UC Berkeley

UC Berkeley Previously Published Works

Title

Identification and semiparametric estimation of a finite horizon dynamic discrete choice model with a terminating action

Permalink

<https://escholarship.org/uc/item/9c35w8sq>

Journal

Quantitative Marketing and Economics, 14(4)

ISSN

1570-7156

Authors

Bajari, P
Chu, CS
Nekipelov, D
[et al.](#)

Publication Date

2016-12-01

DOI

10.1007/s11129-016-9176-3

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-ShareAlike License, available at <https://creativecommons.org/licenses/by-nc-sa/4.0/>

Peer reviewed

Identification and Semiparametric Estimation of a Finite Horizon Dynamic Discrete Choice Model with a Terminating Action¹

Patrick Bajari, University of Washington and NBER

Chenghuan Sean Chu, Facebook

Denis Nekipelov, University of Virginia

Minjung Park, University of California, Berkeley

This version: August 29, 2016

Abstract

We study identification and estimation of finite-horizon dynamic discrete choice models with a terminal action. We first demonstrate a new set of conditions for the identification of agents' time preferences. Then we prove conditions under which the per-period utilities are identified for all actions in the agent's choice-set, without having to normalize the utility for one of the actions. Finally, we develop a computationally tractable semiparametric estimator. The estimator uses a two-step approach that does not use either backward induction or forward simulation. Our methodology can be implemented using standard statistical packages without the need to write specialized computational routines, as it involves linear (or nonlinear) projections only. Monte Carlo studies demonstrate the superior performance of our estimator compared with existing two-step estimation methods. Monte Carlo studies further demonstrate that the ability to identify the per-period utilities for all actions is crucial for counterfactual predictions. As an empirical illustration, we apply the estimator to the optimal default behavior of subprime mortgage borrowers, and the results show that the ability to identify the discount factor, rather than assuming an arbitrary number as typically done in the literature, is also crucial for obtaining correct counterfactual predictions. These findings highlight the empirical relevance of key identification results of the paper.

Keywords: Finite Horizon Optimal Stopping Problem, Time Preferences, Semiparametric Estimation

JEL Classifications: C14, C18, C50

¹We are grateful to the editor and anonymous reviewers for their insightful comments and constructive suggestions. The paper has also benefited from helpful comments by seminar participants at Chicago Booth, Olin Business School, Stanford GSB, Berkeley ARE, IO fest, Cirpée Conference on Industrial Organization, and Conference on "Recent Contributions to Inference in Game Theoretic Models" at University College London. All remaining errors are our own. Correspondence: bajari@uw.edu; seanchu@gmail.com; denis.nekipelov@gmail.com; mpark@haas.berkeley.edu.

1 Introduction

In this paper, we study finite-horizon dynamic discrete choice problems in which the agent’s set of potential choices includes a terminating action that ends the decision problem. We provide identification results and a multi-step estimation procedure for this important subset of dynamic discrete choice models.

Our first result provides conditions for the identification of agents’ time preferences, which is critical to understanding many economic and marketing problems, including consumer behavior in credit markets, purchases of durable goods, and firms’ investment decisions. In general, it is not possible to identify time preferences in dynamic discrete choice models; doing so rather requires special conditions to hold (Rust, 1994; Magnac and Thesmar, 2002). Some researchers have obtained identification using experimental data (see Frederick, Loewenstein and O’Donohue (2002) for a review of the literature). Likewise, Dubé, Hitsch and Jindal (2014) discuss a survey design that enables joint identification of utility and discount functions.

In the non-experimental literature, identification has relied on identification-at-infinity arguments, exclusion restrictions on variables that affect the state transition or future payoffs but not the current utility (Hausman, 1979; Magnac and Thesmar, 2002; Fang and Wang, 2015), or observing agents’ final-period behavior in finite horizon problems. The last approach is obviously only feasible without data truncation. In this case, the period utilities can be identified from the static decision problem in the final period, allowing the discount factor to be identified based on the remaining temporal variation in observed behavior. For the intuition in the case of continuous choice variables, see Duffie and Singleton (1997), Yao et al. (2012) and Chung, Steenburgh, and Sudhir (2014). Our paper advances this argument further by proving identification even if the agents’ final-period decisions are unobserved, for the case of discrete choice variables. The key idea is that a finite horizon model is “intrinsically” non-stationary even if all primitive objects are time-homogeneous, allowing us to identify the discount factor based on the variation over time in agents’ choice probabilities.

Our second identification result shows that, under certain conditions (including availability of final period data), the actual levels of agents’ payoffs from various actions—and not just the differences in utility between them—are identified alongside the discount factor. In such cases, there is no need to normalize the payoff from one of the actions, which the literature has demonstrated is not innocuous in dynamic settings (Bajari, Hong, and Nekipelov, 2013). In particular, Aguirregabiria (2010), Norets and Tang (2014) and Kalouptsi et al. (2016) showed that certain counterfactual conditional choice probabilities (CCPs) are not identified when only the differences in utility, but not the levels, are known. Our identification result thus has practical importance, given that counterfactual analysis is often the

ultimate goal when researchers employ structural models.

To the best of our knowledge, this is a novel identification result, as the prior literature on dynamic discrete choice models has typically either normalized or relied on additional data to pin down the per-period payoff function of one of the actions, typically the terminating action (e.g., Heckman and Navarro, 2007; Kalouptsi, 2014). The logic behind our result is as follows. In any period before the last, the agent’s continuation value from choosing a non-terminal action (but not from choosing the terminal action) includes an option value of being able to choose the terminal action in a future period. This option value depends on the level of the utility associated with the terminal action, and also diminishes toward zero in the final period because there are no remaining periods at that point. Therefore, by examining how the relative choice probabilities between the non-terminating actions and the terminating action differ in the final period compared with earlier periods, we can identify the option value of delayed termination, and thus the level of the utility associated with the terminating action (along with utilities of other choices).

In addition to providing identification results, we propose an estimation procedure that does not require backward induction or data from the final periods. The procedure is conceptually distinct from multi-step estimation procedures for infinite horizon problems such as Pesendorfer and Schmidt-Dengler (2008), as we must take into account the non-stationarity of agents’ optimal behavior due to the presence of a final period. Our estimation method exploits Hotz and Miller’s (1993) intuition that, when there is a terminating action, the continuation value can be represented as a function of the choice probabilities one period ahead. We build on that intuition by proposing a nonparametric estimator of the agents’ expectations about their continuation values, and use the estimates to recover agents’ preferences within a regression framework.² The resulting estimator involves only linear (or nonlinear) projections and does not require simulation. Our method is thus computationally light, easily scalable to data-rich settings, and can be implemented using predefined procedures from standard statistical software such as R, STATA or MATLAB.

We use a Monte Carlo study to numerically illustrate the implications of our identification results for counterfactual predictions and compare the performance of our estimator with existing two-step estimation methods. In these experiments, we simulate data based on a stylized model capturing mortgage borrowers’ default decisions. The Monte Carlo exercises show that the utility of the terminating action, whose value is often arbitrarily chosen in empirical work, could have a significant impact on counterfactual predictions. The Monte Carlos also demonstrate that our estimator is significantly faster computationally

²Our approach of nonparametrically estimating agents’ expectations and using the estimates to recover their preferences is closely related to Ahn and Manski (1993) and Manski (1991, 1993, and 2000), who examine agents’ responses to their expectations in models with uncertainty or endogenous social effects.

than existing two-step estimation methods. The results also indicate that, for the considered setting, our estimator is more robust than the two-step estimator proposed by Bajari, Benkard, and Levin (2007; BBL henceforth), which we find to be sensitive to the choice of alternative policies used in constructing the objective function as well as the choice of initial values.

Finally, we demonstrate our estimation method using real-world data on the default behavior of subprime mortgage borrowers. Our empirical analysis employs panel data on a sample of subprime mortgage borrowers whose loans were originated between 2000 and 2007. The data contain detailed borrower-level information from the loan application, including the terms of the contract, the loan-to-value ratio, the level of documentation, and the borrower’s credit score at the time of origination. We also observe the month-by-month stream of payments made by the borrower and whether the mortgage goes into default or is prepaid. To track movements in home prices, we merge the mortgage data with zip-code level home price indices. Our structural estimate of the discount factor implies a 4.9% monthly discount rate, which is significantly higher than the average contractual interest rate of the mortgages in the data, and even higher than the “risk-free” interest rate. In the literature, a typical approach to dealing with the discount factor when it cannot be estimated is to fix it to some economy-wide rate of return on assets. We demonstrate that (incorrectly) imposing a discount factor that corresponds to the market interest rate and using the resulting structural estimates for counterfactual analysis could lead to significantly biased predictions. This highlights the importance of the ability to recover the discount factor, one of the key identification results of the paper.

Although our main motivational example examines the default decisions of mortgage borrowers, our approach can be applied to many other settings. One application that has been studied extensively is students’ decision to drop out from school (e.g., Eckstein and Wolpin, 1999). The decision to drop out is the terminating action, and the graduation date is the final period of the dynamic decision problem. Another example is the decision to retire from the labor force (Rust and Phelan, 1997). In this case, retiring is the terminating action, and the final period is the terminal age at which the probability of death is 1. Our identification results and estimation method can also be applied to consumer decisions with respect to buying goods that have deadlines, such as flight- or event tickets.

The paper makes contributions to the growing theoretical literature on identification of dynamic discrete choice models, which started with Rust (1994) and Magnac and Thesmar (2002). Arcidiacono and Miller (2015a) consider the identification of dynamic discrete choice models when the time horizon for agents extends beyond the length of the data. They focus on the identification of utilities (up to a normalization) and counterfactual choice probabilities with a discrete state space and known discount

factor. Aguirregabiria (2010) examines nonparametric identification of behavioral and welfare effects of counterfactual policy interventions. Norets and Tang (2014) discuss the effect on counterfactual predictions of normalizing the utility from one of the actions. Kalouptsi et al. (2016) explore conditions for the identification of counterfactual behavior and welfare.

Our proposed estimator expresses the continuation value as a function of one-period-ahead CCPs. In this sense, we add to the literature exploiting the “finite dependence” property for estimation (Hotz and Miller, 1993; Altug and Miller, 1998; Arcidiacono and Miller, 2011; Joensen, 2009; Scott, 2013; Arcidiacono, Bayer, Blevins, and Ellickson, 2015; Aguirregabiria and Magesan, 2013; Beauchamp, 2015; Arcidiacono and Miller, 2015b). The novelty of our estimator is in its use of linear projection instead of simulation in constructing the continuation value function, and the use of OLS (for the case of linear utility) in estimating the structural parameters. By combining these features with the use of a simplified expression for the continuation value, we develop an estimator that is computationally attractive, easy to implement, and robust.

The rest of this paper proceeds as follows. In Section 2, we present our model and discuss identification of the model primitives, including the discount factor and utility levels. Section 3 discusses our estimation methodology and its step-by-step implementation. In Section 4, we present results from the Monte Carlo experiments to explore implications of our identification results for counterfactual predictions and also study the applicability and performance of our estimator. Section 5 presents an analysis of subprime mortgage defaults as an empirical illustration. Section 6 concludes.

2 Model

In this section, we set up a single-agent, finite-horizon, dynamic discrete choice model, in which one of the agent’s possible actions is terminal. For clarity of exposition, we specify our model by addressing mortgage borrowers’ default decisions, and we prove our identification results within the context of that model. However, our identification results hold more generally and do not rely on any specific feature of the mortgage setting. In our model, each agent is a borrower that enters a mortgage contract lasting T time periods, where T corresponds to the maturity date of the mortgage and is common across all agents.

2.1 Actions

At each time period t over the life of borrower i ’s loan, the borrower chooses an action $a_{i,t}$ from the finite set $A = \{0, 1, \dots, K\}$.³ For illustrative purposes, we focus on the case of a borrower choosing from three

³Note that we use t to denote the loan’s age, not calendar time.

possible actions $A = \{0, 1, 2\}$, but all the results and proofs below trivially extend to any finite K . The possible actions in A are to default ($a_{i,t} = 0$), to prepay (refinance) the mortgage ($a_{i,t} = 1$), or to make just the regularly scheduled monthly payment, which we refer to as “paying” ($a_{i,t} = 2$). We assume that there is no interaction among borrowers, so our setup is a single-agent model rather than a game. We assume that default is a terminating action: once a borrower defaults, there is no further decision to be made and no further flow of utility starting from the next period.⁴

2.2 Period Utility and State Transition

For notational simplicity, from now on we drop the index i for borrowers, except where necessary for disambiguation. We assume that, at loan age t , the period utility of a borrower taking an action a_t depends on a vector of state variables $s_t \in \mathcal{S}$, observed by both the borrower and the econometrician. $\mathcal{S}$ is an m -dimensional product space with either discrete or continuous subspaces. The borrower is also characterized by a time-dependent vector of idiosyncratic shocks associated with each action $\varepsilon_t = (\varepsilon_{0,t}, \varepsilon_{1,t}, \varepsilon_{2,t})$ (unobserved by the econometrician). Each element of ε_t is assumed to have a continuous support on the real line. We assume that the period utility takes an additively separable form:

$$\begin{aligned} U(a_t, s_t) &= u(a_t, s_t) + \varepsilon_{a_t,t}, & \text{for } t < T, \\ U(a_t, s_t) &= u_T(a_t, s_t) + \varepsilon_{a_t,t}, & \text{for } t = T. \end{aligned}$$

As specified, the period utility has a deterministic component, $u(\cdot, \cdot)$ for all periods $t < T$ and $u_T(\cdot, \cdot)$ in the final period. The state vector s_t may include the borrower’s characteristics, the current home value, monthly payments, etc. Some of these components are time-varying, but by assumption the period utility does not depend on the loan age t itself (or equivalently, the remaining time to maturity). We allow the final-period utility $u_T(\cdot, \cdot)$ to differ from all other periods. As motivation, consider that, in the context of mortgage default, the borrower obtains full ownership of the house once the mortgage is fully paid off at maturity, which we might think of as adding a lump-sum boost to the per-period utility in the final period. We make the following assumption regarding the marginal distributions of the random variables.

ASSUMPTION 1

- (i) *Independence of idiosyncratic payoff shocks: $s_t \perp \varepsilon_t$.*

⁴To be precise, we must assume that the lifetime value of choosing default does not depend on the sequence of optimal actions (of the considered model) to be taken in the future periods. The most straightforward case where this assumption would be satisfied is when agents make no further decisions once default is chosen in the current period. However, we need not interpret the terminating action as literally resulting in the agent not having to make any future decisions. Rather, the utility from default can be regarded as the “reduced-form” representation of the ex ante value function from a *separate* dynamic decision problem the borrower solves after foreclosure.

- (ii) *Conditional independence over time of idiosyncratic payoff shocks:* $\varepsilon_t \mid (\varepsilon_{t-1}, a_{t-1}) \sim \varepsilon_t \mid a_{t-1}$.
- (iii) *Markov transition of state variables:* s_t follows a reversible Markov process conditional on a_t , which is homogeneous with respect to time t .
- (iv) *Full support of idiosyncratic payoff shocks:* the distribution of the idiosyncratic shocks, $F_\varepsilon(\cdot)$, has a full support with the density strictly positive on $\mathbb{R}^3$.

Assumptions (i) and (ii) are standard assumptions in the discrete-choice literature. Assumption (iii) ensures that the density of the time-homogeneous Markov transition process is bounded away from zero everywhere on the support. Assumption (iv) ensures a positive probability for each of the actions given any realization of s_t . In our Monte Carlo experiments and empirical application we will use a conventional specification for the distribution of the idiosyncratic shocks by assuming that ε_t has a type I extreme value distribution (EVD), and is i.i.d. across agents and over time. This assumption is not essential, and we prove our identification results for an arbitrary continuous distribution of random shocks satisfying Assumption 1. The transition of some of the state variables may be influenced by the current action, a_t . We also allow the state variables potentially to follow a higher-order Markov process. This structure allows for greater realism, as certain important state variables may exhibit lag dependence.

