimate discount factors. It should also be pointed out that even though we make several simplifications in our empirical application, the obtained estimates nonetheless are plausible.

[28] Specifically, the state transition matrices for the Monte Carlo experiments are generated as follows. For simplicity, we let

$$\pi\left(x_r', x_e' | x_r, x_e, i\right) = \pi\left(x_e' | x_e\right) \pi\left(x_r' | x_r, x_e, i\right).$$

The $|\mathcal{X}_e| \times |\mathcal{X}_e| = 4 \times 4$ matrix Π_e summarizing $\pi(x_e'|x_e)$ has diagonal given by (0.6, 0.7, 0.8, 0.9) and the remaining elements randomly generated from a uniform [0, 1] and then properly normalized. (The diagonal of Π_e is chosen to ensure that x_e are reasonably persistent, a feature we encounter in the data.) To generate the two matrices representing $\pi(x_r'|x_r, x_e, i = 0)$ and $\pi(x_r'|x_r, x_e, i = 1)$, we first generate two random matrices M_0 and M_1 each with dimension $(|\mathcal{X}_r| \times |\mathcal{X}_e|) \times |\mathcal{X}_r| = 24 \times 6$, with each entry a random number generated from a uniform [0, 1] distribution. We then normalize the entry in each row by its row sum to ensure a proper probability matrix. The resulting matrices are denoted by $\hat{\Pi}_0$ and $\hat{\Pi}_1$. The choice-specific state transition probabilities, denoted as Π_0 and Π_1, are then constructed from Π_e, $\hat{\Pi}_0$, and $\hat{\Pi}_1$ appropriately. The matrices Π_1 and Π_0 are assumed to be known by the decision maker and are directly taken to be the state transition probabilities in our Monte Carlo exercise reported in Table 1.

[29] The sample size in our empirical application reported in Section 5 is about 11,000.

5.1. *Background on Mammography.* Among American women, breast cancer is the third most common cause of death and the second leading cause of cancer death. According to the American Cancer Society, from birth to age 39, one woman in 231 will get breast cancer (< 0.5% risk); from age 40−59, the chance is 1 in 25 (4% risk); from age 60−79, the chance is 1 in 15 (nearly 7%). Assuming that a woman lives to age 90, the chance of getting breast cancer over the course of an entire lifetime is one in seven, with an overall lifetime risk of 14.3%.

Breast cancer takes years to develop. Early in the disease, most breast cancers cause no symptoms. When breast cancer is detected at a localized stage before it spreads to the lymph nodes, the 5-year survival rate is 98%. If the cancer has spread to nearby lymph nodes (regional disease), the rate drops to 81%. If the cancer has spread (metastasized) to distant organs such as the lungs, bone marrow, liver, or brain, the 5-year survival rate is 26%.

A screening—mammography—is the best tool available to find breast cancer before symptoms appear. Mammography can often detect a breast lump before it can be felt and therefore save lives by finding breast cancer as early as possible. For women over the age of 50, mammography has been shown to lower the chance of dying from breast cancer by 35% (American Cancer Society). Leading experts, the National Cancer Institute, the American Cancer Society, and the American College of Radiology recommend annual mammography for women over 40.[30] The guideline issued by the U.S. Preventive Services Task Force in 2002 also recommended mammography screening for women beginning at age 40 every 12–24 months to reduce the risk of death from breast cancer.[31]

5.2. *Data.* The data used in this analysis are from the Health and Retirement Study (HRS).[32] The HRS is a nationally representative biennial panel study of birth cohorts 1931 through 1941 and their spouses as of 1992. The initial sample includes 12,652 persons in about 7,600 households who have been interviewed every two years since 1992. The most recent available data are for year 2010 (wave 10). The survey history and design are described in more details in Juster and Suzman (1995). Because the HRS only started asking women questions about their usage of mammography in 1996, our sample is limited to women interviewed from 1996 onwards.

We focus on the age group 51–64 and exclude those observations with missing values for any of the critical variables.[33] We also exclude those who have ever been diagnosed of (breast) cancer,[34] since those who are diagnosed of cancer might not make decisions on mammography or any other preventive health care the same way as do others.

For estimation purposes, to be included in our sample, a woman must be interviewed in (at least) two consecutive waves t and $t+1$ and must have reported her mammography decision in wave t, though not necessarily her mammography decision in wave $t+1$. Such a woman can be considered an observation in our two-period short panel. We look for such qualified observations from each wave of the HRS from 1996 and pool all the two-period short panels together to get our sample. Our final sample consists of 11,450 observations for 6,493 individuals.

[30] See American Cancer Society Web site: http://www.cancer.org/Cancer/BreastCancer/DetailedGuide. Women age 40 and older should have a screening mammogram every year and should continue to do so for as long as they are in good health.

[31] The new guideline by the U.S. Preventive Service Task Force issued in November of 2009 changed its recommendation for women in their 40s. Now the guideline recommends that women in their 40s should not get routine mammograms for early detection of breast cancer. For women ages 50–74, it recommends routine mammography screenings every two years. The American Cancer Society maintains its original recommendation that women from their 40s should seek annual mammography.

[32] We use the Version K of the HRS data released by Rand Corporation.

[33] The key variables are education status, non-Hispanic white, self-reported health, insurance status, whether father and mother are alive or died after age 70, father and mother's education, income, and mammography usage (see Table 2 for their summary statistics).

[34] For those who entered in 1992 (i.e., the first wave), the HRS provides information on breast cancer diagnoses as of 1992. But for those who entered HRS in later waves, we only know whether there is *any* cancer diagnosis.

TABLE 2
SUMMARY STATISTICS OF KEY VARIABLES IN THE ESTIMATION SAMPLE

Variable	Mean	Std. Dev.	Minimum	Maximum	Obs.
Mammogram	0.763	0.426	0	1	11,447
Bad health	0.217	0.412	0	1	11,447
Married	0.714	0.452	0	1	11,447
White (non-Hispanic)	0.798	0.401	0	1	11,447
High school or higher	0.796	0.402	0	1	11,447
Age	57.82	3.95	51	64	11,447
Death	0.014	0.117	0	1	11,447
Insurance	0.721	0.449	0	1	11,447
Household income ($1,000)	50.95	67.29	0.101	2136	11,447
Log of household income	10.354	1.053	4.615	14.575	11,447
Mother still alive or died after age 70	0.768	0.422	0	1	11,447
Mother education (high school or higher)	0.431	0.495	0	1	11,447
Father still alive or died after age 70	0.625	0.484	0	1	11,447
Father education (high school or higher)	0.404	0.491	0	1	11,447
Bad health ($t+1$)	0.237	0.425	0	1	11,289
Household income ($t+1$) ($1,000)	50.297	184.589	0.103	17,600	11,289
Log of household income ($t+1$)	10.307	1.036	4.638	16.684	11,289

NOTES: The last three variables in the table are observed only for those who survive to the second period.