The assumption of time homogeneity for the per-period utility is a critical assumption that allows us to exploit non-stationarity in the CCPs for identification.⁵ We believe that this assumption is sensible in many, but certainly not all, applications. For instance, in the classic example of married couples' contraceptive choice examined in Hotz and Miller (1993), the time homogeneity assumption implies that the period utility from choosing to become sterilized does not directly depend on a couple's age, conditioning on health status, income level, the stock of existing children, etc. Obviously the continuation values would vary over time, but it seems a reasonable approximation to think of the period utility from choosing sterilization—which presumably includes monetary costs as well as the physical and mental discomfort associated with the procedure—as being independent of age after conditioning on the above state vector.

The assumption also seems sensible in the case of purchasing a product with a deadline, such as flight tickets. The period utility of buying a ticket (which can be considered the terminating action) presumably depends on ticket costs as well as the utility from consuming the good (appropriately discounted to reflect the fact that consumption will occur in the future). This period utility does not however depend on when the ticket is purchased, conditional on the price of the ticket and other state variables. Likewise, it seems

⁵Alternatively, a functional form assumption on how the utilities depend on time t would also allow us to exploit non-stationarity in the CCPs for identification. We do not focus on this approach since we are interested in nonparametric identification.

reasonable to assume that the period utility of not purchasing a ticket is zero, regardless of the date and other state variables.

Other authors have also made the assumption of time homogeneity in contexts ranging from workers' retirement decisions (Rust and Phelan, 1997) to students' decisions to drop out from school (Eckstein and Wolpin, 1999) – in each case, the assumption is that the utility of leisure, consumption, and work/school is independent of age after conditioning on factors such as health, wealth, income, and grade.

2.3 Decision Rule and Value Function

We consider the borrower's problem as an optimal stopping problem, and assume that the default decision is irreversible—that is, the borrower cannot “resume” paying off the mortgage following a default. Provided that the default (“stopping”) decision is irreversible, the choice of the default option is equivalent to taking a one-time “payoff” without future utility flows. If the borrower pays or prepays (refinances) his mortgage, he receives the corresponding per-period payoff, plus the expected discounted stream of future utility. We assume that the borrower's inter-temporal preferences exhibit standard exponential discounting, where the parameter β is the single-period discount factor.

The borrower's decision rule D_t for each period t is a mapping from the vector of payoff-relevant variables into actions, $D_t : \mathcal{S} \times \mathbb{R}^3 \mapsto A$. We denote the borrower's decision probabilities by $\sigma_t(k|s_t) = E[\mathbf{1}\{D_t(s_t, \varepsilon_t) = k\} | s_t]$ for $k \in A$. We collect $\sigma_t(k|s_t)$ for all k and t such that $\sigma_t(s_t) = [\sigma_t(k=0|s_t), \sigma_t(k=1|s_t), \sigma_t(k=2|s_t)]$ and $\sigma = (\sigma_1(s_1), \dots, \sigma_T(s_T))$, and refer to σ as the policy function.

The *ex ante value function* at period $t < T$ is the expected discounted utility flow of a borrower who has not defaulted before t :

$$V_{t,\sigma}(s_t) = E_{\sigma,g(s)} \left[\sum_{\tau=t}^T \left(\beta^{\tau-t} U(a_\tau, s_\tau) \prod_{\tau_1=1}^{\tau-1} \mathbf{1}(a_{\tau_1} > 0) \right) | s_t \right],$$

where $g(s)$ represents the state transitions. The term $\prod_{\tau_1=1}^{\tau-1} \mathbf{1}(a_{\tau_1} > 0)$ reflects the fact that, once a borrower defaults, there is no further flow of utility starting from the next period.

The *choice-specific value function*, denoted by $V_{t,\sigma}(a_t = k, s_t)$, is the deterministic component of the borrower's discounted utility flow conditional on choosing action k in period t :

$$\begin{aligned} V_{t,\sigma}(a_t = k, s_t) &= u(a_t = k, s_t) + \beta E[V_{t+1,\sigma}(s_{t+1}) | s_t, a_t = k] \text{ for } t < T, \\ V_{t,\sigma}(a_t = k, s_t) &= u_T(a_t = k, s_t) \text{ for } t = T. \end{aligned}$$

In particular, the choice-specific value of default is equal to the per-period utility of default, i.e., $V_{t,\sigma}(a_t = 0, s_t) = u(a_t = 0, s_t)$ for $t < T$, because default is a terminating action whose continuation value $E[V_{t+1,\sigma}(s_{t+1})|s_t, a_t = 0]$ is zero. The existence and uniqueness of the borrower's optimal decision for the considered model is a standard result in the literature (e.g., Rust (1994) and references therein). For completeness, we provide the proposition and proof in the appendix.

2.4 Semiparametric Identification

In this section we demonstrate that our model is identified from objects observed in the data, namely, the choice probability of each option, conditional on the current state; and the transition distribution for the state variables, characterized by the conditional cdf $G(s_t | s_{t-1}, a_{t-1})$.

The model's three structural elements are: (1) the deterministic component of the per-period payoff function, $u(\cdot, \cdot)$ ⁶; (2) the time preference parameter β ; and (3) the distribution of the idiosyncratic utility shocks, which have cdf $F_\varepsilon(\cdot)$. We shall argue that $u(\cdot, \cdot)$ is nonparametrically identified and that the time preference parameter β is identified, for a given distribution of the idiosyncratic payoff shocks satisfying Assumption 1. We emphasize that our identification results do not rely on the extreme value assumption for the distribution of the idiosyncratic shocks.

We show the model is identified by demonstrating that there exists a unique mapping from the observable distribution of the data to the structural parameters. We start with the case in which the payoff from the default option is known, before relaxing this assumption.

Theorem 1 (*Identification with known default utility*) *Suppose that the payoff from the default option is a known function $u(0, \cdot)$. Also, suppose that for at least two consecutive periods t and $t + 1$, $\sigma_{k,t}(\cdot) \neq \sigma_{k,t+1}(\cdot)$ for $k \in A$. Given Assumption 1:*

- (i) *If the data distribution contains information on at least two consecutive periods and the discount factor β is known, the per-period utilities $u(1, s)$ and $u(2, s)$ are nonparametrically identified. Moreover, if $u(\cdot, \cdot) = u_T(\cdot, \cdot)$ and the observed periods include the final period T , then the discount factor is also identified.*
- (ii) *If the data distribution contains information on at least three consecutive periods (which do not necessarily include the final period T), both the discount factor and the per-period utility functions $u(1, s)$ and $u(2, s)$ are identified.*

⁶We omit discussing the identification of $u_T(\cdot, \cdot)$, as it is obvious that, if $u_T \neq u$, then $u_T(\cdot, \cdot)$ is identified if and only if decisions from the final period are observed.

Proof. Using the joint cdf of the idiosyncratic payoff shocks, we introduce the following functions:

$$\begin{aligned}\sigma_0(z_1, z_2) &= \int \mathbf{1}\{\varepsilon_0 \geq z_1 + \varepsilon_1, \varepsilon_0 \geq z_2 + \varepsilon_2\} F_\varepsilon(d\varepsilon), \\ \sigma_1(z_1, z_2) &= \int \mathbf{1}\{z_1 + \varepsilon_1 \geq \varepsilon_0, z_1 + \varepsilon_1 \geq z_2 + \varepsilon_2\} F_\varepsilon(d\varepsilon), \\ \sigma_2(z_1, z_2) &= \int \mathbf{1}\{z_2 + \varepsilon_2 \geq \varepsilon_0, z_2 + \varepsilon_2 \geq z_1 + \varepsilon_1\} F_\varepsilon(d\varepsilon).\end{aligned}\tag{1}$$

We note that the introduced functions are known (given that we can normalize the distribution of the idiosyncratic shocks) and are monotone and differentiable in their arguments. Using Lemma 1 in the appendix, which establishes that *the system of equations*

$$\begin{aligned}\sigma_0(z_1, z_2) &= \bar{\sigma}_0, \\ \sigma_1(z_1, z_2) &= \bar{\sigma}_1\end{aligned}$$

has a unique solution if and only if $\bar{\sigma}_0 + \bar{\sigma}_1 < 1$, we can show that the model is nonparametrically identified. We introduce the function

$$\begin{aligned}\nu(z_1, z_2) &= \int \left(z_1 \mathbf{1}\{z_1 + \varepsilon_1 \geq \varepsilon_0, z_1 + \varepsilon_1 \geq z_2 + \varepsilon_2\} \right. \\ &\quad \left. + z_2 \mathbf{1}\{z_2 + \varepsilon_2 \geq \varepsilon_0, z_2 + \varepsilon_2 \geq z_1 + \varepsilon_1\} \right) F_\varepsilon(d\varepsilon) + u(0),\end{aligned}$$

where the payoff from the default option $u(0)$ is a fixed known function according to the assumption of the theorem.

The observed probability distribution characterizes the CCPs $\{\sigma_{k,t}(\cdot), k \in A\}$. Given the structure of the optimal solution, there is a direct link between the choice-specific value functions in period t and the choice probability which is expressed through the distribution of the idiosyncratic payoff shocks. In particular, for each $s \in \mathcal{S}$ and each $t \leq T$, we can write the system of identifying equations

$$\begin{aligned}\sigma_0(V_t(1, s) - u(0, s), V_t(2, s) - u(0, s)) &= \sigma_{0,t}(s), \\ \sigma_1(V_t(1, s) - u(0, s), V_t(2, s) - u(0, s)) &= \sigma_{1,t}(s).\end{aligned}$$

Given Lemma 1, we can solve for the choice-specific value functions $V_t(1, s)$ and $V_t(2, s)$ over $\mathcal{S}$.

The conditional distribution $s_{t+1} | s_t$, a_t is observable. As a result, for each $k \in \{1, 2\}$ and each s we

can consider the system of equations in two consecutive periods t and $t + 1 < T$:

$$\begin{aligned} V_t(1, s) &= u(1, s) + \beta E [\nu (V_{t+1}(1, s') - u(0, s'), V_{t+1}(2, s') - u(0, s')) \mid a_t = 1, s], \\ V_t(2, s) &= u(2, s) + \beta E [\nu (V_{t+1}(1, s') - u(0, s'), V_{t+1}(2, s') - u(0, s')) \mid a_t = 2, s], \\ V_{t+1}(1, s) &= u(1, s) + \beta E [\nu (V_{t+2}(1, s') - u(0, s'), V_{t+2}(2, s') - u(0, s')) \mid a_{t+1} = 1, s]. \end{aligned}$$

This is a system of three linear equations with three unknowns $u(1, s)$, $u(2, s)$ and β . We note that to set up this system of equations, we need to have observations for at least three consecutive periods.

Then simple differencing solves for the discount factor:

$$\begin{aligned} \beta &= \left(E [\nu (V_{t+1}(1, s') - u(0, s'), V_{t+1}(2, s') - u(0, s')) \mid a_t = 1, s] \right. \\ &\quad \left. - E [\nu (V_{t+2}(1, s') - u(0, s'), V_{t+2}(2, s') - u(0, s')) \mid a_{t+1} = 1, s] \right)^{-1} \\ &\quad \times [V_t(1, s) - V_{t+1}(1, s)]. \end{aligned}$$

The denominator in this expression is not equal to zero because of the assumption of the theorem that $\sigma_{k,t}(s) \neq \sigma_{k,t+1}(s)$ for at least two consecutive periods t and $t + 1$. We also can recover the per-period utility function as

$$\begin{aligned} u(1, s) &= \left(E [\nu (V_{t+1}(1, s') - u(0, s'), V_{t+1}(2, s') - u(0, s')) \mid a_t = 1, s] \right. \\ &\quad \left. - E [\nu (V_{t+2}(1, s') - u(0, s'), V_{t+2}(2, s') - u(0, s')) \mid a_{t+1} = 1, s] \right)^{-1} \\ &\quad \times \left(V_{t+1}(1, s) E [\nu (V_{t+1}(1, s') - u(0, s'), V_{t+1}(2, s') - u(0, s')) \mid a_t = 1, s] \right. \\ &\quad \left. - V_t(1, s) E [\nu (V_{t+2}(1, s') - u(0, s'), V_{t+2}(2, s') - u(0, s')) \mid a_{t+1} = 1, s] \right). \end{aligned}$$

Similarly, we can explicitly recover the utility $u(2, s)$ for any s in the support. We note that if the discount factor is fixed, then we can identify the utility function from just two periods.

If the per-period payoff function in the final period coincides with the per-period payoff function in the previous periods ($u(\cdot, \cdot) = u_T(\cdot, \cdot)$), then in the final period the choice specific values coincide with the utilities, meaning that $V_T(1, s) = u(1, s)$ and $V_T(2, s) = u(2, s)$. Both these utilities are recovered from just observing the final-period choice probabilities. Then we can re-construct the ex ante value function

of the last period using these estimates. We can take the first equation of the system considered before

$$V_{T-1}(1, s) = u(1, s) + \beta E [\nu(u(1, s') - u(0, s'), u(2, s') - u(0, s')) \mid a_{T-1} = 1, s],$$

where the only remaining unknown is the discount factor. We can recover it as

$$\beta = \frac{V_{T-1}(1, s) - u(1, s)}{E [\nu(u(1, s') - u(0, s'), u(2, s') - u(0, s')) \mid a_{T-1} = 1, s]}.$$

Therefore, in the special case where $u = u_T$, the discount factor as well as the utility functions are identified from just two consecutive periods $T - 1$ and T . ■

This theorem establishes the general result that the considered model is identified (including identification of the time preference parameter) if the payoff from default is a known function. The argument requires the presence of two consecutive time periods over which there is variation in the optimal decision probability, conditional on the state variables. The theoretical justification for why we would expect two such periods to exist stems from the finite horizon, which makes borrowers' default decision dependent on the time remaining until the mortgage maturity date. More generally, the optimal decision rules depend on time t in finite horizon models, even after conditioning on the state variables, thereby satisfying the assumption in the theorem.

The key feature of our identification results is that they exploit the time variation in the optimal choice probabilities that is inherent in finite horizon models. The presence of a finite horizon, combined with the assumed time-homogeneity of the per-period payoff, implies that the time to maturity affects the continuation value without affecting the per-period utility directly. Therefore, we do not need an additional variable to satisfy the exclusion restriction, as in Fang and Wang (2015). Also, our identification does not rely on the availability of observations on the final period. As part (ii) of Theorem 1 states, the discount factor is identified as long as researchers observe data from at least three consecutive periods, even if those periods do not include the final period.

When the default utility is not a known function, we would in general need to normalize it by fixing it to some known function. In the empirical literature on dynamic discrete optimization problems, it has been noted (e.g., see Bajari, Hong and Nekipelov, 2013; Norets and Tang, 2014) that the choice of normalization for the per-period utility associated with one of the choices is not innocuous as it can affect counterfactual predictions. In Theorem 2 below, we formally demonstrate this to be the case in the finite horizon optimal stopping problem.⁷ Moreover, we show that, under stronger requirements on the data,

⁷We are grateful to Günter Hitsch, who encouraged us to present the formal argument supporting this statement.

the structural model is overidentified for a given normalization, such that one can identify the payoffs from all options without the need to normalize any payoff. These two insights can be used to explore the identification of the elements of the structural model, including the payoff from default.

Theorem 2 (*Identification with unknown default utility*) Suppose that for at least two consecutive periods t and $t + 1$, $\sigma_{k,t}(\cdot) \neq \sigma_{k,t+1}(\cdot)$ for $k \in A$. Given Assumption 1:

- (i) If $u(\cdot, \cdot) \neq u_T(\cdot, \cdot)$ or the choices of the borrowers in the final period are not observed, one cannot identify the utilities from all choices, $u(0, s)$, $u(1, s)$, and $u(2, s)$. However, if $u(0, s)$ is normalized to a fixed function, Theorem 1 (ii) applies. In this case, the choice of normalization does not affect the recovered discount factor, but it does affect the recovered differences between the per-period payoffs from payment or prepayment and the per-period payoff from default (i.e., $u(1, s) - u(0, s)$ and $u(2, s) - u(0, s)$).
- (ii) Suppose that $u(\cdot, \cdot) = u_T(\cdot, \cdot)$ and that the choices of the borrowers in the final period T are observed along with the choices from earlier periods. If the data distribution contains information on at least three consecutive periods, then the utilities from all choices, $u(0, s)$, $u(1, s)$, and $u(2, s)$, are identified along with the discount factor β .