TABLE 3
MAMMOGRAM CHOICES BY EACH SHORT PANEL DEFINED BY CONSECUTIVE WAVES OF HRS

Panels	Total	No	Yes
Wave 3→4	3,899	1,059 (27.16%)	2,830 (72.84%)
Wave 4→5	636	135 (21.23%)	501 (78.77%)
Wave 5→6	3,476	771 (22.18%)	2,705 (77.82%)
Wave 6→7	45	7 (15.56%)	38 (84.44%)
Wave 7→8	3,372	741 (21.98%)	2,631 (78.02%)
Wave 8→9	29	5 (17.24%)	24 (82.76%)
Total	11,447	2,718	8,729

5.3. *Descriptive Statistics.* Table 2 provides summary statistics of the key variables for the final analysis sample. The sample of women has an average age of 57.82. A large majority of the sample are non-Hispanic white (79.8%) and with at least high school education (79.8%). The average household income in our sample is about 50,972 dollars; 21.7% of the sample has self-reported bad health, and about 76.2% of the women undertook mammogram in the survey year. About 1.4% of the surveyed women died within two years; 72.1% of our sample reported having at least one insurance plan. Finally, about 76.8% of the mothers (and about 61.1% of the fathers) of the women in our sample are either still alive or died at an age greater than 70 at the time of the interview; 43.1% of the mothers (and about 40.7% of the fathers) of the women in our sample finished at least high school. From one wave to the next wave, there is a slight increase in the fraction of survivors who reported bad health (an increase from 21.7% to 23.7%).

Table 3 provides more detailed information about women's mammogram choices in our sample by wave. Because our short panels are defined by the consecutive waves of HRS, the panel "wave $t \to t+1$" consists of women who were interviewed in waves t and $t+1$ and

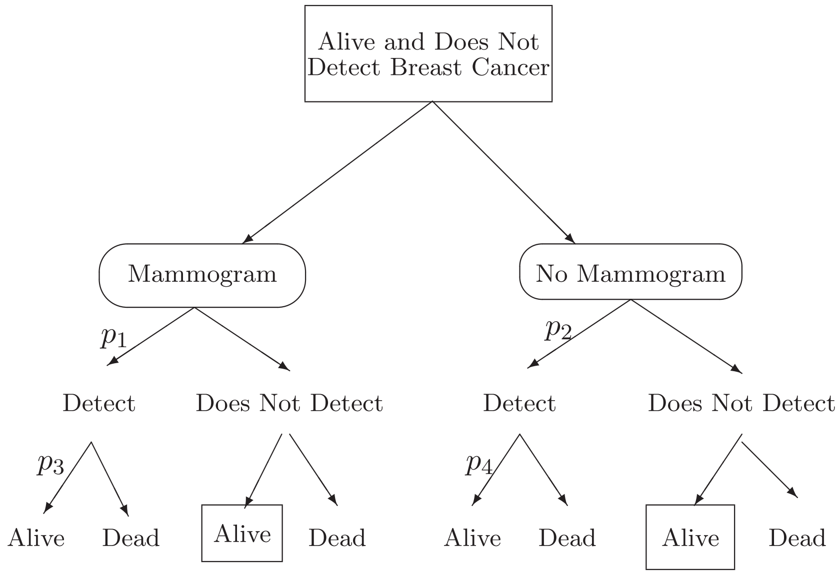

NOTES: (1) $p_1 > p_2$: mammogram can detect breast cancer at its early stage; $p_3 > p_4$: survival rate is higher when breast cancer is detected at earlier stage. (2) The states with rectangular frame box are those in which she will keep making decisions on whether to undertake mammography.

FIGURE 1

THE TIMELINE FOR MAMMOGRAPHY DECISIONS

reported their mammography decisions (at least) in wave *t*. The sizes of the panels change because HRS do not ask all (or the same numbers of) women of their mammogram use every year. For example, in 1998 (wave 4), only those who were added to the HRS were asked.[35] In 2000 and 2004, every woman in the survey was asked, but in 2002 and 2006, only a few women in a special module were asked this question. Though the number of women asked about their mammography decisions changed over time, the fractions of women who did not undertake mammogram stayed relatively stable, ranging from about 16% to about 27%, with an average noncompliance rate of about 24%.

5.4. *Decision Timeline.* Figure 1 depicts the timeline for mammography decisions for women in our sample. As we mentioned earlier, we only consider women who are alive and have not yet been diagnosed with any cancer (thus not breast cancer) in the first period. Given her period-1 state variables, she makes the decision of whether to undertake mammography. Mammography detects breast cancer with very high probability, though not for certain, if the woman has breast cancer. In the event that the woman has breast cancer, early detection of breast cancer will lead to higher survival probability.

To fully capture the diagnostic nature of mammography, we would need to have information about whether the woman has breast cancer at any period and estimate the probability of detecting breast cancer with (p_1 in Figure 1) and without (p_2 in Figure 1) mammogram. However, we do not have access to such data. In HRS, even though we have information on women's mammography choices from 1996 on and we know whether their doctors have told them that

[35] Because the new cohort added in 1998 (wave 4) is the cohort born around the Great Depression (1923–1930), they were older than the age group we consider in our empirical setting. That is why we have fewer observations for the short panel defined by wave 4→5.

they have *any* cancer, we do not have information on which kind of cancer they have been diagnosed with.[36]

Due to these data issues, we decide to go directly from the mammography decision to the live/death outcome and health status if alive (see our empirical specification below), without going through the intermediate step (having breast cancer or not). That is, we simply capture the ultimate effect of undertaking mammography as to lower the probability of dying and to lower the probability of being in bad health status if alive.

5.5. *Empirical Specification.* For this application, we assume that each woman in our sample decides whether or not to undertake mammography each period ($i = 1$ if she does and $i = 0$ if she does not). From the previous section on identification, we need to impose some normalization of the contemporaneous utilities. We normalize the individual's instantaneous utility at the death state to be zero. Note that no decision is necessary if one reaches the death state.

Now we describe the state variables we use in our empirical specification. They include Age (AGE); Education Status (HIGHSCHOOL); Bad Health (BADHEALTH), which indicates whether the individual self reports bad health; Log of Household Income (LOGINCOME); Death (DEATH); and whether her mother is still alive or died at age greater than 70 (MOTHER70).