Proof. For a particular normalization of $u(0, s)$, the following relationship holds:

$$\sigma_{0,t}(s) = \int \mathbf{1}\{u(0, s) + \varepsilon_{0,t} \geq V_t(1, s) + \varepsilon_{1,t}, u(0, s) + \varepsilon_{0,t} \geq V_t(2, s) + \varepsilon_{2,t}\} F_\varepsilon(d\varepsilon).$$

Using the notation in the proof of the previous theorem, we can then express

$$\sigma_{0,t}(s) = \sigma_0 (V_t(1, s) - u(0, s), V_t(2, s) - u(0, s)).$$

We can also provide similar expressions for other choices. Then we can solve the system of equations

$$\begin{aligned} \sigma_{0,t}(s) &= \sigma_0 (V_t(1, s) - u(0, s), V_t(2, s) - u(0, s)), \\ \sigma_{1,t}(s) &= \sigma_1 (V_t(1, s) - u(0, s), V_t(2, s) - u(0, s)) \end{aligned}$$

to recover the choice specific value functions. We note that the recovered differences $V_t(k, s) - u(0, s)$ are invariant to the choice of $u(0, s)$ as they are directly recovered from the data.

The ex ante value function can be recovered as

$$\begin{aligned} V_t(s) &= \int \left((V_t(1, s) - u(0, s)) \mathbf{1}\{V_t(1, s) - u(0, s) + \varepsilon_{1,t} \geq \varepsilon_{0,t}, V_t(1, s) + \varepsilon_{1,t} \right. \\ &\quad \left. \geq V_t(2, s) + \varepsilon_{2,t}\} \right. \\ &\quad + (V_t(2, s) - u(0, s)) \mathbf{1}\{V_t(2, s) - u(0, s) + \varepsilon_{2,t} \geq \varepsilon_{0,t}, V_t(2, s) + \varepsilon_{2,t} \\ &\quad \left. \geq V_t(1, s) + \varepsilon_{1,t}\} \right) F_\varepsilon(d\varepsilon) \\ &\quad + u(0, s) = \mathbf{V}_t(s) + u(0, s), \end{aligned}$$

where the function $\mathbf{V}_t(\cdot)$ is invariant to the choice of the default payoff (being determined by the observed choice probabilities). The system of equations identifying the per-period payoffs takes the form

$$\begin{aligned} V_t(1, s) &= u(1, s) + \beta E \left[\mathbf{V}_{t+1}(s') + u(0, s') \mid a_t = 1, s \right], \\ V_t(2, s) &= u(2, s) + \beta E \left[\mathbf{V}_{t+1}(s') + u(0, s') \mid a_t = 2, s \right], \\ V_{t+1}(1, s) &= u(1, s) + \beta E \left[\mathbf{V}_{t+2}(s') + u(0, s') \mid a_{t+1} = 1, s \right]. \end{aligned}$$

Provided that $u(0, s)$ is normalized to a fixed function, taking the difference between the first and the third equation allows us to express the discount factor as

$$\beta = \frac{V_t(1, s) - u(0, s) - (V_{t+1}(1, s) - u(0, s))}{E[\mathbf{V}_{t+1}(s') \mid a_t = 1, s] - E[\mathbf{V}_{t+2}(s') \mid a_{t+1} = 1, s]}. \quad (2)$$

Provided that the numerator and the denominator are invariant with respect to the choice of the default utility (as they can be directly recovered from the observed choice probabilities), the discount factor will not depend on that choice either. Having obtained the discount factor, we can focus on the first two equations of the considered system and re-cast them to the form

$$\begin{aligned} V_t(1, s) - u(0, s) - \beta E \left[\mathbf{V}_{t+1}(s') \mid a_t = 1, s \right] &= u(1, s) - u(0, s) + \beta E \left[u(0, s') \mid a_t = 1, s \right], \\ V_t(2, s) - u(0, s) - \beta E \left[\mathbf{V}_{t+1}(s') \mid a_t = 2, s \right] &= u(2, s) - u(0, s) + \beta E \left[u(0, s') \mid a_t = 2, s \right]. \end{aligned}$$

The left-hand side of this system is invariant to the choice of the default utility, implying that the right-hand side is as well. In other words, for any two choices $u(0, s)$ and $u'(0, s)$,

$$u(k, s) - u(0, s) + \beta E[u(0, s') \mid a_t = k, s] = u'(k, s) - u'(0, s) + \beta E[u'(0, s') \mid a_t = k, s].$$

Thus, the gap between the utilities from different options depends on the choice of the default utility, as follows:

$$u(k, s) - u(0, s) - (u'(k, s) - u'(0, s)) = \beta E [u(0, s') - u(0, s) \mid a_t = k, s]$$

for $k = 1, 2$.

If the final-period choice probability is observed and $u = u_T$, then we can complement the system of Bellman equations with the expressions for the choice probabilities of the final period (when the problem becomes static). Those choice probabilities identify the differences $u(1, s) - u(0, s)$ and $u(2, s) - u(0, s)$. Thus, we can identify the pair of conditional expectations $E [u(0, s') \mid a_t = 1, s]$ and $E [u(0, s') \mid a_t = 2, s]$ from the system of linear equations that we constructed above:

$$\beta E [u(0, s') \mid a_t = 1, s] = V_t(1, s) - u(0, s) - \beta E [\mathbf{V}_{t+1}(s') \mid a_t = 1, s] - u(1, s) + u(0, s), \quad (3)$$

$$\beta E [u(0, s') \mid a_t = 2, s] = V_t(2, s) - u(0, s) - \beta E [\mathbf{V}_{t+1}(s') \mid a_t = 2, s] - u(2, s) + u(0, s).$$

When the per-period utility functions are not specified parametrically, this system of equations represents a nonparametric instrumental variable problem.⁸ The unknown function on the left-hand-side, $u(0, s)$, is determined by the mean independence condition between the right-hand side terms and a set of instruments consisting of the choice k and state variables s . It has a unique solution and, thus, the utility from default is identified, under the *completeness* condition that is discussed in Newey and Powell (2003), and Chen, Chernozhukov, Lee and Newey (2014). In the simple case where the true default utility is constant, i.e., $u(0, s) = u(0)$, identification comes trivially from the above equation. In parametric settings, the completeness condition reduces to the standard rank condition. If the default utility is further a linear function of the state variables, the rank condition would be satisfied if the autocovariance matrix $E[s_{t+1}s'_t]$ is non-degenerate, i.e., the process of the state transition should be “sufficiently rich.” Therefore, the utilities from all choices $u(0, s)$, $u(1, s)$, and $u(2, s)$ are identified along with the discount factor β in this case. ■

Part (i) of the theorem holds because, whereas only the utility in the current period is shifted by the normalization in the case of default (as the future discounted payoff is zero), the payoffs from payment or prepayment are additionally shifted by an amount equal to the discounted expected payoff from defaulting in a future period. However, β is invariant to the normalization because the tradeoff between current payoffs and future option values is unaffected by the normalization. Importantly, the theorem implies

⁸When the per-period utilities are parametric functions, this problem reduces to a simpler nonlinear instrumental variable setup. Although in our model the per-period utilities can be nonparametrically identified, we chose to adhere to a semiparametric model (with parametrically specified utilities) for inference in the econometric section as well as in our Monte Carlo and application. We believe the semiparametric model to be more common in applied work.

that identification of β does not require knowledge of the default utility, as long as a researcher has data on any three consecutive periods.

Part (ii) of the theorem indicates that observing the final-period choices allows us to identify the actual levels of the payoffs and not just the differences in level, provided the final-period utility function is the same as in the previous periods. The intuition is as follows. In the final period, the borrower’s decision is static and thus there is no option value of delaying default to a future period. Rather, the final-period decision depends only on the differences between the period utilities from payment and prepayment relative to default. By contrast, in any period before the last, the discounted payoff streams from the non-terminal actions (paying or prepaying) additionally include an option value of being able to default in a future period. This asymmetry between paying or prepaying versus default implies that, in all periods other than the last, the normalization of $u(0, s)$ disproportionately affects the payoffs to the non-terminal actions. Therefore, $u(0, s)$ is identified by variation in the probabilities of the non-terminal actions relative to default—both across periods (especially between the final period and earlier periods) and with respect to the state variables.

To the best of our knowledge, this result is novel, as the existing literature has either normalized the per-period payoff to one of the actions, or used additional data to pin down this object (e.g., Heckman and Navarro, 2007; Kalouptsi, 2014). Prior work (Aguirregabiria, 2010; Aguirregabiria and Suzuki, 2014; Norets and Tang, 2014; Kalouptsi et al., 2016) showed that certain counterfactual CCPs are not identified when only the differences in utility, but not the levels, are known. Therefore, this theorem has significant implications for researchers’ ability to conduct counterfactual analysis, which is often the key object of interest for applied researchers estimating structural models.

The theorems require time-homogeneity of the utility functions. However, it is worth emphasizing that we allow $u(\cdot, \cdot)$ to differ from $u_T(\cdot, \cdot)$, except for the cases covered by Theorem 1 (i) and Theorem 2 (ii). In other words, we do not require the time-homogeneity assumption on the utility functions to extend to the final period, unless we are interested in identification of $u(0, s)$ or want to identify β just based on data from the two periods $T - 1$ and T .

3 Econometric Methodology

For convenience, we reexpress the per-period utility as

$$u(a, s) = u(s; \theta(a)),$$

where $a \in A$ and $\theta : A \mapsto \Theta$, with Θ denoting the parameter space. For inference, we focus on the semiparametric setup where the utility functions are parametrically specified.⁹ We propose a plug-in semiparametric estimator that does not rely on backward induction or availability of data on the final periods. Parallel to our identification argument, we nonparametrically estimate the borrowers' policy functions and use them to recover the choice-specific value functions of the borrowers. Using those, we then recover the parameters in the utility function and the time preference parameter. Below, we first characterize the general form of the estimator corresponding to an arbitrary distribution of idiosyncratic payoff shocks satisfying Assumption 1. Then, we discuss a specific implementation with idiosyncratic payoff shocks that are distributed type I extreme value, which we expect to be the most common specification for applied researchers. For this case, estimation reduces to evaluating several linear projections, which does not require costly computations or simulations, and can be implemented using any standard statistical software. Estimation for more general shock distributions may necessitate using more advanced computational tools.

3.1 Nonparametric Estimation of CCPs

First, we nonparametrically estimate the CCPs of the borrowers. Suppose the data represent a panel of loans $i = 1, \dots, J$ observed in periods $\tau = 1, \dots, T^*$, with the latter indices denoting calendar time. The panel is unbalanced due to defaults and issuance of new loans. We use $T_{i,\tau}$ to denote the time elapsed from the period of mortgage origination for a loan observed in period τ . We estimate the policy functions by evaluating the conditional distribution of observed actions at each observed loan age t . It is important to recover policy functions separately for each t in order to take into account the non-stationarity of agents' optimal behavior due to the presence of a final period.

Our estimation procedure uses a projection on orthogonal polynomials $q^L(s) = (q_1(s), \dots, q_L(s))'$, where L is the highest degree polynomial used in the projection. The components of $q_l(s)$ are vectors corresponding to the tensor products of orthogonal series. Typical choices of orthogonal polynomials in cases where the support of state variables is unbounded include Hermite polynomials. In applied work in natural sciences the researchers have successfully used high-dimensional basis functions such as Bessel functions or spherical functions. However, they have not been sufficiently explored in application to Marketing or Economic problems. In the appendix we provide technical conditions that lead to the choice of the highest degree of polynomials as a function of the sample size. Since these conditions are usually difficult to verify, a possible approach includes cross-validation. In that case, we take two random

⁹Under additional technical conditions, our estimation procedure can be extended to the case where utility functions are nonparametrically specified.

subsamples of the data and estimate the model with a varying power L using one of the subsamples. Then we compute the fit of the predicted values using the second subsample. The choice of L corresponds to the best fit of the predictions. We note that once the polynomial degree L is chosen, the estimation procedure becomes identical to the standard parametric GMM or, with linear specification of utility functions, to the standard linear IV. We consider the orthogonal representation of the choice probability as

$$\log \frac{\sigma_t(k, s)}{\sigma_t(0, s)} = \sum_{l=1}^{\infty} r_l(t, k) q_l(s), \text{ for } k = 1, 2,$$

where $r_l(t, k)$ are coefficients of the series representation. This representation will provide a uniform approximation for the decision rule if the choice probability ratio is continuous in s and the state space $\mathcal{S}$ is a compact set, by the Weierstrass theorem. We construct our estimator by replacing the infinite sum with a finite sum for some (large) number L . The parameters are estimated by forming a quasi-likelihood:

$$\begin{aligned} \hat{Q}(r^L(t, 1), r^L(t, 2)) = & \frac{1}{J} \sum_{\tau=1}^{T^*} \sum_{i=1}^J \mathbf{1}\{T_{i,\tau} = t\} \left[\mathbf{1}\{a_{i,t} = 1\} \sum_{l=1}^L r_l(t, 1) q_l(s_{i,t}) \right. \\ & + \mathbf{1}\{a_{i,t} = 2\} \sum_{l=1}^L r_l(t, 2) q_l(s_{i,t}) \\ & \left. - \log \left(1 + \exp \left(\sum_{l=1}^L r_l(t, 1) q_l(s_{i,t}) \right) + \exp \left(\sum_{l=1}^L r_l(t, 2) q_l(s_{i,t}) \right) \right) \right], \end{aligned}$$

where $r^L(t, k) = (r_1(t, k), \dots, r_L(t, k))$. Then, we obtain the estimator as

$$(\hat{r}^L(t, 1), \hat{r}^L(t, 2)) = \operatorname{argmax}_{r^L} \hat{Q}(r^L(t, 1), r^L(t, 2)). \quad (4)$$

The estimated policy functions correspond to the fitted values based on the estimated parameters:

$$\begin{aligned} \hat{\sigma}_t(k, s) = & \frac{\exp \left(\sum_{l=1}^L \hat{r}_l(t, k) q_l(s) \right)}{1 + \exp \left(\sum_{l=1}^L \hat{r}_l(t, 1) q_l(s) \right) + \exp \left(\sum_{l=1}^L \hat{r}_l(t, 2) q_l(s) \right)} \text{ for } k = 1, 2, \\ \hat{\sigma}_t(0, s) = & 1 - \hat{\sigma}_t(1, s) - \hat{\sigma}_t(2, s) \end{aligned}$$

The number of series terms is a function of the total sample size, with $L \rightarrow \infty$ as $J \rightarrow \infty$. In the appendix we give concrete technical conditions that ensure the consistency of our estimates as well as requirements for the choices of the number of basis functions and its interaction with the sample size. As we show in the appendix, for our asymptotic distribution results to be valid (and, thus, for the first-stage estimation

error to have no impact on the convergence rate for the estimated structural parameters), it is sufficient to find an estimator for the choice probabilities with a uniform convergence rate of at least $J^{-1/4}$. Such estimators exist if the choice probability is a sufficiently smooth function of the state.

3.2 Estimation of Structural Parameters

In the case of an arbitrary (known) distribution of idiosyncratic shocks, we must perform a functional inversion to recover the choice-specific and ex ante value functions from the estimated CCPs. As Lemma 1 in the appendix states, the functions in (1) are well-defined, and if the default probability is strictly between zero and one for almost all values of the state variables, the system determining the choice probabilities is everywhere invertible for $z_1, z_2 \in \mathbb{R}$, assuming a large support for the idiosyncratic payoff shocks (Hotz and Miller, 1993; Norets and Takahashi, 2013).