In our empirical specification, we assume that the instantaneous payoff from undertaking mammogram, $u_1(x)$, relative to the instantaneous payoff from not undertaking mammogram, $u_0(x)$, depends on the individual's health status ("BADHEALTH") and log income ("LOGINCOME"). "BADHEALTH" is a binary variable indicating whether the agent is in bad health or not at time t; "LOGINCOME" denotes the logarithm of household income of the agent at time t. In our empirical analysis, we treat log income as a continuous variable in the first-step estimations of the probabilities of choosing mammography and the probabilities of dying (see Tables 4 and 5).[37] For the estimation of the state transitions and the estimation of the utility function parameters, we discretize log incomes into 10 intervals.[38]

We implement the estimation method proposed in Section 4 to nonparametrically estimate utility values corresponding to $u_1(x) - u_0(x)$. However, in order to save space, in the *reporting* of these nonparametrically estimated utility values, we report the regression coefficients of these estimated values on BADHEALTH and LOGINCOME, i.e., we report α_0, $\alpha_{1,}$ and α_2 in the following regression:

$$u_1(x) - u_0(x) = \alpha_0 + \alpha_1 \text{ BADHEALTH} + \alpha_2 \text{LOGINCOME}. \tag{29}$$

It is important to remark that even though the other state variables, AGE, WHITE, and HIGHSCHOOL, do not show up in our specification of $u_1(x) - u_0(x)$, it does not mean that these variables do not affect the instantaneous utility of the individual; what it means is that these variables affect the instantaneous utility under action 1 (mammogram) and action 0 (no mammogram) in exactly the same way (see Remark 7).

The agents make their decisions about whether to get mammography by comparing the expected summations of current and discounted future utilities from each choice. Individuals are uncertain about their future survival probabilities and, if alive, the transition probabilities of future health and income. These probabilities depend on their choices about whether or not to get mammography; time-variant state variables including their lagged health status, their lagged income, and their age, denoted as AGE; and time-invariant state variables including their race, denoted by a binary variable WHITE, their education status HIGHSCHOOL, and the

[36] Even for those over age 65 with matching Medicare claim data, the number of observations is not big enough for us to get the needed probabilities of being diagnosed with breast cancer with and without mammography stratified by all the state variables we want to control in our model such as age, race, health status, income, and education status.

[37] We then use the estimated Logit to predict the probabilities of choosing mammography and probabilities of dying corresponding to each log income interval.

[38] Such discretization is necessary to implement our nonparametric estimation method because we need a sufficient number of observations in each cell (defined by the state vector) to obtain precise estimates of the state transitions.

TABLE 4
DETERMINANTS OF MAMMOGRAPHY DECISIONS: THE CHOICE PROBABILITIES FROM LOGIT REGRESSION

Variable	(1)	(2)	(3)	(4)	(5)	(6)
Bad health	0.146**	0.113*	0.161***	0.099*	0.095	0.079
	(0.054)	(0.063)	(0.057)	(0.056)	(0.058)	(0.059)
Log income	0.253***	0.264***	0.252***	0.290***	0.287***	0.324***
	(0.027)	(0.028)	(0.026)	(0.025)	(0.026)	(0.026)
Married	0.116***	0.106*	0.086	0.063	0.080	0.127**
	(0.056)	(0.058)	(0.053)	(0.053)	(0.055)	(0.055)
White	−0.265***	−0.206***	−0.269***	−0.212***	−0.270***	−0.244***
	(0.060)	(0.063)	(0.056)	(0.056)	(0.059)	(0.059)
Insurance	0.537***	0.513***	0.519***	0.558***	0.563***	
	(0.055)	(0.056)	(0.052)	(0.052)	(0.053)	
High school	0.344***	0.324***	0.370***			0.414***
	(0.061)	(0.063)	(0.056)			(0.060)
Mother70	−0.090*		−0.058			−0.089*
	(0.054)		(0.051)			(0.054)
MotherHighSchool	0.112**				0.179***	0.115**
	(0.050)				(0.048)	(0.049)
Father70		0.089*		0.082*		
		(0.048)		(0.044)		
FatherHighSchool		0.081				
		(0.051)				
Age	0.012**	0.011*	0.010*	0.010*	0.011**	0.009
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Constant	−2.607***	−2.800***	−2.460***	−2.718***	−2.733***	−2.918***
	(0.448)	(0.467)	(0.426)	(0.430)	(0.442)	(0.446)
Pseudo-R^2	0.043	0.042	0.040	0.037	0.040	0.034

NOTES: (1) Robust standard errors are in parenthesis; (2) the included variables in each specification correspond to those in Table 6; (3) *, **, and *** represent statistical significance at 10%, 5%, and 1%, respectively.

longevity of their mothers, denoted by the binary variable MOTHER70, which takes value 1 if the mother is still alive or died at the age older than 70 and 0 otherwise, and MOTHERHIGHSCHOOL (a binary variable which takes value 1 if mother finished high school or more and 0 otherwise), or analogously, FATHER70 and FATHERHIGHSCHOOL (see Table 5 for the specifications).

Recall that in our setup described in Section 2, we assumed infinite horizon and that the state space $\mathcal{X}$ is finite. To make our empirical setup consistent with the above setup, in our empirical application, we include "age" as one of the state variables, but we group ages 64 and above as one state labeled as "64+." This allows our empirical specification to be consistent with all three assumptions in our setup, namely, infinite horizon, stationarity, and finite state space.

5.5.1. *Exclusion variables.* From the assumptions for exclusion restriction variables (see Assumption 5 and Remark 7), we know that any variables that do not enter the relevant instantaneous payoff (after normalization), i.e., $u_1(x) - u_0(x)$, but affect the transition of payoff relevant state variables can qualify as exclusion variables. Since only the variables BADHEALTH, LOGINCOME, and DEATH enter the instantaneous payoff functions (after normalization), the other variables that affect the transition of the above three variables can all qualify as potential exclusion restriction variables, including MOTHER70, WHITE, MARRIED, HIGHSCHOOL, AGE, or any other variables that one may find to have important effects on the transition of the state variables relevant to the instantaneous payoff functions, but do not directly affect the instantaneous payoff function (e.g., we experimented with insurance status in some of our specifications).[39] In what follows, we report estimation results under six different sets of exclusion variables. The sets of exclusion restriction variables are listed in the last panel in Table 6, and they are also

[39] See, e.g., Ayanian et al. (1993) and Decker (2005) for the relationship between health insurance and health outcomes for women with breast cancer.