Thus we can use the system of equations in (1) to recover the difference between the choice-specific value of each non-default choice and the utility from default (captured by z_1 and z_2 in the expression), for each period t and each value of the state variables s . Specifically, we reexpress the choice probabilities as functions of these differences and equate them to their empirical analogues (which were recovered in the first step). This yields

$$\begin{aligned}\hat{\sigma}_{0,t}(s) &= \sigma_0(V_t(1, s) - u(s; \theta(0)), V_t(2, s) - u(s; \theta(0))), \\ \hat{\sigma}_{1,t}(s) &= \sigma_1(V_t(1, s) - u(s; \theta(0)), V_t(2, s) - u(s; \theta(0))).\end{aligned}$$

The solution to the above system can be expressed as:

$$\hat{V}_t(k, s) = u(s; \theta(0)) + F_k(\hat{\sigma}_{0,t}(s), \hat{\sigma}_{1,t}(s)), \quad k = 1, 2,$$

where the function $F_k(\cdot, \cdot)$ is the solution of the inversion problem (1) for the argument z_k . We can characterize the ex ante value function as

$$\hat{V}_t(s) = u(s; \theta(0)) + F(\hat{\sigma}_{0,t}(s), \hat{\sigma}_{1,t}(s)),$$

where the function $F(\cdot, \cdot)$ is determined by the solutions $F_1(\cdot, \cdot)$ and $F_2(\cdot, \cdot)$ and the distribution of the idiosyncratic payoff shocks.

We then substitute the obtained expressions into the Bellman equation for the borrower, obtaining a

system of nonlinear simultaneous equations

$$E \left[F_k(\hat{\sigma}_{0,t}(s_t), \hat{\sigma}_{1,t}(s_t)) - u(s_t; \theta(k)) + u(s_t; \theta(0)) - \beta u(s_{t+1}; \theta(0)) \right. \\ \left. - \beta F(\hat{\sigma}_{0,t+1}(s_{t+1}), \hat{\sigma}_{1,t+1}(s_{t+1})) \mid a_t = k, s_t \right] = 0, \quad k = 1, 2. \quad (5)$$

We can estimate this system using nonlinear IV methodology, with functions of the current state s_t and current action a_t as the instruments. Specifically, we define the instruments as $Z_t = \{W_m(s_t; a_t), m \in \mathcal{M}\}$ for some set of functions indexed by $\mathcal{M}$. For example, these could be a finite list of orthogonal polynomials of the state variables. For nonlinear IV estimation, it is convenient to replace the system of conditional moment equations with a system of unconditional moment equations. Define:

$$\epsilon_{kt} = F_k(\hat{\sigma}_{0,t}(s_t), \hat{\sigma}_{1,t}(s_t)) - u(s_t; \theta(k)) + u(s_t; \theta(0)) - \beta u(s_{t+1}; \theta(0)) \\ - \beta F(\hat{\sigma}_{0,t+1}(s_{t+1}), \hat{\sigma}_{1,t+1}(s_{t+1})), \quad k = 1, 2.$$

Note that ϵ_{kt} can be constructed by plugging in the realized values of the state variables s_t and s_{t+1} and plugging in the expressions $F_k(\cdot, \cdot)$ and $F(\cdot, \cdot)$ evaluated at the recovered choice probabilities $\hat{\sigma}_{0,t}(s_t)$, $\hat{\sigma}_{1,t}(s_t)$, $\hat{\sigma}_{0,t+1}(s_{t+1})$, and $\hat{\sigma}_{1,t+1}(s_{t+1})$ at the realized values of s_t and s_{t+1} . Define $\epsilon_t = (\epsilon_{1t}, \epsilon_{2t})'$ and $Z_t = (Z_{1t}, Z_{2t})$. Then the system of unconditional moment equations takes the familiar form

$$E[\epsilon_t Z_t'] = 0. \quad (6)$$

Provided that the model we analyze is smooth with respect to the nonparametrically estimated components, such as choice probabilities, we can apply the results in Chen, Linton, and van Keilegom (2003) and Mammen, Rothe, and Schienle (2012) to establish the impact of the first-stage estimation error on the standard errors of the structural estimates. In Appendices 3 and 4, we analyze the properties of this two-step estimator for the case of a general distribution of the idiosyncratic payoff shocks.¹⁰

Alternatively, the system (5) can be estimated as a standard regression by treating F_k as the “outcome variable” and the rest of the terms as right-hand side variables with attached parameters to be estimated (i.e., the unknown utilities parametrized by $\theta(0)$ and $\theta(k)$ and the time preference parameter β). Under this approach, one would need to construct $E \left[F(\hat{\sigma}_{0,t+1}(s_{t+1}), \hat{\sigma}_{1,t+1}(s_{t+1})) \mid a_t = k, s_t \right]$ so that it can

¹⁰While our focus in the paper is on a single-agent finite-horizon dynamic discrete choice model with a terminating action, we can extend our estimation framework straightforwardly to multi-agent settings (assuming identification, which requires additional conditions for games compared to single agent models) as well as the finite dependence problems. For instance, for an extension to the finite dependence problems, we have to increase the number of terms that need to be included in the second stage IV regression, because the presence of the finite dependence makes the choice specific value function be determined by sequence of discounted future utilities in the “dependence window” as well as the expected ex ante value function.

be plugged in as one of the regressors. In the prior literature, this construction is typically accomplished via forward simulation based on estimates of the state transition process for $s_{t+1}|s_t, a_t$. We instead propose to construct the expression using a nonparametric regression. Nonparametric estimation of the conditional expectation of the continuation value function can be performed using a variety of standard methods, such as kernel smoothers, sieves or orthogonal polynomials. A rich literature that includes Härdle (1990), Ullah and Pagan (1999) and Chen (2007) exhaustively describes the implementation of each of these methods.

One feature that aids implementation of our estimator is the smoothness of the above conditional expectation with respect to the continuous components of the state variables.¹¹ In practice, implementation of the nonparametric estimator requires trading off bias and variance, and reducing bias (e.g., through choosing a higher degree for the approximating polynomials) can make the small-sample behavior of the estimator unstable. Appendix 3 formalizes the asymptotic properties of our nonparametric estimator. Our convergence result (Theorem 3 in the appendix) merely states conditions for the estimator to converge uniformly at the rate faster than $J^{-1/4}$ and thus does not require a high degree of bias reduction for a given sample size.¹²

3.3 OLS Estimation under Type I EVD and Linear Utility Function

This subsection describes, step-by-step, how to implement our proposed estimator for the case of a linear utility function and type I extreme-value-distributed idiosyncratic payoff shocks, which is a common specification in applied research.

When the idiosyncratic payoff shocks have an i.i.d. type I extreme value distribution, the elements of the derived system of conditional moment equations can be expressed in closed form:

$$F(\sigma_0(s), \sigma_1(s)) = \log(\sigma_0(s)^{-1}).$$

¹¹In our Monte Carlo exercises, we specify the per period utility function to be a linear function of the state variables, and assume a type I extreme value distribution for the idiosyncratic payoff shocks. Because the type I extreme value distribution has an infinitely differentiable density, the continuation value function in period $T-1$ is infinitely differentiable. Continuing the backward induction, we can establish that, in the context of our specification, all continuation value functions are infinitely differentiable. This implies that the (optimal) convergence rate of the nonparametric estimator approaches the parametric convergence rate.

¹²In particular if we use a degree K polynomial to estimate the nonparametric regression, then the bias corresponds to the magnitude of the residual from numerical approximation of the estimated function via such a polynomial. For infinitely differentiable functions, this residual can be majorized by a power function of the degree K , e.g., c^{-K} , where c depends on the support of the state space (see Judd, 1998). At the same time, the variance of the estimator has the magnitude $O(K/J)$. This permits choosing the degree of the approximating polynomials to be up to $K \propto J/\log J$. In practice, this choice can be made using cross-validation.

And the functions $F_1(\cdot)$ and $F_2(\cdot)$ in this case take the form:

$$F_k(\sigma_0(s), \sigma_1(s)) = \log(\sigma_k(s)/\sigma_0(s)), \text{ for } k = 1, 2.$$

The choice-specific and ex ante value functions in this case can be recovered directly from the estimated choice probabilities (up to the normalization of $u(s; \theta(0))$), without requiring any iterations or complicated inversions.

If we assume, further, that the choice utilities are linearly parametrized, one possible estimation method is OLS. Although OLS is not semiparametrically efficient, it has the advantage of being simple to estimate and intuitive to understand. In this case, the entire estimation procedure reduces to three simple regressions.

1. We nonparametrically estimate the borrower's choice probabilities $\sigma_{k,t}(\cdot)$. This can be done in standard statistical software such as R, STATA or MATLAB by running a multinomial logit regression separately for each loan age t , using orthogonal polynomials of the state variables as basis functions. We perform this estimation separately for each t in order to account for the non-stationarity of the problem. The fitted values from this regression form the estimated CCPs. Using the estimated CCPs, we recover the choice specific value function ($\log(\hat{\sigma}_{k,t}(s_t)/\hat{\sigma}_{0,t}(s_t))$) and ex ante value function ($F_t = \log(\hat{\sigma}_{0,t}(s_t)^{-1})$) up to the normalization of $u(s; \theta(0))$ for each borrower in each time period.
2. We estimate the continuation value conditional on the current state and action, $E[\log(\hat{\sigma}_{0,t+t}(s_{t+1})^{-1}) | a_t, s_t]$, through nonparametric estimation. For instance, the estimation can be performed by running a linear regression of $\log(\hat{\sigma}_{0,t+1}(s_{t+1})^{-1})$ (the fitted value from the previous step) on orthogonal polynomials of the state variables, separately for each combination of a_t and t . The fitted values from this regression form the variable

$$\hat{E}[\log(\hat{\sigma}_{0,t+t}(s_{t+1})^{-1}) | a_t = k, s_t] \text{ for } k = 1, 2,$$

which we construct for each borrower in each time period. Similarly, we construct $\hat{E}[s_{t+1}(0) | s_t, a_t = k]$ for $k = 1, 2$ by nonparametrically regressing the realized values of the state variables of default utility in period $t + 1$ ($s_{t+1}(0)$) on s_t separately for each k and computing the fitted values from the regression.

3. We estimate the structural parameters using OLS, using the continuation values constructed in the

previous step as a regressor. The recovered choice specific value function $\log(\hat{\sigma}_{k,t}(s_t)/\hat{\sigma}_{0,t}(s_t))$ for $k = 1, 2$ provides the “outcome” variables $Y_{k,t}$. Because the per-period utility is represented by a linear index of the state variables, the vector of “regressors” comprises:

$$X_{k,t} = \left(s_t(k), -s_t(0), \hat{E}[s_{t+1}(0)|s_t, a_t = k], \hat{E}[\log(\hat{\sigma}_{0,t+t}(s_{t+1})^{-1}) | a_t = k, s_t] \right)'.$$

The components of this vector correspond to $u(s_t; \theta(k))$, $-u(s_t; \theta(0))$, $u(s_{t+1}; \theta(0))$ and $F(\hat{\sigma}_{0,t+1}(s_{t+1}), \hat{\sigma}_{1,t+1}(s_{t+1}))$ in (5), respectively. Estimating the coefficient δ_k in the regression

$$Y_{k,t} = \delta_k' X_{k,t} + \epsilon_{k,t} \quad (7)$$

yields the structural parameters of interest. In fact, by construction of Y and X , we get $\delta_k = (\theta(k), \theta(0), \theta(0) \cdot \beta, \beta)'$. This estimation can be performed using nonlinear least squares procedures from standard statistical software, imposing the constraint that the parameter on the third term, $\theta(0) \cdot \beta$, is the product of the parameters on the second ($\theta(0)$) and fourth terms (β). We can improve inference by simultaneously estimating the equations for $k = 1$ and 2 as a system. Depending on the normalization we choose for $u(s; \theta(0))$, the estimation can be further simplified. For instance, if we normalize $u(s; \theta(0))$ to a constant, which is a common choice in empirical work, we only need to run a linear regression to estimate the structural parameters. If we normalize $u(s; \theta(0))$ to be a function of time-invariant state variables that only influence the default utility but not the per-period utilities of the other options, we again only need to run a linear regression, by pooling $-s_t(0)$ and $\hat{E}[s_{t+1}(0)|s_t, a_t = k]$ into one regressor.

Note that, in the second regression (Step 2), we recover the agents' expectations nonparametrically via projection of the continuation value onto the current-period state variables, conditional on a_t and t . This is a key innovation of our estimator. Our projection approach contrasts with the commonly used alternative of simulating future state variables using the estimated transition functions in order to construct the continuation value function, such as in Hotz, Miller, Sanders and Smith (1994) or Bajari, Benkard and Levin (2007). The key advantage of the projection-based approach compared to the simulation-based approach is computational, a point that we detail in the next section.

Another key feature of our estimator is that we exploit the ability to represent the continuation value as a function of one-period-ahead choice probabilities, an intuition borrowed from Hotz and Miller (1993). This is advantageous when heavy data truncation makes it infeasible to construct the continuation value via forward simulation without relying on extrapolation of policy functions to unobserved periods.

4 Monte Carlo Experiments

In this section, we run a series of Monte Carlo experiments in order to numerically illustrate our key identification results and their implications for counterfactual analysis. Furthermore, we compare the performance of our estimator with existing estimation methods.¹³ For each experiment, we generated 500 data sets that simulate the behavior of mortgage borrowers. While we label our simulations as a “mortgage borrower’s problem” to be consistent with the model and the empirical illustration later in the paper, our framework generalizes to any finite-horizon discrete choice model with a terminating action. We assume that mortgages mature after twenty periods ($T = 20$) and that, for all $t < T$, the per-period utilities for prepaying and paying are given by $u(1, s) = \theta_1 + \alpha_1^1 s_1 + \alpha_1^2 s_2 + \varepsilon_1$ and $u(2, s) = \theta_2 + \alpha_2^1 s_1 + \alpha_2^2 s_2 + \varepsilon_2$, respectively. For the first few experiments, we specify the default utility to be $u(0, s) = \theta_0 + \varepsilon_0$, which we later modify to allow for dependence on the state variables. We choose parameter values for our data-generating process such that the generated default patterns vary meaningfully over time, a condition for identification in the theorems. Specifically:

- The state variables s_1 and s_2 are independent AR(1) processes with AR parameter 0.2 and a variance of 0.24 for the white noise. At time 0, s_1 and s_2 are drawn from independent normal distributions with mean 0 and variance 0.25.¹⁴
- ε_1 and ε_2 are drawn from the type 1 extreme value distribution.
- $\theta_0 = -4$, $\theta_1 = -1$, $\alpha_1^1 = 1$, $\alpha_1^2 = 0$, $\theta_2 = -2$, $\alpha_2^1 = 0$, $\alpha_2^2 = 1$, and $\beta = 0.9$.
- In the final period T , the per-period utilities of paying or prepaying are shifted up by 2 units, which captures a one-time boost to utility that the borrower would enjoy from having full ownership of the house once the mortgage is fully paid off at maturity. The per-period utility of default is the same as in all other periods.

Figures 1 and 2 graphically illustrate how the discount factor and utility of the terminating action are identified. Figure 1 plots the default probability (in percentage terms) as a function of time for various values of β , fixing the remaining parameters at their chosen values for the Monte Carlos and fixing the state variables at $(s_1 = 0.9, s_2 = 0)$. Figure 2 shows similar plots but instead varying the default utility $u(0, s)$ while fixing the remaining parameters. The figures show that both greater impatience (lower β) and a greater value of $u(0, s)$ are associated with a greater probability of default at all points in time.

¹³MATLAB code for the Monte Carlo exercises is provided as an online appendix.

¹⁴We chose these specific parameter values in order for the state variables to have constant variance over time, although this feature is certainly not important to our method.

However, while the impact of a lower β on default probability is non-monotonic in time, the impact of a greater $u(0, s)$ on default probability monotonically decreases in time. These contrasting implications are what allow the two parameters to be separately identified.

[Figures 1 and 2 about here]

Illustration of Key Identification Results

The first set of Monte Carlos numerically illustrate key identification results of the paper, namely identification of β and $u(0, s)$. For this purpose we consider four scenarios. Scenario (a) assumes that the researcher does not observe the final period, but rather, only has data on the first 14 periods. This scenario is chosen to illustrate our claim in Theorem 1 (ii) that the discount factor is identified in finite-horizon models even without data from the final period. Scenario (b) is the same as (a) but here we impose an incorrect normalization for the default payoff. We examine this scenario to illustrate our claim in Theorem 2 (i) that the choice of normalization for the default utility does not affect the recovered discount factor but does affect the recovered differences between the per-period payoff from payment or prepayment and the per-period payoff from default. In scenarios (c) and (d), we consider the case in which data from the final period are observed and $u(\cdot, \cdot) = u_T(\cdot, \cdot)$.¹⁵ These scenarios are designed to illustrate the claim in Theorem 2 (ii) that one can recover the utilities from all choices, along with the discount factor, without normalizing any of the payoffs, under the stated conditions. Scenario (c) considers the case where the default utility is a constant, whereas scenario (d) considers the case where the default utility is a function of the state variables s_1 and s_2 .

For each scenario, we estimate the model using various sample sizes. There are several reasons to expect that a large sample is necessary in order to precisely estimate the CCPs. First, the state variables are assumed to be drawn from continuous distributions, and the CCPs must be reasonably estimated over the entire support. Furthermore, non-stationarity of agents' optimal behavior implies that CCPs must be estimated separately for each time period, decreasing the effective sample size for estimation. Finally, given our chosen parameter values, defaults are rare close to the final period, necessitating a large sample in order to estimate the default probabilities close to the final period T to a reasonable degree of precision.

¹⁵To be more precise, we set $u(0, s) = u_T(0, s)$ and $u(k, s) + 2 = u_T(k, s)$ for $k = 1, 2$. Although the specified condition in the theorem states $u(k, s) = u_T(k, s)$ for $k = 0, 1, 2$, in fact the theorem holds as long as the researcher knows the exact relationship between $u(\cdot, \cdot)$ and $u_T(\cdot, \cdot)$ (which can be easily seen in the proof of the theorem). In the Monte Carlo study, since the researcher is assumed to know that the final period's utility from prepayment or payment is higher than earlier period's utility from the same action by 2 and impose that knowledge during estimation, the theorem still holds.

Table 1 summarizes the results from scenarios (a) and (b). Standard errors for the parameter estimates are computed as the standard deviation of the estimates across the 500 samples. The results from scenario (a) show that our proposed estimator performs reasonably well for all sample sizes. The discount factor estimate shows the greatest improvement in unbiasedness and precision as the sample size increases. The reason is that the discount factor is identified by variation in the continuation value, which we approximate based on the estimated CCPs. Estimation error introduced by CCP estimation leads to attenuation bias (biasing the estimate of β toward zero) through a familiar mechanism associated with errors in right-hand side variables. Nevertheless, for all sample sizes and all parameters, the implied 95% confidence interval from the estimator includes the true value of the parameter, and the estimates are also quite precise. These results numerically confirm that the structural parameters, including in particular the time preference parameter, are identified even if the researcher lacks information on agents' behavior in the final several periods.

[Table 1 about here]

The results from scenario (b) show that imposing an incorrect assumption about the level of the default utility has no effect on the estimate of the discount factor β , which is numerically equivalent between scenarios (a) and (b) in Table 1 (along with the “slope” parameters α_1^1 , α_1^2 , α_2^1 , and α_2^2). However, it does affect the recovered differences $u(2, s) - u(0, s)$ and $u(1, s) - u(0, s)$. For instance, if the correct normalization is imposed with a sample size of 80,000 (scenario (a)), the estimated value of $u(2, s) - u(0, s)$ is $-1.9890 + 0.0038s_1 + 1.0020s_2 + 4$, which is very close to the true difference of $2 + s_2$. However, the estimated difference becomes $-1.5401 + 0.0038s_1 + 1.0020s_2$ when we use an incorrect normalization (scenario (b), sample size 80,000). These results numerically confirm our identification results in Theorem 2 (i).

Scenarios (c) and (d), shown in Table 2, go one step further and confirm the identification results in Theorem 2 (ii). Here, we demonstrate our estimator's ability to correctly recover the utilities from all choices as well as the discount factor, without needing to normalize any of the payoffs, when the final period is observed and $u(\cdot, \cdot) = u_T(\cdot, \cdot)$. In scenario (c) we assume that the researcher is aware that $u(0, s)$ only has an intercept term and, accordingly, constrains the estimate of $u(0, s)$ not to depend on the state variables. In scenario (d) we assume that $u(0, s)$ is a function of the state variables s_1 and s_2 so that the researcher needs to recover the slopes as well as the intercept for $u(0, s)$. For scenario (d), we modify the specification of $u(0, s)$ to $u(0, s) = \theta_0 + \alpha_0^1 s_1 + \alpha_0^2 s_2 + \varepsilon_0$, where $\theta_0 = -1$, $\alpha_0^1 = 0.5$ and $\alpha_0^2 = -0.5$.

[Table 2 about here]

The results from Table 2 numerically illustrate identification of $u(0, s)$ for both scenarios (c) and (d). In particular, the results imply that we can identify the default utility even when the default utility depends on the same state variables that enter the non-default utilities. Although the estimates would be more precise, for a given sample size, when the default utility is a constant (scenario (c) of Table 2) or if it were a function of state variables that are excluded from the non-default utilities, the results illustrate that we can achieve identification of $u(0, s)$ even in the “worst case scenario” where the same state variables enter utilities of all actions, confirming our identification results in Theorem 2 (ii).

It is clear from scenario (d) in Table 2 that the data requirements for identification of $u(0, s)$ when it is state-dependent are much higher than for scenarios (a)-(c). As discussed earlier, identification of $u(0, s)$ stems from differences in the relative probabilities of the terminating versus non-terminating actions being chosen in the final period as compared with earlier periods, and from co-movement of these differences with state variables. In other words, the information obtained from the final period is crucial for identification of $u(0, s)$. Given our choice of parameters in simulating the data, only a small fraction of borrowers reach the final period and, among those borrowers, only a small fraction default in the final period, which prevents precise estimation of CCPs for the final period unless the sample size is very large. Noisy CCPs for the final period affect researchers’ ability to obtain precise estimates. The problem is especially severe when $u(0, s)$ is a function of the same state variables that enter $u(1, s)$ and $u(2, s)$ because we then need rich variation in the final period’s data as a function of the state variables, which is obtainable only with a very large data set.

Note that having precise CCP estimates is necessary for any two-step estimator. In order to meaningfully compare the data requirements of our approach against existing two-step approaches, we must assess how much *additional* noise stems from the projection and regression steps of our proposed estimator, excluding noise from CCP estimation. To this end, in Table 3, we re-estimate scenarios (a) and (d) by plugging the “true” (rather than estimated) CCPs into estimation. Therefore, by comparing the sample sizes required for a given degree of estimation precision in Table 1 (a) vs. Table 3 (a) and Table 2 (d) vs. Table 3 (d), we see that the projection and regression steps of our estimator do not impose a heavy requirement on sample size: as shown in Table 3, the bias is already negligible when the sample size is only 4,000, even when we need to estimate state-dependent $u(0, s)$.

[Table 3 about here]

Overall, the Monte Carlo results in Tables 1 and 2 show two things. First, they numerically illustrate the identification results of the paper. Second, they show that the data requirements of our proposed estimator are heavy because of the need for nonparametrically estimating the first-stage reduced form

policy function. Such heavy data requirements are common to any two-step estimator, and our particular approach is no exception. As such, our proposed method is at a disadvantage compared with full-solution methods based on maximum likelihood. The feasibility of two-step estimators, including ours, is constrained by the availability of a sufficiently large sample.¹⁶

Counterfactual Analysis

This subsection examines the implications of our identification results regarding $u(0, s)$ for counterfactual predictions.¹⁷ Even if the identification of $u(0, s)$ is of theoretical interest, it would not be empirically relevant if an incorrect normalization of $u(0, s)$ did not significantly affect counterfactual predictions, which are often the key object of interest for applied researchers.

We do not focus here on formal arguments for the identification of counterfactual predictions, which are examined more carefully in recent work by Aguirregabiria (2010), Norets and Tang (2014), Aguirregabiria and Suzuki (2014), Arcidiacono and Miller (2015a), Kalouptsi et al. (2016), and others. Instead, we demonstrate the empirical relevance of our main identification results by examining a few select counterfactual scenarios and demonstrating that making an incorrect assumption about $u(0, s)$ has a material effect on some counterfactual predictions.

The first counterfactual simulates the effects of increasing the discount factor β from 0.9 to 0.95. Because our model does not explicitly model liquidity constraints, we can think of β as capturing the true time preferences as well as the impact of liquidity constraints, which can cause borrowers to behave as though they discounted the future more heavily. This scenario thus attempts to capture (in an admittedly crude way) the effect of an exogenous government policy that relaxes liquidity constraints, such as by increasing access to credit.

To start, we simulate the fraction of borrowers that default by the final period T when we set β to first the actual value of 0.9 and then to the counterfactual value of 0.95, respectively (keeping all the other parameters the same between the two cases). Comparing predictions based on the two values of β indicates the expected impact of a relaxation in liquidity constraints. Next, we compute the counterfactual predictions (i.e., setting β to 0.95) when the researcher incorrectly assumes $u(0, s)$ is -2 instead of the true value of -4. The key thing to note is that here we appropriately adjusted the values of the intercepts in $u(1, s)$ and $u(2, s)$ in order to produce exactly the same CCPs under the actual scenario, even with

¹⁶Data size is likely to become less of a binding constraint in many increasingly popular empirical settings in which large consumer panels are available.

¹⁷We will investigate the implications of the identification regarding β for counterfactual predictions in our application, since estimation of β is the main focus of the application.

the incorrect normalization.¹⁸ In other words, choosing an incorrect normalization has no effect on our ability to rationalize the actual CCPs.

However, choosing an incorrect normalization does affect the counterfactual predictions. As shown in scenario (1) of Table 4, relaxing liquidity constraints (increasing the effective discount factor from 0.9 to 0.95) implies a reduction in the share of borrowers that default from 35.29% to 27.10%. The results further show that choosing the incorrect normalization for $u(0, s)$ leads to a prediction of 13.55% of borrowers defaulting under the counterfactual scenario, rather than 27.10%, resulting in a severe overestimate of the effects of the policy. This outcome highlights that a normalization that is seemingly innocuous in terms of explaining observed outcomes could very well lead to biased counterfactual predictions.

The second counterfactual simulates the effects of decreasing the discount factor β from 0.9 to 0, which allows us to examine how consumers' optimal default decisions would differ if they behaved myopically. We simulate the fraction of borrowers that default by the final period T when we set β to first the actual value of 0.9 and then to the counterfactual value of 0, respectively. Next, we compute the counterfactual predictions (i.e., setting β to 0) when the researcher incorrectly assumes $u(0, s)$ is -2 instead of the true value of -4. As before, we made appropriate adjustments to the values of the intercepts in $u(1, s)$ and $u(2, s)$ to ensure that the normalization has no effect on the factual CCPs.

The results for scenario (2), reported in Table 4, again show that an incorrect normalization of $u(0, s)$ biases the counterfactual predictions. The counterfactual prediction of the share of consumers defaulting given myopic behavior is 50.24% when the correct normalization is used. By comparison, the predicted share is 97.69% when the incorrect normalization is used, significantly overestimating the difference in default probability implied by myopic versus forward-looking behavior.

Whereas the previous examples highlight cases in which the normalization of $u(0, s)$ can alter counterfactual predictions, it is also illuminating to examine a case in which it does not. Scenario (3) considers the impact of an exogenous upward shift in the cost of default by 0.5. Such a shift may arise, for example, through a change in government regulations or lenders' policies.

We first simulate the share of consumers that default (by the final period T) based on the actual value of $u(0, s) = -4$, then again after increasing $u(0, s)$ by 0.5 to reflect the counterfactual policy shift. Then, we “mistakenly” normalized the actual value of $u(0, s)$ to -2, adjusted the remaining parameters such that the implied CCPs match the actual CCPs, and simulated outcomes after increasing $u(0, s)$ by 0.5 relative to the incorrect normalization. The results in Table 4 show that the counterfactual predictions for Scenario (3) are unaffected by the misspecification of $u(0, s)$ —whether or not we correctly specify the

¹⁸To be precise, when we increase $u(0, s)$ by c (in all periods), also adding $(1 - \beta)c$ to $u(1, s)$ and $u(2, s)$ for $t < T$ and adding c to $u(1, s)$ and $u(2, s)$ for the final period T ensures that the CCPs remain unaffected.

actual value of $u(0, s)$, in the counterfactual scenario in which $u(0, s)$ is incremented by 0.5, we predict that 24.26% of borrowers default.

[Table 4 about here]

In sum, the choice of normalization for $u(0, s)$ influences counterfactual predictions when the counterfactual scenario involves a shift in intertemporal tradeoffs (which may be due to a change in discounting, as we suppose in scenarios 1 or 2, or due to a change in the state transition process). However, for at least certain other types of counterfactuals, the choice of normalization is innocuous. While our goal is not to provide a comprehensive analysis of when normalization matters and when it does not, the results from our selected set of counterfactuals seem consistent with the results in Aguirregabiria (2010), Norets and Tang (2014), Aguirregabiria and Suzuki (2014), Arcidiacono and Miller (2015a), and Kalouptsi et al. (2016). While the considered settings, focus, and the results on (non)identification vary across the papers and one cannot summarize their results in one sentence without losing subtle distinctions across them, a somewhat simplistic summary of those papers is that counterfactuals that involve linear changes to the payoffs are identified even when the model is not (i.e., those counterfactuals are not affected by the choice of normalization), while counterfactuals that involve nonlinear changes to the per-period payoffs, changes in state transition or changes in β tend to be unidentified if the model is not identified (i.e., those counterfactuals tend to be affected by the choice of normalization). This is indeed similar to what we find in Table 4. While the earlier papers take non-identification of model primitives as given and consider identification of counterfactuals, a major contribution of our paper is to show that the model primitives, including the levels of utility, can be identified under a certain set of conditions, thereby allowing “all counterfactuals” to be identified.

Comparison with Alternative Estimators

Finally, we present Monte Carlos that compare our estimator’s performance against estimators based on existing methods in the literature. Specifically, we focus on adaptations of Hotz and Miller (HM henceforth) and BBL to our setting.¹⁹ Because our estimator is a multi-step estimator, we do not compare it against full-solution methods (e.g., Rust, 1987).

¹⁹To be more precise, we focus on Hotz and Miller estimator augmented with conditional choice simulator proposed by Hotz, Miller, Sanders and Smith (1994). The original Hotz and Miller estimator requires computation of conditional choice probability for all feasible future paths, which could become impractical when, for example, some observable state variables are continuous. To address this problem, Hotz, Miller, Sanders and Smith (1994) proposed use of forward simulation so that only path of simulated future choices needs to be considered. Since the use of forward simulation clearly speeds up estimation, we use the Hotz and Miller estimator augmented with forward simulation, instead of the original Hotz and Miller estimator, in our comparison. For simplicity, we still call the estimator in the comparison group as Hotz and Miller.

Before proceeding, it is useful to summarize the key similarities and differences between our approach, HM, and BBL. The main idea in HM is to estimate a structural equation while including, as a control, an approximation of the continuation value constructed from the CCPs. The structural estimation objective functions for both HM and our approach are thus ultimately based on comparing model-predicted and observed CCPs (or, equivalently, comparing the implied choice-specific value functions). By contrast, structural estimation in BBL uses an objective function based on violations of inequalities formed from comparing the true value function (constructed using the estimated policy function) and alternative value functions (constructed using perturbed policy functions).