TABLE 5
DETERMINANTS OF PROBABILITY OF DYING IN TWO YEARS FROM LOGIT REGRESSIONS

Variable	(1)	(2)	(3)	(4)	(5)	(6)
Mammogram	−0.459***	−0.446***	−0.435***	−0.442***	−0.436***	−0.446***
	(0.172)	(0.181)	(0.159)	(0.161)	(0.170)	(0.170)
Bad health	1.750***	1.725***	1.696***	1.631***	1.698***	1.780***
	(0.187)	(0.189)	(0.177)	(0.173)	(0.184)	(0.186)
Log income	−0.208**	−0.225***	−0.173**	−0.161**	−0.166**	−0.230***
	(0.085)	(0.082)	(0.081)	(0.079)	(0.085)	(0.079)
Married	0.092	0.136	−0.067	−0.072	0.093	0.052
	(0.189)	(0.195)	(0.173)	(0.171)	(0.186)	(0.188)
White	−0.137	−0.088	−0.148	−0.162	−0.202	−0.110
	(0.187)	(0.199)	(0.168)	(0.170)	(0.181)	(0.186)
Insurance	−0.151	−0.151	−0.122	−0.066	−0.139	
	(0.190)	(0.196)	(0.177)	(0.178)	(0.184)	
High school	0.273	0.446**	0.116			0.251
	(0.198)	(0.207)	(0.173)			(0.192)
Mother70	−0.273		−0.239			−0.286
	(0.179)		(0.165)			(0.178)
MotherHighSchool	0.122				0.154	0.099
	(0.190)				(0.185)	(0.188)
Father70		−0.257		−0.165		
		(0.172)		(0.154)		
FatherHighSchool		−0.254				
		(0.204)				
Age	0.075***	0.050**	0.073***	0.072***	0.079***	0.079***
	(0.021)	(0.022)	(0.020)	(0.020)	(0.021)	(0.021)
Constant	−6.890***	−5.412***	−6.891***	−6.905***	−7.528***	−6.962***
	(1.534)	(1.586)	(1.391)	(1.454)	(1.552)	(1.516)
Pseudo-R^2	0.106	0.101	0.106	0.099	0.102	0.107

NOTES: (1) Robust standard errors are in parenthesis; (2) the included variables in each specification correspond to those in Table 6; (3) ** and *** represent statistical significance at 5% and 1%, respectively.

reflected in the specifications of the first-step estimates reported in Tables 4 and 5 (as well as Figure 2), where each column corresponds to a different set of exclusion variables. As is clear from the last panel in Table 6, we always include WHITE, AGE, and MARRIED as exclusion variables for all six specifications, but they differ in whether we include the woman's education and insurance status as well as whether we use MOTHER70 and MOTHERHIGHSCHOOL combination or FATHER70 and FATHERHIGHSCHOOL combination.

5.6. *Estimation Results.*

5.6.1. *First-step estimates.* As we noted earlier, our estimation strategy has two steps. In the first step, we need to use the data to estimate choice probabilities and the state transitions. Here, we report these first-step estimation results. The choice probabilities and the death probability are estimated using Logit regressions, but the transition of BADHEALTH and LOGINCOME is estimated nonparametrically.

Logit estimates of the probability of choosing mammography as a function of the state variables. Table 4 produces the reduced form Logit regression results for the determinants of whether a woman will undertake mammogram in a given year, with each column corresponding to one of the six specifications on the exclusion variables. Column (1) uses MOTHER70 and MOTHERHIGHSCHOOL together with other exclusion variables, whereas column (2) uses FATHER70 and FATHERHIGHSCHOOL instead. In the specification reported in column (1) where we use MOTHER70 and MOTHERHIGHSCHOOL together with other exclusion variables, we find that women who are married with higher household income, with high school education, and with

TABLE 6

PARAMETER ESTIMATES FOR THE INSTANTANEOUS UTILITY FUNCTION AND TIME PREFERENCE PARAMETERS UNDER SIX DIFFERENT SETS OF EXCLUSIVE RESTRICTION VARIABLES

	(1)	(2)	(3)	(4)	(5)	(6)
Panel (A) Instantaneous Utility Function Parameters						
Bad health	−0.434***	−0.724***	−0.138***	−0.913***	−0.335***	−0.472***
	(0.114)	(0.260)	(0.061)	(0.120)	(0.053)	(0.103)
Log income	1.177***	1.167***	1.346***	1.153***	1.265***	1.280***
	(0.031)	(0.104)	(0.062)	(0.106)	(0.032)	(0.072)
Constant	−0.811***	−0.928***	−2.732***	−0.926***	−1.722***	−2.014***
	(0.256)	(0.199)	(0.064)	(0.163)	(0.054)	(0.091)
Panel (B) Time Preference Parameters						
δ	0.681***	0.792***	0.741***	0.947***	0.759***	0.764***
	(0.123)	(0.098)	(0.058)	(0.100)	(0.020)	(0.058)
β	0.679***	0.791***	0.679***	0.508***	0.578***	0.762***
	(0.187)	(0.193)	(0.298)	(0.109)	(0.074)	(0.185)
$\tilde{\beta}$	1.000***	1.000***	0.984***	1.000***	1.000***	1.000***
	(0.282)	(0.247)	(0.496)	(0.105)	(0.027)	(0.281)
Panel (C) Hypothesis Tests						
$H_0 : \beta = 1$	Reject	Reject	Reject	Reject	Reject	Reject
$H_0 : \tilde{\beta} = \beta$	Reject	Reject	Reject	Reject	Reject	Reject
Exclusion Variables:						
White	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes
Married	Yes	Yes	Yes	Yes	Yes	Yes
HighSchool	Yes	Yes	Yes	No	No	Yes
Insurance	Yes	Yes	Yes	Yes	Yes	No
Mother70	Yes	No	No	No	No	Yes
MotherHighSchool	Yes	No	Yes	No	Yes	Yes
Father70	No	Yes	No	Yes	No	No
FatherHighSchool	No	Yes	No	No	No	No

NOTES: (1) The last panel indicates the exclusive restriction variables used in the specification in that column, with "Yes" meaning the variable is used and "No" otherwise; (2) standard errors for parameter estimates are in parenthesis, and *** represents statistical significance at 1%; (3) for hypothesis tests reported in panel C, all are rejected with p-value less than 0.01.

health insurance are more likely to undertake mammograms, and white women are less likely than black women to undertake mammogram. All these coefficient estimates are significant at 1%. Interestingly, we also find that women whose mothers are still alive or died after age 70 are less likely, but those whose mothers have at least high school are more likely to undertake mammograms. Finally, older women and women with bad health are more likely to undertake mammogram, and these coefficients are statistically different at the 5% and 10% level, respectively.

In the specification reported in column (2), where we use the information about women's father as the exclusion variables, none of the other coefficient estimates change qualitatively. Interestingly, we found that FATHER70 and FATHERHIGHSCHOOL both positively affect women's probability of undertaking mammogram. The estimated coefficients for BADHEALTH and LOGINCOME are qualitatively very similar in the specifications reported in columns (3)–(6) when we use different set of exclusion variables.[40]

[40] Note that the coefficient estimates from such reduced form regressions, while informative, do not shed light on the mechanisms under which the observed relationships between a variable and the mammogram decision arise. For example, we see in Table 4 that women whose mother is still alive or died after age 70 are less likely to undertake mammogram. But it is not clear from the table why. The structural analysis we undertake below will help us to achieve a better understanding of the observed reduced form relationship reported in Table 4.