Our approach represents the continuation value as a function of the one-period-ahead choice probabilities, which is possible due to the presence of a terminating action. When the horizon is finite and data from the final period are unavailable, the “simplified” representation of the continuation value (i.e., relative to one obtained through simulating for all future periods) obviates the need for extrapolating the estimated policy function to unobserved time periods, which is potentially invalid. Furthermore, in case the state transition function is misspecified, relying on forward simulation for all future periods would worsen estimation bias because the amount of misspecification will be amplified towards the end period, whereas using the simplified representation would limit the extent of bias.

Both HM and BBL rely on forward simulation to approximate the continuation value, and use a nonlinear objective function to recover the structural parameters. In contrast, we rely on a linear projection to do the approximation, and use OLS to estimate the structural parameters. Our projection-based approach should significantly improve computational performance, and our use of OLS—which has a fast analytic solution—should also improve the robustness and computational performance of our estimator.

To see how each of these features affects the performance of our estimator relative to existing methods, we compare our approach against several alternative adaptations of HM and BBL, which are summarized in Table 5. Each variant highlights a feature of our estimator that contributes to its computational performance and robustness. These adaptations are modifications of the original HM and BBL estimators that incorporate the key features of our estimator. Thus, although we label them as adaptations of HM and BBL, they should be viewed as a hybrid between those estimators and our proposed estimator.

[Table 5 about here]

The first set of adaptations (which we label “HM1” and “BBL1”) does not exploit the ability to represent the continuation value in terms of the one-period-ahead choice probabilities and thus requires forward simulation all the way through the final period. As such, HM1 and BBL1 require extrapolating the policy functions to unobserved time periods ($t= 15$ through 20), because the optimal policy function

depends on t . Our approach is to include the time index t as a state variable when estimating the policy function for the pooled data from $t = 1$ through 14. We then predict the policy function for unobserved periods $t = 15$ through 20 by replacing the value of t in the estimated policy function.

We also compare the performance of our estimator against adaptations of HM and BBL (labeled “HM2” and “BBL2”) that exploit the simplified representation of the continuation value but still rely on simulation in order to compute its approximation. In addition, we compare the performance of our estimator against an adaptation of BBL (“BBL3”) that exploits the simplified representation for the continuation value as well as uses projection instead of simulation to construct the continuation value function. In each of these exercises, we will use scenario (a). Using scenario (a) for the comparison will allow us to examine the performance of the estimators for the frequently encountered case in which the researcher does not have access to data from the final periods.

The results are reported in Tables 6 and 7.^{20,21} These tables compare the mean estimates and computational time, respectively, across the various methods using different sample sizes. In testing the HM and BBL variants, we found the resulting estimates to be sensitive to the initial values in the structural estimation step. To illustrate this sensitivity, we report results under two different initial values for each of HM1, HM2 and BBL1; these are denoted as HM1(a)-HM1(b), HM2(a)-HM2(b), and BBL1(a)-BBL1(b), respectively.²² The issue of choosing starting values does not apply to our method because it uses OLS and hence does not require initial values to be specified. We also report two versions for each of BBL2 and BBL3. For the first version in each case, BBL2(a) and BBL3(a), the set of alternative policy functions is carefully chosen to produce sensible estimates, which involved a significant amount of manual tuning through trial and error. For the second version, BBL2(b) and BBL3(b), the set of alternative policy functions is similar to the ones prescribed in the original BBL paper (2007), which adds a random noise to each period’s policy function. For estimators that require simulation (BBL1, BBL2, HM1 and HM2), we use 250 simulation draws.²³

Table 6 shows the mean parameter estimates obtained using each method. The estimates for BBL1(a) are very sensible even when the sample size is small. Our estimator produces similarly good estimates for larger sample sizes, but for the smallest sample size, the BBL1(a) estimator is less biased than ours. However, the BBL1 estimator is less biased than ours only when the choice of initial values is “lucky.”

²⁰In these tables, we assume that the researcher uses correctly specified state transitions in case of HM and BBL.

²¹We used Hermite polynomials of the third degree to estimate the CCPs for our proposed estimator as well as for our implementation of HM and BBL. For our proposed estimator, we also used the same basis functions to compute the projection of the continuation value.

²²We find that results for BBL2 and BBL3 are also sensitive to the choice of initial values, but we do not present separate cases depending on the initial values for those due to space constraints. Results are available upon request.

²³We decided to use 250 draws for simulation since we found that using less than 250 simulation draws occasionally resulted in implausible estimates.

As shown by the BBL1(b) estimates, BBL1 produces rather biased results under different initial values. Since it is impossible to determine a priori which initial values will work and which values will not work,²⁴ these results indicate lack of robustness of the BBL estimator for the considered model.

[Table 6 about here]

The estimates for BBL2 and BBL3 depend strongly on how the alternative policy functions are specified in constructing the objective function. For BBL2(a) and BBL3(a), we used a set of alternative policy functions that we found, through much trial and error, to produce estimates that were reasonably close to the true parameters. This amounts to considering a special restricted class of deviations from the agents’ observed policy functions, where the restrictions take into account the specific decay of the optimal default probabilities over time. Even so, the estimates are rather poor for the smallest sample size. And even with the larger samples, the estimate of the discount factor is more biased than using our method or BBL1(a) with the same sample size. In actual empirical research, the true parameters are, naturally, unobserved. Thus, we consider BBL2(a) and BBL3(a) to be essentially “infeasible” estimators. The fact that even these infeasible estimates do not perform well raises a concern about BBL2 and BBL3’s performance in settings considered in this paper. Furthermore, as shown in BBL2(b) and BBL3(b), we obtain very poor estimates when we construct the alternative policy functions by simply adding random noise to each period’s optimal policy. In sum, comparing BBL2 and BBL3 with BBL1 suggests that incorporating some key features of our estimator—use of the simplified representation of the continuation value and use of projection—within the BBL framework adds little value (except through computational savings, discussed below) and can rather hurt the estimator’s performance. Overall, we conclude that none of the various versions of BBL estimator considered in the paper is preferable to our estimator in terms of unbiasedness and robustness.

Examining HM1, we see that, similar to BBL1, the estimates are again sensitive to the choice of initial values. The estimates are reasonable for a certain set of initial values (HM1(a)), but not for others (HM1(b)). This makes the HM1 estimator undesirable for the considered model. We find that the HM2 estimates, which exploit the simplified representation of the continuation value but still rely on simulation (rather than projection) to generate them, are also somewhat sensitive to the choice of starting values, albeit less so than the other estimators that we examined. Overall, the HM2 estimator seems to perform the best among all of the adaptations of HM and BBL, in terms of robustness and unbiasedness.

²⁴We found that initial guesses close to the true parameter values typically lead to sensible estimates, but not always.

The overall finding from Table 6 is that our estimator performs the best in terms of unbiasedness and robustness.²⁵²⁶ A close second is HM2, but we find HM2 to be somewhat sensitive to the initial values. All variants of BBL are quite sensitive to either the choice of initial values or to the choice of alternative policy functions. Thus, BBL may not be an appropriate methodology for empirical settings similar to the one we explore. The HM estimator that does not use the simplified representation also does not perform very well in our setting, as it is also highly sensitive to the choice of initial values. The finding that HM2 is the best-performing among all variants of HM and BBL is in fact not surprising, because once we use the simplified representation of the continuation value within the HM estimator, that approach only differs from ours in terms of the method for constructing the continuation value and structural estimation objective function.

To examine how an alternative extrapolation method would affect the results, we re-estimated HM1 and BBL1 while assuming that the policy function for $t = 15$ through 20 is the same as the policy function from the last observed period $t = 14$. The results, not reported but available upon request, show that the estimates become much less sensible when using this rather naive extrapolation method. This outcome illustrates that using the simplified representation of the continuation value is helpful for consistent estimation when the horizon is finite and data from the final periods are unavailable.

Our estimator has much better computational performance than the competing methods. Table 7 reports how long it takes (in minutes) to run the Monte Carlo simulations 500 times using each estimator. BBL1 takes 25-30 times longer to run than our estimator even for “good” initial values. HM1 also takes approximately 20 times longer than our estimator. Thus, computational performance gives an additional reason why our approach may be an appealing choice in other similar empirical settings. Comparing the computational times for BBL1 and BBL2, and comparing HM1 and HM2, demonstrates that the use of the simplified expression for the continuation value reduces computational time by about 90% for both BBL and HM. The use of projection instead of simulation reduces the computational time by an additional 60% for BBL (BBL2 versus BBL3). In sum, we find that the key features of our estimator help to significantly improve computational performance.

[Table 7 about here]

The results clearly demonstrate that, for the setting that we consider (namely, a finite-horizon discrete decision model with a terminating action and lack of data for final periods), our proposed estimator

²⁵The results on mean bias and standard errors, reported in the online appendix, also show that our proposed estimator outperforms the competing methods in terms of mean bias and standard error.

²⁶The fact that our estimator is easier to code, reducing the risk of analyst error, is hard to incorporate into a Monte Carlo, but would likely increase the precision and reduce the bias of our estimator relative to its rivals in applications.

performs better than the alternatives derived from existing approaches in the literature, as measured by unbiasedness, robustness and computational time.

5 Empirical Illustration: Subprime Mortgage Default

In this section, we illustrate our proposed estimation methodology using data from CoreLogic on subprime mortgages originated between January 2000 and September 2007 and securitized in the private-label market. For each loan, we observe the loan terms and borrower characteristics reported at the time of loan origination, such as the type of mortgage (fixed rate, adjustable rate, etc.), the initial contract interest rate, the level of loan documentation (full, low, or none), the appraised value of the property, the loan-to-value ratio, the location of the property by zip code, and the borrower’s FICO score. As we explain below, we do not observe behavior in the final period for any of the borrowers, and thus we focus primarily on estimating the discount factor β , which we can still identify in the presence of data truncation.

To simplify the exposition of our estimation methodology, we restrict our sample in a number of ways in order to make it more homogeneous. Thus, we do not claim that our sample is representative of the entire subprime population. Specifically, we focus on 30-year fixed-rate mortgages, the most common mortgage type. We also restrict our sample to first-lien loans located in 20 major Metropolitan Statistical Areas (MSAs).²⁷ Furthermore, we exclude “cash-out” refinance loans,²⁸ focusing on loans for home purchases or refinances with no cash out. Finally, we restrict our sample to the small proportion (3.5%) of loans for which we are able to impute the borrower’s income.²⁹ With these restrictions, we have approximately 11,500 borrowers in the sample.

The data track each loan over the course of its life, showing the outstanding balance, delinquency status, and scheduled payment in each month. Three possible actions (default, prepay, and pay) are available to each borrower in every period. We define default as occurring if the bank takes possession of the home or if the loan has been delinquent for 90 days or more, a commonly used definition of default in the mortgage literature. Default is a terminal event, so if a loan defaults in month t , it drops from the sample starting at $t + 1$. We define prepayment as occurring if the loan balance is observed going to zero before maturity, presumably because the borrower has paid off the loan in full. We do not observe

²⁷ Atlanta, Boston, Charlotte, Chicago, Cleveland, Dallas, Denver, Detroit, Las Vegas, Los Angeles, Miami, Minneapolis, New York, Phoenix, Portland, San Diego, San Francisco, Seattle, Tampa, and Washington D.C.

²⁸ In a cash-out refinance, the borrower pays off the balance on the existing loan with a larger loan, and receives cash for the difference.

²⁹ Specifically, for these loans, we observe the “front-end debt-to-income” ratio—defined as the ratio of monthly mortgage-related payments to the borrower’s income—from which we can impute the borrower’s income based on the fact that we also observe the numerator of the ratio.

the new loan used to refinance the original loan, because our data identify individual loans but not the identities of specific borrowers. Thus, when the borrower prepays, we assume that the borrower refinances into a new loan that matures at the same maturity date as the old loan and has an interest rate equal to the current market interest rate. Finally, if the borrower neither defaults nor prepays in month t , the borrower continues to make just the regularly scheduled payment, and the loan survives into the next month. Our data allow us to track the status of each loan in the sample through March 2013. Due to the relatively short span of the sample compared with the mortgage term (30 years), we do not have data on the behavior of borrowers whose mortgages are close to maturity. Theorem 2 (*i*), which deals with the case of truncated data, thus applies to our empirical setting.

The full list of state variables is described in Table 8. Some are static, varying across borrowers but not over time—namely, the previously described loan terms and borrower characteristics reported at the time of origination; and MSA dummies for the property location. We also include state variables that vary both over borrowers and over time: (1) a proxy for the property value, which we impute from the appraised property value at the time of origination, adjusted by a zipcode-level home price index; (2) the borrower’s net equity, defined as the imputed property value net of the outstanding loan balance; (3) a proxy for the current market interest rate available to the borrower; (4) the county-level unemployment level (by month), which proxies for potential income fluctuations affecting the borrower over time. One major limitation of the data is lack of information on individual level employment or income shocks. Although we control for various borrower characteristics as well as the county-level unemployment rate, these controls are unlikely to fully capture the idiosyncratic factors affecting individual borrower risk over time. This limitation raises an important caveat in interpreting our results below.

[Table 8 about here]

Before we present the estimates, we discuss whether the key features of the model, which are required for identification, are satisfied in our application. One of the key assumptions of the model is time homogeneity of utility functions. The utility from default likely depends on the cost of damaged credit, the cost of having to move and find another place to live, “stigma” and possible psychological costs associated with default, etc. It seems sensible to assume that these factors do not systematically vary with the loan age. The period utility from prepayment and payment likely depends on the utility from residing in the house, the disutility from making payments, etc. Again, it seems unlikely that these factors systematically vary with the loan age (for all periods $t < T$) after conditioning on the state variables. Therefore, we think the assumption that the utility functions do not directly depend on the loan age, conditional on state variables, is sensible in our application for all periods except for the final period. In

the last period, the period utility of choosing to “pay” presumably includes a lump-sum boost due to the homeowner’s finally paying off the loan and obtaining full ownership of the house. Therefore, it seems unreasonable to assume that time homogeneity of the utility functions extends to the final period.

The assumption of $u(\cdot, \cdot) = u_T(\cdot, \cdot)$ is required for identification results in Theorems 1 (i) and 2 (ii), which rely on availability of data from the final period. In our empirical application, we do not observe behavior in the final period for any of the borrowers. Thus we focus primarily on estimating the discount factor without the use of final-period data, and the assumption of $u(\cdot, \cdot) = u_T(\cdot, \cdot)$ is not required in our empirical application.

It also seems reasonable to assume that the state transitions do not depend directly on the loan age. Movements in housing prices or market interest rates—being market-wide variables—obviously do not depend on the age of a particular loan. Similarly, although a borrower’s net equity and loan balance evolve systematically over the course of a loan’s life, our state transition process already explicitly controls for all of the relevant factors that determine their evolution—namely, the existing loan balance, monthly payment amount, and house price changes. Therefore, we consider the assumption of time-homogeneous state transition to be satisfied in our setting.

Another key requirement is the presence of a finite horizon. In our empirical application, we focused on 30 year fixed rate mortgages. Because the maturity of a given mortgage is fixed, we regard the assumption of an exogenously fixed T to be appropriate.³⁰

Also required is the presence of a terminating action. For default to be a terminating action in our application, the default utility should not depend on a future sequence of optimal actions (of the considered model) following default. Empirically, there are cases in which a borrower misses several payments and then becomes current again by making back-payments. In order for default to be a terminating action, we must preclude such cases. For this reason, we define default as occurring if the bank takes possession of the home or if the loan has been delinquent for 90 days or more, because it is very rare for mortgages to become current again after they have been delinquent for so long. Thus, default defined in this way can be reasonably thought of as being a terminal action.

We follow the estimation procedure outlined in Section 3. To estimate the CCPs, we employ a sieve logit with splines of the state variables in order to flexibly model borrowers’ choice probabilities. The sieve basis includes restricted cubic splines for the continuous state variables, interpolating between 3 equally spaced percentiles of each state variable’s marginal distribution. It also includes interactions

³⁰It is possible that the maturity of a loan is negotiated in case the loan goes through a loan modification program. However, loan modifications occur very rarely, most of the modifications involve rate reduction or principal reduction, but not an extension of a loan term, and, importantly, whether a loan modification is granted or not is at the discretion of lenders, not borrowers. Therefore, it seems reasonable to assume that T is seen as exogenously fixed from the perspective of a borrower.

among the state variables.

For the structural estimation, we must normalize the utility of default because we lack data on the final period. We assume that a borrower’s utility from default depends on the borrower’s credit score and MSA of residence. The borrower’s credit score enters the utility of default because the amount of damage to the borrower’s credit caused by a default conceivably depends on the borrower’s existing credit quality. The MSA is relevant because mortgages are recourse loans in some states but non-recourse loans in others.³¹ We assume that the FICO score and MSA do not directly affect the utility from payment or prepayment, and therefore include these variables only in the utility function of default. Due to these exclusion restrictions, we can identify the coefficients on the FICO score and on the MSA dummies in the expression for the default utility, even though we do not observe behavior in the final period. However, we still cannot identify the intercept of the default utility separately from the intercept of the payment utility or prepayment utility. Therefore, we normalize the intercept of the default utility to zero. It is worth mentioning that, following Theorem 2 (*i*), this normalization does not affect the estimate of the discount factor, which is a key parameter of interest.

We jointly estimate the equations for prepayment and payment (7) using a seemingly unrelated regression, with the discount factor and the parameters of the default payoff constrained to be the same across equations. We also constrain the coefficients that capture the disutility from monthly payments and the utility from housing services (expressed as a percentage of housing value) to be the same across the prepayment and payment equations. We parametrize the utility functions to be linear in the state variables. Table 9 reports our estimates of the structural parameters.

[Table 9 about here]

Our results for the per-period payoff from default and the per-period payoff from paying are sensible. Borrowers with higher FICO scores have a lower utility of default, consistent with the notion that default causes greater damage to the credit of borrowers with good existing credit. Higher home value and lower monthly payments increase the utility of continuing to pay. Higher local unemployment, a proxy for the likelihood of job loss, reduces the utility of continuing to pay. Borrowers with low-documentation loans, who are presumably riskier, have lower utility from paying. These indicators of greater borrower risk may proxy for a higher probability of binding liquidity constraints, which effectively raises the cost of making the monthly payments and thus decreases the probability of choosing the “pay” option.

The intuition for the prepayment equation estimates is somewhat more nuanced. The probability of prepayment depends on both the borrowers’ willingness to refinance and their ability to do so. In

³¹Under non-recourse, lenders cannot go after a defaulter’s assets other than the mortgage collateral (i.e., the house), which lowers the perceived cost of default to the borrower.

particular, borrowers may be unable to refinance if lenders deem them to be too risky. Thus, it makes sense that the coefficient estimates imply that higher local unemployment as well as having characteristics that may proxy for greater borrower risk, such as low-doc loans and multiple liens, decrease the propensity to prepay.

The estimate of the discount factor, the key focus of our analysis, is 0.953 (monthly), which is statistically significantly less than 1, as one might expect. Our estimation does not restrict the magnitude of the discount factor, so it is reassuring that our estimate of the discount factor is a plausible number. Interestingly, the monthly interest rate implied by the estimated discount factor is much higher than the average monthly interest rates these borrowers pay on their mortgages in the data, and even higher than the “risk-free” rate. In other words, these borrowers seem to discount the future at a higher rate than the discount rate implied by market interest rates. In the literature, a typical approach to dealing with the discount factor when it cannot be estimated is to fix it to some economy-wide rate of return on assets. The significance of our finding is that such an approach would be incorrect in our setting, potentially leading to bias in the other structural estimates, which might have a significant impact on counterfactual predictions of interest.

To investigate the extent to which imposing an incorrect assumption about β can bias counterfactual predictions, we analyze a “static choice” counterfactual scenario—that is, a scenario in which borrowers are assumed to be myopic and make a statically optimal decision in every period. For this purpose, we generate predictions based on two sets of counterfactual simulations. For each set of simulations, we set β to 0 and predict the share of borrowers that would default or prepay, respectively, at some point before the end of the sample period.

For the first set of simulations (shown in the upper left panel of Table 10), we set the remaining parameters to the unbiased estimates reported in Table 9. For the second set of simulations (shown in the upper right panel of Table 10), we set the remaining parameters to estimates obtained while imposing the incorrect constraint $\beta = 0.993$.³² Note that imposing an incorrect constraint also biases the remaining parameter estimates.

By comparing the two sets of predictions, we can determine how incorrectly specifying β in estimation affects the counterfactual predictions. As shown in the table, incorrectly specifying β introduces significant bias in the counterfactual predictions. Specifically, with the unbiased parameter estimates, we predict that, on average, the fractions of borrowers choosing to default at some point, prepay at some point, or make just the regularly scheduled payments through the end of the sample period are 3.7%,

³²The value 0.993 corresponds to $1/(1 + \bar{r})$, where $\bar{r}$ is the empirical mean of the monthly interest rate that borrowers pay on their mortgages.

57.9%, and 38.4%, respectively. The predicted number of defaults, 422 defaults out of 11523 loans, is thus significantly lower than implied by predictions for the factual case based on the actual estimate of β (shown in the bottom left panel of Table 10).

This reduction in default among myopic borrowers compared to forward-looking borrowers is likely due to the market conditions during the examined period. The examined period includes the period during which the housing market plummeted, leading to expectations of home price decline in the future. This pessimistic view on the future housing prices would make default more desirable among forward-looking borrowers, but not among myopic borrowers. This could explain why the predicted probability of default is lower under the static model than under the dynamic model.

[Table 10 about here]

When we use the biased parameter estimates (i.e., those obtained under the incorrect assumption about the value of β), the counterfactual simulation erroneously implies that there would be almost no default (6.2 defaults out of 11523 loans), significantly overestimating the difference in default probability implied by myopic versus forward-looking behavior. Thus, misspecification of β in estimation leads to misleading counterfactual conclusions. This outcome highlights the empirical relevance of being able to correctly estimate the discount factor, rather than having to assume an arbitrary number.

Importantly, imposing an incorrect assumption about β does not lead to significantly different predictions for the “actual scenario,” relative to predictions based on unbiased parameter estimates (as seen by comparing the bottom left and bottom right panels of Table 10). Both predictions are reasonably close to the actual behavior seen in the data (not shown). That is, imposing an incorrect assumption about β has little effect on the within-sample fit and an unsuspecting researcher would see little evidence against a misspecified value of β .

6 Conclusion

In this paper we study identification and construct a computationally efficient estimation method for finite horizon optimal stopping problems, an important subset of dynamic discrete choice models. We first prove that we can identify the discount factor in this kind of setting without the availability of data from the final period. Next, we provide new results on identifying the utility of all of the agent’s potential choices, without having to normalize the utility from one of them. Finally, we propose an intuitive and easily implementable estimation method for finite horizon discrete choice problems with a terminating action.

Our Monte Carlo exercises numerically illustrate our identification results and their implications for counterfactual analysis, and also allow us to compare our proposed estimator against existing approaches in the literature. The results show that, for the setting considered in this paper, our proposed method outperforms the existing two-step estimation methods in terms of unbiasedness, robustness and computational time. Thus, our method could be an appealing choice for studying empirical settings in which there is a finite horizon discrete decision problem with a terminating action. An empirical illustration of our methodology using default decisions of subprime mortgage borrowers provides insights on time preferences of an economically important segment of consumers.

References

- [1] **Aguirregabiria, Victor**. 2010. “Another Look at the Identification of Dynamic Discrete Decision Processes: An Application to Retirement Behavior.” *Journal of Business & Economic Statistics*, 28(2): 201-218.
- [2] **Aguirregabiria, Victor and Arvind Magesan**. 2013. “Euler Equations for the Estimation of Dynamic Discrete Choice Structural Models.” *Advances in Econometrics*, 31: 3-44.
- [3] **Aguirregabiria, Victor and Junichi Suzuki**. 2014. “Identification and Counterfactuals in Dynamic Models of Market Entry and Exit.” *Quantitative Marketing and Economics*, 12(3): 267-304.
- [4] **Ai, Chunrong and Xiaohong Chen**. 2003. “Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions.” *Econometrica*, 71(6): 1795-1843.
- [5] **Ahn, Hyungtaik, and Charles Manski**. 1993. “Distribution Theory for the Analysis of Binary Choice under Uncertainty with Nonparametric Estimation of Expectations.” *Journal of Econometrics*, 56(3): 291-321.
- [6] **Altug, Sumru and Robert Miller**. 1998. “The Effect of Work Experience on Female Wages and Labour Supply.” *Review of Economic Studies*, 65(1): 45-85.
- [7] **Andrews, Donald**. 1991. “Asymptotic Normality of Series Estimators for Nonparametric and Semiparametric Regression Models.” *Econometrica*, 59(2): 307-345.
- [8] **Arcidiacono, Peter, Pat Bayer, Jason Blevins, and Paul Ellickson**. 2015. “Estimation of Dynamic Discrete Choice Models in Continuous Time with an Application to Retail Competition.” forthcoming in *Review of Economic Studies*.
- [9] **Arcidiacono, Peter and Robert Miller**. 2011 “CCP Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity.” *Econometrica*, 79: 1823-1867.
- [10] **Arcidiacono, Peter and Robert Miller**. 2015a. “Identifying Dynamic Discrete Choice Models off Short Panels.” working paper.
- [11] **Arcidiacono, Peter and Robert Miller**. 2015b. “Nonstationary Dynamic Models with Finite Dependence.” working paper.
- [12] **Bajari, Patrick, Lanier Benkard and Jonathan Levin**. 2007. “Estimating dynamic models of imperfect competition.” *Econometrica*, 75(5): 1331-1370.

- [13] **Bajari, Patrick, Han Hong and Denis Nekipelov.** 2013. "Game Theory and Econometrics: A Survey of Some Recent Research." in *Advances in Economics and Econometrics: Tenth World Congress of Econometric Society*, Vol. 3, Cambridge University Press.
- [14] **Beauchamp, Andrew.** 2015. "Regulation, Imperfect Competition, and the U.S. Abortion Market." *International Economic Review*, 56(3): 963-996.
- [15] **Chen, Xiaohong.** 2007. "Large Sample Sieve Estimation of Semi-nonparametric Models." *Handbook of Econometrics*, 7: 5549-5632.
- [16] **Chen, Xiaohong, Oliver Linton and Ingrid van Keilegom.** 2003. "Estimation of Semiparametric Models when the Criterion Function Is Not Smooth." *Econometrica*, 71(5): 1591-1608.
- [17] **Chen, Xiaohong, Victor Chernozhukov, Sokbae Lee and Whitney Newey.** 2014. "Local Identification of Nonparametric and Semiparametric Models." *Econometrica*, 82(2): 785-809.
- [18] **Chung, Doug, Thomas Steenburgh and K. Sudhir.** 2014. "Do Bonuses Enhance Sales Productivity? A Dynamic Structural Analysis of Bonus-Based Compensation Plans." *Marketing Science*, 33(2): 165-187.
- [19] **Dubé, Jean-Pierre, Günter Hitsch and Pranav Jindal.** 2014. "The Joint Identification of Utility and Discount Functions from Stated Choice Data: An Application to Durable Goods Adoption." *Quantitative Marketing and Economics*, 12: 331-377.
- [20] **Duffie, Darrell, and Kenneth J. Singleton.** 1997. "An Econometric Model of the Term Structure of Interest-Rate Swap Yields." *Journal of Finance*, 52(4): 1287-1321.
- [21] **Eckstein, Zvi and Kenneth Wolpin.** 1999. "Why Youths Drop Out of High School: The Impact of Preferences, Opportunities, and Abilities." *Econometrica*, 67(6): 1295-1339.
- [22] **Fang, Hanming, and Yang Wang.** 2015. "Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions." *International Economic Review*, 56(2): 565-596.
- [23] **Frederick, Shane, George Loewenstein and Ted O'Donohue.** 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, 40: 350-401.
- [24] **Härdle, Wolfgang.** 1990. *Applied Nonparametric Regression*. Cambridge University Press.
- [25] **Hausman, Jerry.** 1979. "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables." *Bell Journal of Economics*, 10(1): 33-54.
- [26] **Heckman, James and Salvador Navarro.** 2007. "Dynamic Discrete Choice and Dynamic Treatment Effects." *Journal of Econometrics*, 136 (2): 341-396.
- [27] **Hotz, Joseph, and Robert Miller.** 1993. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *Review of Economic Studies*, 60(3): 497-529.
- [28] **Hotz, Joseph, Robert Miller, Seth Sanders and Jeffrey Smith.** 1994. "A Simulation Estimator for Dynamic Models of Discrete Choice." *Review of Economic Studies*, 61(2): 265-289.
- [29] **Joensen, Juanna Schrøter.** 2009. "Academic and Labor Market Success: The Impact of Student Employment, Abilities, and Preferences." working paper.
- [30] **Judd, Kenneth.** 1998. *Numerical Methods in Economics*. MIT Press.
- [31] **Kalouptside, Myrto.** 2014. "Time to Build and Fluctuations in Bulk Shipping." *American Economic Review*, 104(2): 564-608.
- [32] **Kalouptside, Myrto, Paul Scott and Eduardo Souza-Rodrigues.** 2016. "Identification of Counterfactuals in Dynamic Discrete Choice Models." Working Paper.

- [33] **Magnac, Thierry, and David Thesmar.** 2002. "Identifying Dynamic Discrete Decision Processes." *Econometrica*, 70(2): 801-816.
- [34] **Mammen, Enno, Christoph Rothe, and Melanie Schienle.** 2012. "Nonparametric Regression With Nonparametrically Generated Covariates." *Annals of Statistics*, 40(2): 1132-1170.
- [35] **Manski, Charles.** 1991. "Nonparametric Estimation of Expectations in the Analysis of Discrete Choice under Uncertainty." *Nonparametric and Semiparametric Methods in Econometrics and Statistics: Proceedings of the Fifth International Symposium in Economic Theory and Econometrics*, Cambridge University Press.
- [36] **Manski, Charles.** 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies*, 60(3): 531-542.
- [37] **Manski, Charles.** 2000. "Identification Problems and Decisions under Ambiguity: Empirical Analysis of Treatment Response and Normative Analysis of Treatment Choice." *Journal of Econometrics*, 95(2): 415-442.
- [38] **Newey, Whitney.** 1997. "Convergence Rates and Asymptotic Normality for Series Estimators." *Journal of Econometrics*, 79: 147-168.
- [39] **Newey, Whitney and James Powell.** 2003. "Instrumental Variable Estimation of Nonparametric Models." *Econometrica*, 71(5): 1565-1578.
- [40] **Norets, Andriy and Satoru Takahashi.** 2013. "On the Surjectivity of the Mapping between Utilities and Choice Probabilities." *Quantitative Economics*, 4: 149-155.
- [41] **Norets, Andriy and Xun Tang.** 2014. "Semiparametric Inference in Dynamic Binary Choice Models." *Review of Economic Studies*, 81(3): 1229-1262.
- [42] **Pagan, Adrian and Aman Ullah.** 1999. *Nonparametric Econometrics*. Cambridge University Press.
- [43] **Pesendorfer, Martin and Philipp Schmidt-Dengler.** 2008. "Asymptotic Least Squares Estimators for Dynamic Games." *Review of Economic Studies*, 75(3): 901-928.
- [44] **Pollard, David.** 1984. *Convergence of Stochastic Processes*, Springer-Verlag.
- [45] **Rust, John.** 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica*, 55(5): 999-1033.
- [46] **Rust, John.** 1994. "Structural Estimation of Markov Decision Processes." in Robert Engle and Daniel McFadden eds., *Handbook of Econometrics*, Vol. 4. Amsterdam: North-Holland.
- [47] **Rust, John and Christopher Phelan.** 1997. "How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets." *Econometrica*, 65(4): 781-831.
- [48] **Scott, Paul.** 2013. "Dynamic Discrete Choice Estimation of Agricultural Land Use." working paper.
- [49] **van der Vaart, Aad and Jon Wellner.** 1996. *Weak Convergence and Empirical Processes*, Springer.
- [50] **Wong, Wing Hung and Xiaotong Shen.** 1995. "Probability Inequalities for Likelihood Ratios and Convergence Rates of Sieve MLES." *Annals of Statistics*, 23(2): 339-362.
- [51] **Yao, Song, Carl F. Mela, Jeongwen Chiang, and Yuxin Chen.** 2012. "Determining Consumers' Discount Rates With Field Studies." *Journal of Marketing Research*, 49(6): 822-841.