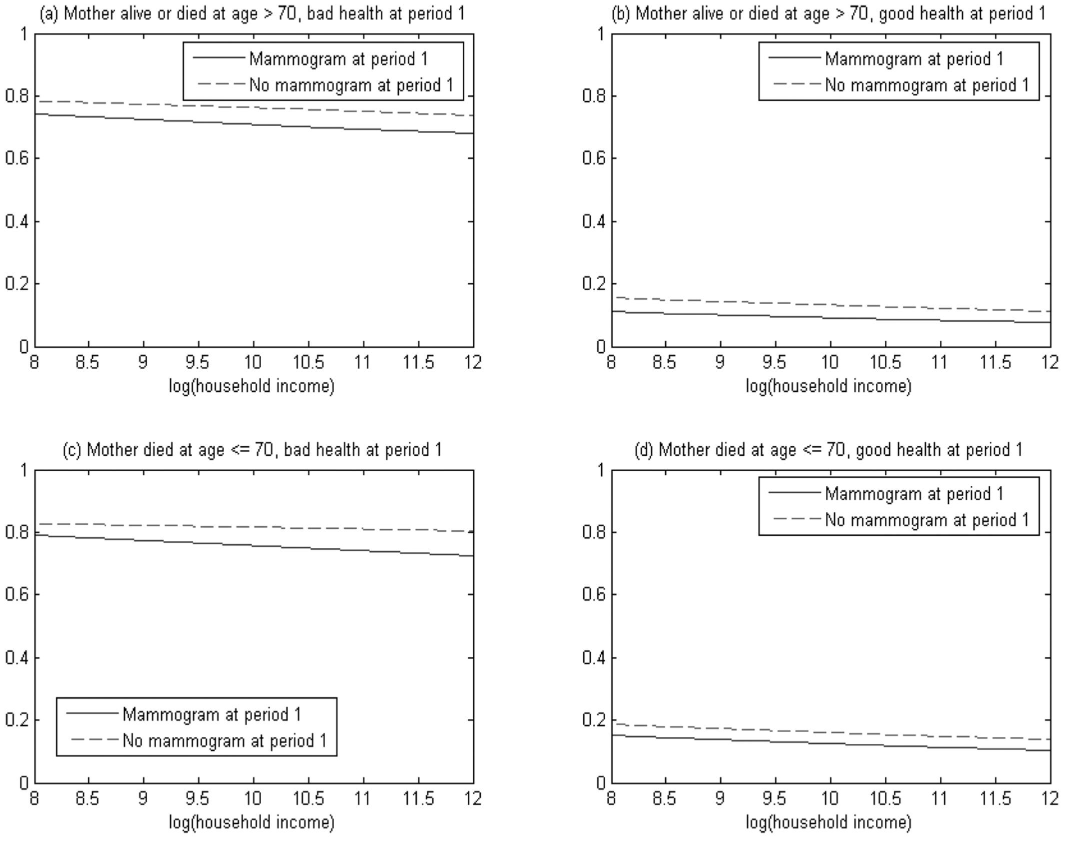

FIGURE 2

NONPARAMETRIC ESTIMATE OF THE PROBABILITY OF BAD HEALTH AS A FUNCTION OF LOGINCOME, CONDITIONAL ON MOTHER70, BADHEALTH, AND MAMMOGRAM IN PREVIOUS PERIOD

Transition probabilities of the payoff relevant state variables: Determinants of the probability of dying in two years. Table 5 reports the Logit regression results for the probability of dying in two years, again using six different sets of exclusion variables corresponding to the six columns. In all six specifications of the exclusion variables, the coefficient estimates show that undertaking mammogram significantly lowers the probability of death (notice that the average age of the sample is about 58 years). Not surprisingly, women with bad health are more likely to die, but mammogram reduces the probability of dying conditional on bad health; moreover, white, richer, and younger women are less likely to die. Also note that women whose mothers are either still alive or died after age 70 as well as women whose fathers are either still alive or died after age 70 are less likely to die even though the estimates are not statistically significant. This suggests a possible genetic link of longevity between daughters and their mothers as well as fathers.

Transition probabilities of the payoff relevant state variables: Evolution of bad health in two years. Figure 2 depicts a subset of the results from the nonparametric estimation of the evolution of BADHEALTH for a selective combinations of the other state variables and the mammogram choice for our exclusion restriction variable specification listed in column (1) in Table 6. For example, panel A shows that how the probability of having bad health in period 2 is affected by mammogram decision in the previous period and LOGINCOME for women whose mothers are alive or died after age 70 and had bad health in the previous period. It shows that the probability of bad health decreases with Log Income and is lower if one undertakes mammogram in the previous period. A similar negative relationship between the probability of bad health and income is also shown in other panels. Across all panels, we see that mammogram always reduces the probability of bad health, good health in the previous period predicts a higher probability of good health this period, and mother living or having died after 70 reduces the probability of bad health. A similar nonparametric estimation of LogIncome is also conducted.

5.6.2. *Second-step estimates: Utility parameters and discount factors.* Table 6 reports the estimation results for the parameters in the instantaneous utility function specification (29) and the identified discount factors δ, β, and $\tilde{\beta}$ as well as standard errors calculated from the asymptotic distributions of the maximum pseudolikelihood estimator. As presented in Tables 4 and 5, we report results from six specifications using different sets of exclusion variables.[41] The set of exclusion variables used in each specification is reported in the last panel. For example, in the specification reported in column (1), we include only White, Age, Married, HighSchool, Insurance, Mother70, and MotherHighSchool as the exclusion restriction variables, whereas in the specification reported in column (2), we replace Mother70 and MotherHighSchool by Father70 and FatherHighSchool. These are our two main specifications; the other specifications reported in columns (3)–(6) provide robustness checks.

The parameter estimates, both for the instantaneous payoff function parameters and the time preference parameters, are quite similar both qualitatively and quantitatively across all six specifications. In panel A of Table 6, we find that having bad health lowers the utility of undertaking mammography relative to not undertaking mammography. This may seem counter-intuitive in light of our finding in Table 4 that individuals with bad health are less likely to undertake mammography. This is due to the fact that undertaking mammogram significantly reduces the probability of bad health next period, more so if one is currently in bad health, as shown in Figure 2. We also found that the relative utility of undertaking mammography increases with income, consistent with the finding in Table 4 that women with higher incomes are more likely to undertake mammography. The estimated negative constant term is consistent with the idea that preventative health care typically involves huge one-time instantaneous cost, which once combined with present bias and naivety about present bias might lead to undesirable health-related decisions.

In panel B of Table 6, we present our estimates of the time preference parameters. Depending on specifications, our estimates of the standard discount factor δ range from 0.681 (column (1)) to 0.947 (column (4)); our estimates of the present bias factor β range from 0.508 (column (4)) to 0.791 (column (2)). Also interestingly, in five out of the six specifications, our estimate of the partial naivety parameter $\tilde{\beta}$ is 1 with the exception of specification (3).[42] In our two main specifications reported in columns (1) and (2), the present bias factor β is estimated to be about 0.68 and 0.79, respectively, suggesting that women exhibit substantial present bias ($\beta < 1$) as well as naivety about their present bias ($\tilde{\beta} > \beta$) when making mammography decisions.

In panel C of Table 6, we report the results from testing two hypotheses; the first null is no present bias, i.e., $\beta = 1$, and the second null is no naivety, i.e., $\tilde{\beta} = \beta$. In all six specifications, both nulls are rejected with p-value less than 0.01.

It may seem surprising that we found substantial naivety in the setting of women's mammography decisions, which is a not a high-frequency decision. The generally accepted view is that naivety is more likely to manifest itself in high-frequency decisions (see, e.g., O'Donoghue and

[41] We imposed the restrictions that $\delta \in (0, 1]$, $\beta \in (0, 1]$, and $\tilde{\beta} \in [\beta, 1]$ in the estimation results reported in Table 6. In online Appendix G (Supporting Information), we also report in Table G2 estimation results when $\tilde{\beta}$ is relaxed to $(0, 1]$, and in Table G3 we report the estimation results when β and $\tilde{\beta}$ are both relaxed to be $(0, 2]$. In both cases, we show that the estimates do not change much except for standard errors, and we cannot formally reject the hypothesis that $\tilde{\beta} = 1$.

[42] Because the estimated $\tilde{\beta}$ is on the boundary of the parameter space, we follow Andrews (1999, 2001) and in particular Moran (1971) to compute the correct standard errors. See online Appendix F.1 (Supporting Information) for more details.

TABLE 7
MAMMOGRAPHY COMPLIANCE RATES PREDICTED BY THE MODEL AND IMPLIED BY DIFFERENT COUNTERFACTUAL EXPERIMENTS

	(1)	(2)	(3)	(4)	(5)	(6)
Data	0.76236	0.76307	0.75892	0.75841	0.76152	0.76179
Model	0.76440	0.76603	0.75781	0.76015	0.76100	0.76259
Counterfactual experiments:						
[1] No naivety: $\tilde{\beta} = \beta[= \hat{\beta}]$	0.76442	0.76607	0.75786	0.76061	0.76104	0.76262
[2] No naivety and no present bias: $\tilde{\beta} = \beta = 1$	0.78673	0.79121	0.78796	0.85860	0.79025	0.78498

NOTES: (1) The sample sizes slightly vary as we change the set of the exclusive restriction variables, which explains the changes in the mammography compliance rates in the data; (2) the exclusive variables used in each column correspond to those of the same column in Table 6.

Rabin, 1999a). However, in the mammography setting, present bias and naivety affect women's mammography decision both because there is an immediate cost of undertaking mammography, as indicated by the estimated parameters of $u_1(x) - u_0(x)$ reported in panel A of Table 6, and also because the benefit of mammography, which is the reduced probability of bad health and ultimately low probability of dying, can be very delayed as well.

It is also interesting to compare our estimates for $\beta\delta$ with what is in the literature. Using the estimates of β and δ reported in Table 6, we have that our estimates for $\beta\delta$ range from 0.44 (column 5) to about 0.50 (columns (1), (3), and (4)) to about 0.58 (column (6)) and 0.63 (column (2)). This can be compared with the estimate in Fang and Silverman (2009) where they estimate β to be 0.338 and δ to be 0.88, with $\beta\delta \approx 0.30$ for a group of single mothers with dependent children. It is important to note that our sample period is two years, whereas Fang and Silverman's (2009) sample period is one year. Also the sample of women in this article are older and have very different social economic status (e.g., education and income) from the sample in Fang and Silverman (2009).

5.7. *Counterfactual Experiments*. Table 7 reports the mammography compliance rates predicted by the model and implied by two counterfactual experiments where in experiment 1, we assess the mammography rates predicted by the model if we hypothetically set $\tilde{\beta}$ equal to the estimated $\hat{\beta}$, and in experiment 2, we set both $\tilde{\beta}$ and β to 1. Experiment 1 allows us to assess the impact of naivety on mammography take-up rate, whereas experiment 2 allows us to assess the impact of both present bias and the naivety about present bias.

We find that in experiment 1, when women in our sample are present biased but fully sophisticated, the mammography compliance rate only increases slightly in all six specifications. For example, in the specification reported in column (1), we find that the mammography compliance rate ranges from 76.440% in the baseline model prediction to 76.442% in the counterfactual experiment with full sophistication, and both are very similar to the compliance rate in the data (76.236%). The qualitative conclusions for this experiment are the same in other specifications.[43,44]

[43] It is somewhat surprising that eliminating naivety in our setting does not seem to make much of a difference in terms of choices. Indeed preference reversal is the standard source of identification for both present bias and naivety parameters (Ainslie, 1992; Loewenstein and Elster, 1992). However, preference reversals can only be observed in laboratory settings where subjects are asked to make hypothetical choices for the distant future and then to make the same choice again as the future looms closer. We use observational data in our analysis, and as a result the main source of identification is not preference reversals, but how exclusion variables affect choices over time.

[44] In order to understand whether the small effect of no naivety is due to (i) removing naivety (i.e., change $\tilde{\beta}$ from 1 to $\tilde{\beta} = \beta$) has a relatively small impact on perceived future outcomes or (ii) even a relatively large change in perceived future outcomes has a relatively small impact on one's optimal choice, we conducted counterfactual experiments where we assume that the effect of LOGINCOME on future instantaneous utility $u_1 - u_0$ is five percentage points higher than the estimated effect (i.e., adding 0.05$*$LOGINCOME to the estimated effect as reported in Table 6). We found in these experiments that removing naivety generates a much larger effect on the mammography rate (about 89% with sophisticated hyperbolic discounter vs. about 84% with naive hyperbolic discounter).

In experiment 2, where the agents are exponential discounters, we find that the mammography compliance rate goes up more substantially. Depending on the specifications, the mammography compliance rate increases by 2.333 [= 78.673 − 76.440] percentage points in specification 1, 2.518 percentage points in specification 2, and up to 9.845 percentage points in specification 4. These represent substantial reductions in the mammography *noncompliance* rate, ranging from a 9.90% reduction in specification 1 to 10.789% reduction in specification 2 and up to 41.05% reduction specification 4. Thus, time-inconsistent preferences caused by present bias and naivety about present bias indeed have significant policy implications for the low compliance rates of preventive health care in the United States.

6. CONCLUSION AND DISCUSSION

This article extends the semiparametric identification and estimation method for dynamic discrete choice models using Hotz and Miller's (1993) conditional choice probability approach to the setting where individuals may have hyperbolic discounting time preferences and may be naive about their time inconsistency.

Our analysis showed that the three discount factors, the present bias factor β, the standard discount factor δ, and the perceived present bias factor $\tilde{\beta}$ for naive agents can be separately identified. The key identifying restriction is that there exist variables that do not directly enter the instantaneous utility function but affect the transition of other payoff relevant state variables.

We proposed an estimation strategy based on the identification argument and implemented the proposed estimation method to the decisions of undertaking mammography to evaluate the importance of present bias and naivety in the underutilization of mammography. Our estimates are consistent with the presence of both present bias and naivety about present bias. In our counterfactual experiments, we found that time-inconsistent preferences caused by present bias and naivety about present bias indeed have significant policy implications for mammography take-up rates.

We assumed in the main text that the model is stationary. In many applications, such a stationarity assumption may not be valid. However, we show in online Appendix H (Supporting Information) that the identification arguments, properly modified, still work for the case of finite horizon models with nonstationary state transitions. Of course, estimating finite horizon models with nonstationary state transitions requires longer panels. When one has longer panel data sets, it may also be important to allow for habit effects in individuals' choice model. Such habit effects can be incorporated in our framework by including a habit stock variable summarizing individuals' past choices.

In this article, we assume, as in most (if not all) of the existing literature, both the state space and the control variables are discrete.[45] Conceptually the discrete state space assumption can be easily extended to allow for continuous state variables. For continuous state variables, the state transition matrix needs to be replaced by state transition densities, and almost all of the analysis will continue to hold. If the control variables are continuous, then the optimal choice will be characterized by first-order conditions, even for potentially naive hyperbolic discounting agents. The identification and estimation of dynamic programming models with continuous controls for exponential discounting agents have been studied in a recent paper by Schrimpf (2011), where he shows that under some exclusion restrictions similar to those in our article, together with the assumption that one of the partial derivatives of the payoff function is known, the payoff function is nonparametrically identified by the observed distribution of states and controls. We believe that similar extension to hyperbolic discounting agents for dynamic programming problems with continuous controls can be obtained.

In this article, we assumed that other than the choice-specific idiosyncratic payoff shocks, which we assume to be serially and cross-sectionally independent, we do not allow for any

[45] For example, Rust (1994) and Magnac and Thesmar (2002) both made the same set of assumptions. See Aguirregabiria and Mira (2010) for a recent survey of the literature.

observed heterogeneity or unobserved state variables among individuals. Recent results by Kasahara and Shimotsu (2009) and Hu and Shum (2012) show that the conditional choice probabilities, which are crucial for implementing Hotz–Miller type estimators, can be identified in the presence of unobserved heterogeneity or unobserved state variables in the context of dynamic discrete choice models with exponential discounting.[46] We conjecture that their arguments can be generalized to our setting of dynamic discrete choice models with hyperbolic discounting preferences. Once conditional choice probabilities for any given state, including both observed and unobserved state variables, can be identified, the estimation methods we proposed here can then be applied. We will examine this important extension in future research.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Online Appendix A–H
Table A1: Summary of Key Notations.
Table G2: Parameter Estimates for the Instantaneous Utility Function and Time Preference Parameters Under Six Different Sets of Exclusive Restriction Variables with $\tilde{\beta} \in [0, 1]$.
Table G3: Parameter Estimates for the Instantaneous Utility Function and Time Preference Parameters Under Six Different Sets of Exclusive Restriction Variables with $\beta \in [0, 1]$ and $\tilde{\beta} \in [0, 1]$.

References

AGUIRREGABIRIA, V., "Pseudo Maximum Likelihood Estimation of Structural Models Involving Fixed-Point Problems," *Economics Letters* 84 (2004), 335–40.
———, AND P. MIRA, "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models," *Econometrica* 70 (2002), 1519–43.
———, AND ———, "Sequential Estimation of Dynamic Discrete Games," *Econometrica* 75 (2007), 1–53.
———, AND ———, "Dynamic Discrete Choice Structural Models: A Survey," *Journal of Econometrics* 156(1) (2010), 38–67.
AINSLIE, G., *Picoeconomics: The Strategic Interaction of Successive Motivational States Within the Person* (Cambridge, U.K.: Cambridge University Press, 1992).
ANDREWS, D. W. K., "Estimation When a Parameter Is on a Boundary," *Econometrica* 67(6) (1999) 1341–83.
———, "Testing When a Parameter Is on the Boundary of the Maintained Hypothesis," *Econometrica* 69(3) (2001), 683–734.
ARCIDIACONO, P., AND R. MILLER, "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity," *Econometrica* 79(6) (2011), 1823–68.
AYANIAN, J. Z., B. A. KOHLER, T. ABE, AND A. M. ESPTEIN, "The Relation between Health Insurance Coverage and Clinical Outcomes among Women with Breast Cancer," *The New England Journal of Medicine* 329(5) (1993), 326–31.
BAJARI, P., C. L. BENKARD, AND J. LEVIN, "Estimating Dynamic Models of Imperfect Competition," *Econometrica* 75(5) (2007), 1331–70.
———, H. HONG, AND D. NEKIPELOV, "Game Theory and Econometrics: A Survey of Some Recent Research," in D. Acemoglu, M. Arellano, and E. Dekel, eds., *Advances in Economics and Econometrics: Tenth World Congress (Volume III, Econometrics)* (Cambridge, U.K.: Cambridge University Press, 2013), 3–52.
BARRO, R. J., "Ramsey Meets Laibson in the Neoclassical Growth Model," *Quarterly Journal of Economics* 114 (1999), 1125–52.
CARRILLO, J. D., AND T. MARIOTTI, "Strategic Ignorance as a Self-Disciplining Device," *Review of Economic Studies* 67 (July 2000), 529–44.

[46] Arcidiacono and Miller (2011) proposed an EM algorithm to estimate such models with unobserved heterogeneity using conditional choice probabilities.

CHUNG, D., T. J. STEENBURGH, AND K. SUDHIR, "Do Bonuses Enhance Sales Productivity? A Dynamic Structural Analysis of Bonus-Based Compensation Plans," Mimeo, Yale School of Management, 2009.

DECKER, S. L., "Medicare and the Health of Women with Breast Cancer," *Journal of Human Resources* 40(4) (2005), 948–68.

DEGNAN, D., R. HARRIS, J. RANNEY, D. QUADE, J. A. EARP, AND J. GONZALEZ, "Measuring the Use of Mammography: Two Methods Compared," *American Journal of Public Health* 82(10) (1992), 1386–88.

DELLA VIGNA, S., AND MALMENDIER, U., "Contract Design and Self-Control: Theory and Evidence," *Quarterly Journal of Economics* 119(2) (2004), 353–402.

———, AND ———, "Paying Not to Go to the Gym," *American Economic Review* 96(3) (2006), 694–719.

———, AND D. PASERMAN, "Job Search and Impatience," *Journal of Labor Economics* 23(3) (July 2005), 527–88.

FANG, H., AND D. SILVERMAN, "On the Compassion of Time-Limited Welfare Programs," *Journal of Public Economics* 88 (2004), 1445–70.

———, AND ———, "Distinguishing between Cognitive Biases: Belief vs. Time Discounting in Welfare Program Participation," in E. J. McCaffery and J. Slemrod, eds., *Behavioral Public Finance* (New York: Russell Sage Foundation, 2006), 47–81.

———, AND ———, "Time-inconsistency and Welfare Program Participation. Evidence from the NLSY," *International Economic Review* 50(4) (2009), 1043–76.

GINÉ, X., J. GOLDBERG, D. SILVERMAN, AND D. YANG, "Revising Commitments: Time Preference and Time-Inconsistency in the Field," Mimeo, University of Michigan, 2011.

———, KARLAN, D., AND J. ZINMAN, "Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation," *American Economic Journal: Applied Economics* 2(4) (2010), 213–35.

GRUBER, J., AND B. KOSZEGI, "Is Addiction "Rational"? Theory and Evidence," *Quarterly Journal of Economics* 116(4) (2001), 935–58.

HAUSMAN, J. A., "Individual Discount Rates and the Purchase and Utilization of Energy-using Durables," *Bell Journal of Economics* 10(1) (1979), 33–54.

HAYASHI, F., *Econometrics* (Princeton: Princeton University Press, 2000).

HOTZ, J., AND R. MILLER, "Conditional Choice Probabilities and Estimation of Dynamic Models," *Review of Economic Studies* 60 (1993), 497–529.

HU, Y., AND M. SHUM, "Nonparametric Identification of Dynamic Models with Unobserved State Variables," *Journal of Econometrics* 171(1) (November 2012), 32–44.

JUSTER, F. T., AND R. SUZMAN, "An Overview of the Health and Retirement Study," *Journal of Human Resources* 30(5) (1995), S7–S56.

KASAHARA, H., AND K. SHIMOTSU, "Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices," *Econometrica* 77(1) (2009), 135–75.

KRUSELL, P., AND SMITH, A., "Consumption-Saving Decisions with Quasi-Geometric Discounting," *Econometrica* 71(1) (2003), 365–75.

———, B. KURUŞÇU, AND A. SMITH, "Equilibrium Welfare and Government Policy with Quasi-Geometric Discounting," *Journal of Economic Theory* 105 (July 2002), 42–72.

LAIBSON, D., "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics* 112(2) (1997), 443–77.

———, A. REPETTO, AND J. TOBACMAN, "Estimating Discount Functions with Consumption Choices over the Lifecycle," Mimeo, Harvard University, 2007.

LOEWENSTEIN, G., AND J. ELSTER, *Choice over Time* (New York: Russell Sage Foundation, 1992).

MAGNAC, T., AND D. THESMAR, "Identifying Dynamic Discrete Decision Processes," *Econometrica* 20(2) (2002), 801–16.

MAHAJAN, A., AND A. TAROZZI, "Time Inconsistency, Expectations and Technology Adoption: The Case of Insecticide Treated Nets," Mimeo, Stanford and Duke University, 2011.

MAS-COLELL, A., *The Theory of General Equilibrium: A Differentiable Approach* (Cambridge: Cambridge University Press, 1985).

MILLER, R., "Job Matching and Occupational Choice," *Journal of Political Economy* 92(6) (1984), 1086–120.

MORAN, P. A. P., "Maximum-Likelihood Estimation in Non-Standard Conditions," *Proceedings of the Cambridge Philosophical Society* 70 (1971), 441–50.

NORETS, A., AND X. TANG, "Semiparametric Inference in Dynamic Binary Choice Process," *Review of Economic Studies* 81 (2014), 1229–62.

O'DONOGHUE, T., AND M. RABIN, "Doing It Now or Later," *American Economic Review* 89(1) (1999a), 103–24.

———, AND ———, "Addiction and Self-Control," in J. Elster, ed., *Addiction: Entries and Exits* (New York: Russell Sage Foundation, 1999b), 169–206.

PAKES, A., "Patents as Options. Some Estimates of the Value of Holding European Patent Stocks," *Econometrica* 54 (1986), 755–85.

———, AND P. MCGUIRE, "Computing Markov Perfect Nash Equilibrium: Numerical Implications of a Dynamic Differentiated Product Model," *RAND Journal of Economics* 25(4) (1994), 555–89.

———, M. OSTROVSKY, AND S. BERRY, "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Samples)," *RAND Journal of Economics* 38(2) (2007), 373–99.

PASERMAN, M. D., "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation," *Economic Journal* 118(531) (2008), 1418–52.

PEETERS, R., "Stochastic Games with Hyperbolic Discounting," Mimeo, Maastricht University, 2004.

PESENDORFER, M., AND P. SCHMIDT-DENGLER, "Asymptotic Least Squares Estimators for Dynamic Games," *Review of Economic Studies* 75(3) (2008), 901–28.

PHELPS, E. S., AND R. A. POLLAK, "On Second-Best National Saving and Game-Equilibrium Growth," *Review of Economic Studies* 35(2) (1968), 185–99.

RUST, J., "Optimal Replacement of GMC Bus Engines. An Empirical Model of Harold Zurcher," *Econometrica* 55(5) (1987), 999–1033.

———, "Estimation of Dynamic Structural Models, Problems and Prospects: Discrete Decision Processes," in C. Sims and J. J. Laffont, eds., *Proceedings of the 6th World Congress of the Econometric Society* (Cambridge: Cambridge University Press, 1994a), 119–70.

———, "Structural Estimation of Markov Decision Processes," in R. Engle and D. McFadden, eds., *Handbook of Econometrics*, Vol. IV (Amsterdam: North-Holland, 1994b), 3082–139.

SCHRIMPF, P., "Identification and Estimation of Dynamic Games with Continuous States and Controls," Mimeo, University of British Columbia, 2011.

STROTZ, R. H., "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies* 23(3) (1956), 165–80.

WARNER, J. T., AND S. PLEETER, "The Personal Discount Rate: Evidence from Military Downsizing Programs," *American Economic Review* 91(1) (2001), 33–53.

WOLPIN, K. I., "An Estimable Dynamic Stochastic Model of Fertility and Child Mortality," *Journal of Political Economy* 92 (1984), 852–74.

WOOLDRIDGE, J. M., *Econometric Analysis of Cross Section and Panel Data* (Cambridge: MIT Press, 2002).