Figure 1: Time Trend of Default Probability for Various Values of β

Figure 2: Time Trend of Default Probability for Various Values of u_0

Table 1: Monte Carlo Scenarios (a) and (b)

Scenario (a): data not available for the final few periods								
		θ_1 (-1)	α_1^1 (1)	α_1^2 (0)	θ_2 (-2)	α_2^1 (0)	α_2^2 (1)	β (0.9)
$N =$	Mean	-0.9215	1.0831	0.0260	-1.9245	0.0772	1.0331	0.7086
5,000	SE	0.0408	0.0664	0.0568	0.0411	0.0677	0.0613	0.1069
$N =$	Mean	-0.9584	1.0394	0.0147	-1.9598	0.0355	1.0180	0.8018
10,000	SE	0.0267	0.0454	0.0365	0.0275	0.0462	0.0377	0.0764
$N =$	Mean	-0.9759	1.0183	0.0058	-1.9766	0.0167	1.0068	0.8488
20,000	SE	0.0166	0.0297	0.0241	0.0169	0.0306	0.0253	0.0517
$N =$	Mean	-0.9845	1.0092	0.0044	-1.9846	0.0078	1.0051	0.8752
40,000	SE	0.0113	0.0202	0.0164	0.0116	0.0205	0.0171	0.0350
$N =$	Mean	-0.9888	1.0043	0.0019	-1.9890	0.0038	1.0020	0.8878
80,000	SE	0.0075	0.0136	0.0119	0.0077	0.0142	0.0125	0.0246
Scenario (b): researcher incorrectly normalizes θ_0 to 0 when true θ_0 is -4								
		θ_1 (-1)	α_1^1 (1)	α_1^2 (0)	θ_2 (-2)	α_2^1 (0)	α_2^2 (1)	β (0.9)
$N =$	Mean	0.2441	1.0831	0.0260	-0.7590	0.0772	1.0331	0.7086
5,000	SE	0.4660	0.0664	0.0568	0.4658	0.0677	0.0613	0.1069
$N =$	Mean	-0.1654	1.0394	0.0147	-1.1668	0.0355	1.0180	0.8018
10,000	SE	0.3306	0.0454	0.0365	0.3310	0.0462	0.0377	0.0764
$N =$	Mean	-0.3711	1.0183	0.0058	-1.3718	0.0167	1.0068	0.8488
20,000	SE	0.2225	0.0297	0.0241	0.2223	0.0306	0.0253	0.0517
$N =$	Mean	-0.4853	1.0092	0.0044	-1.4854	0.0078	1.0051	0.8752
40,000	SE	0.1507	0.0202	0.0164	0.1508	0.0205	0.0171	0.0350
$N =$	Mean	-0.5400	1.0043	0.0019	-1.5401	0.0038	1.0020	0.8878
80,000	SE	0.1055	0.0136	0.0119	0.1055	0.0142	0.0125	0.0246

500 Monte Carlo runs; True values are reported in the parentheses next to the parameters; Other than the incorrect normalization, the model specifications in scenario (b) are the same as in scenario (a).

Table 2: Monte Carlo Scenarios (c) and (d)

Scenario (c): identification of constant $u(0, s)$											
		θ_0 (-4)	α_0^1 (0)	α_0^2 (0)	θ_1 (-1)	α_1^1 (1)	α_1^2 (0)	θ_2 (-2)	α_2^1 (0)	α_2^2 (1)	β (0.9)
$N =$	M	-4.3657	.	.	-1.0397	1.0951	0.0375	-2.0426	0.0876	1.0455	0.8054
5,000	SE	0.6194	.	.	0.1972	0.0871	0.0651	0.1973	0.0864	0.0675	0.0754
$N =$	M	-4.1536	.	.	-1.0003	1.0492	0.0171	-2.0019	0.0456	1.0207	0.8561
10,000	SE	0.2977	.	.	0.0548	0.0488	0.0392	0.0549	0.0498	0.0402	0.0492
$N =$	M	-4.0672	.	.	-0.9919	1.0224	0.0082	-1.9931	0.0204	1.0105	0.8771
20,000	SE	0.1597	.	.	0.0176	0.0316	0.0249	0.0181	0.0327	0.0257	0.0288
$N =$	M	-4.0276	.	.	-0.9912	1.0106	0.0034	-1.9914	0.0095	1.0044	0.8884
40,000	SE	0.1072	.	.	0.0097	0.0226	0.0181	0.0101	0.0228	0.0187	0.0188
$N =$	M	-4.0091	.	.	-0.9916	1.0055	0.0020	-1.9917	0.0052	1.0025	0.8945
80,000	SE	0.0723	.	.	0.0057	0.0141	0.0127	0.0059	0.0143	0.0130	0.0134
Scenario (d): identification of state-dependent $u(0, s)$											
		θ_0 (-1)	α_0^1 (0.5)	α_0^2 (-0.5)	θ_1 (-1)	α_1^1 (1)	α_1^2 (0)	θ_2 (-2)	α_2^1 (0)	α_2^2 (1)	β (0.9)
$N =$	M	-0.2051	0.6729	-0.8008	-0.1633	1.2244	-0.2203	-1.1729	0.2025	0.8036	0.6059
5,000	SE	1.0742	1.7614	1.6652	1.0608	1.6096	1.4993	1.0612	1.6101	1.5014	0.1488
$N =$	M	-0.6578	0.5705	-0.5246	-0.6387	1.1042	0.0126	-1.6435	0.0902	1.0277	0.7303
10,000	SE	0.3023	0.8841	0.9137	0.2910	0.7716	0.7895	0.2906	0.7717	0.7889	0.0968
$N =$	M	-0.8133	0.5270	-0.5378	-0.8043	1.0474	-0.0150	-1.8071	0.0408	0.9924	0.7969
20,000	SE	0.1600	0.5031	0.4870	0.1590	0.4293	0.4121	0.1590	0.4302	0.4135	0.0687
$N =$	M	-0.9209	0.5053	-0.5194	-0.9181	1.0167	-0.0070	-1.9196	0.0137	0.9957	0.8504
40,000	SE	0.0873	0.3540	0.3399	0.0897	0.2973	0.2853	0.0900	0.2974	0.2855	0.0472
$N =$	M	-0.9634	0.4940	-0.4992	-0.9618	1.0019	0.0062	-1.9624	0.0000	1.0080	0.8731
80,000	SE	0.0522	0.2092	0.2016	0.0541	0.1751	0.1672	0.0537	0.1755	0.1675	0.0307

500 Monte Carlo runs; True values are reported in the parentheses next to the parameters. Since precise estimation of $u(0, s)$ crucially hinges on data from the final period, we oversample individuals who reach the final period during estimation to ensure that we have enough observations to precisely estimate CCPs for the final period.

Table 3: Sample Size Requirements if True CCPs Are Known

Scenario (a)					θ_1 (-1)	α_1^1 (1)	α_1^2 (0)	θ_2 (-2)	α_2^1 (0)	α_2^2 (1)	β (0.9)
$N =$	M	.	.	.	-0.9766	1.0081	0.0033	-1.9766	0.0081	1.0033	0.8501
500	SE	.	.	.	0.0084	0.0113	0.0106	0.0084	0.0113	0.0106	0.0209
$N =$	M	.	.	.	-0.9851	1.0045	0.0017	-1.9851	0.0045	1.0017	0.8752
1,000	SE	.	.	.	0.0064	0.0081	0.0078	0.0064	0.0081	0.0078	0.0164
$N =$	M	.	.	.	-0.9887	1.0027	0.0006	-1.9887	0.0027	1.0006	0.8870
2,000	SE	.	.	.	0.0048	0.0055	0.0058	0.0048	0.0055	0.0058	0.0117
$N =$	M	.	.	.	-0.9908	1.0012	0.0003	-1.9908	0.0012	1.0003	0.8936
4,000	SE	.	.	.	0.0033	0.0042	0.0039	0.0033	0.0042	0.0039	0.0082
Scenario (d)											
		θ_0 (-1)	α_0^1 (0.5)	α_0^2 (-0.5)	θ_1 (-1)	α_1^1 (1)	α_1^2 (0)	θ_2 (-2)	α_2^1 (0)	α_2^2 (1)	β (0.9)
$N =$	M	-0.7412	0.5525	-0.5320	-0.7414	1.0671	-0.0207	-1.7414	0.0669	0.9794	0.7642
500	SE	0.0638	0.0803	0.0764	0.0635	0.0744	0.0703	0.0635	0.0745	0.0703	0.0272
$N =$	M	-0.8990	0.5197	-0.5137	-0.8991	1.0281	-0.0078	-1.8991	0.0280	0.9922	0.8288
1,000	SE	0.0381	0.0470	0.0478	0.0379	0.0437	0.0438	0.0379	0.0438	0.0438	0.0201
$N =$	M	-0.9401	0.5062	-0.5143	-0.9402	1.0120	-0.0084	-1.9402	0.0119	0.9917	0.8605
2,000	SE	0.0291	0.0337	0.0293	0.0290	0.0309	0.0270	0.0290	0.0309	0.0270	0.0159
$N =$	M	-0.9758	0.5083	-0.5127	-0.9759	1.0114	-0.0091	-1.9759	0.0114	0.9909	0.8785
4,000	SE	0.0202	0.0227	0.0216	0.0202	0.0209	0.0197	0.0202	0.0209	0.0197	0.0116

For scenario (a), which covers cases of truncated data, the parameters for the default utility are not estimated since it is impossible to identify them when the data do not extend to the final period. We instead assume that the research knows the default utility and imposes that knowledge during estimation. This is what we did in Table 1 as well.

Table 4: Counterfactual Analysis

Scenario (1): counterfactual of relaxed liquidity constraint			
	Actual	Prediction under $\dot{u}(0, s)$	Prediction under $\ddot{u}(0, s)$
% Default	35.29%	27.1%	13.55%
Scenario (2): counterfactual of static choice			
	Actual	Prediction under $\dot{u}(0, s)$	Prediction under $\ddot{u}(0, s)$
% Default	35.29%	50.24%	97.69%
Scenario (3): counterfactual of higher default costs			
	Actual	Prediction under $\dot{u}(0, s)$	Prediction under $\ddot{u}(0, s)$
% Default	35.29%	24.26%	24.26%

Scenario (1): Actual $\beta=0.9$; Counterfactual $\beta=0.95$; Correct $u(0, s) = -4 (\dot{u}(0, s))$; Incorrect $u(0, s) = -2 (\ddot{u}(0, s))$

Scenario (2): Actual $\beta=0.9$; Counterfactual $\beta=0$; Correct $u(0, s) = -4 (\dot{u}(0, s))$; Incorrect $u(0, s) = -2 (\ddot{u}(0, s))$

Scenario (3): Actual $\beta=$ Counterfactual $\beta=0.9$; Counterfactual $u(0, s) = u(0, s) - 0.5$; Correct $u(0, s) = -4 (\dot{u}(0, s))$; Incorrect $u(0, s) = -2 (\ddot{u}(0, s))$

Table 5: Overview of Key Differences between Our Estimator and Competing Estimators

Estimator	Objective function based on	Estimation based on*	Simplified representation	Construction of continuation value
This paper	diff. btw prediction and data on CCPs	OLS objective	Yes	Projection
HM1	diff. btw prediction and data on CCPs	nonlinear objective	No	Fwd simulation
HM2	diff. btw prediction and data on CCPs	nonlinear objective	Yes	Fwd simulation
BBL1	violation of inequality conditions	nonlinear objective	No	Fwd simulation
BBL2	violation of inequality conditions	nonlinear objective	Yes	Fwd simulation
BBL3	violation of inequality conditions	nonlinear objective	Yes	Projection

Comparison of our estimator to various adaptations of HM and BBL. By simplified representation, we mean expressing continuation value function using one-period-ahead choice probabilities only. * once CCPs and continuation value are constructed.

Table 6: Comparison of Mean Estimates

		θ_1 (-1)	α_1^1 (1)	α_1^2 (0)	θ_2 (-2)	α_2^1 (0)	α_2^2 (1)	β (0.9)
This paper	$N=5,000$	-0.9192	1.0739	0.0297	-1.9213	0.0686	1.0371	0.7067
	$N=40,000$	-0.9841	1.0076	0.0048	-1.9842	0.0070	1.0051	0.8729
	$N=80,000$	-0.9885	1.0046	0.0005	-1.9883	0.0040	1.0002	0.8881
BBL1(a)	$N=5,000$	-0.9530	1.0844	0.0798	-1.9234	0.1214	1.0501	0.8399
	$N=40,000$	-0.9763	1.0685	0.0240	-1.9638	0.0718	1.0161	0.8687
	$N=80,000$	-0.9763	1.0715	0.0228	-1.9696	0.0683	1.0194	0.8709
BBL1(b)	$N=5,000$	-0.8056	1.0533	0.0141	-1.7668	0.0889	0.9739	0.674
	$N=40,000$	-0.7514	1.0871	-0.0695	-1.7212	0.0728	0.9192	0.6117
	$N=80,000$	-0.7545	1.0743	-0.0807	-1.7264	0.0538	0.9093	0.6211
BBL2(a)	$N=5,000$	-1.0736	1.0073	0.0160	-1.9947	0.0977	0.9373	0.5240
	$N=40,000$	-1.0168	1.0039	0.0249	-1.9875	0.0334	0.9983	0.7706
	$N=80,000$	-1.0014	1.0106	0.0147	-1.9820	0.0278	0.9970	0.8193
BBL2(b)	$N=5,000$	-1.4604	0.9474	0.0058	-2.3750	0.0437	0.9195	0.0164
	$N=40,000$	-0.9249	1.0967	0.0531	-1.8816	0.1435	1.0088	0.0064
	$N=80,000$	-0.8324	1.1051	0.0666	-1.7978	0.1425	1.0301	0.0047
BBL3(a)	$N=5,000$	-1.0404	1.0615	-0.0124	-1.9511	0.1494	0.9018	0.4734
	$N=40,000$	-1.0099	1.0233	0.0166	-1.9775	0.0560	0.9876	0.7884
	$N=80,000$	-1.0059	1.0089	0.0249	-1.9866	0.0239	1.0072	0.8201
BBL3(b)	$N=5,000$	0.0980	0.9183	0.5438	-2.0684	0.0023	0.5254	0.2296
	$N=40,000$	0.0150	0.7946	0.2543	-2.6541	0.0031	0.5037	0.2492
	$N=80,000$	0.0171	0.7855	0.2275	-2.6354	0.0033	0.4692	0.2548
HM1(a)	$N=5,000$	-0.9891	0.9962	-0.0052	-1.9883	0.0062	0.9986	0.8958
	$N=40,000$	-0.9906	0.9900	-0.0041	-1.9890	-0.0010	0.9988	0.8964
	$N=80,000$	-0.9908	0.9880	-0.0056	-1.9893	-0.0023	0.9967	0.8969
HM1(b)	$N=5,000$	-0.9627	0.8584	0.0355	-1.8941	0.0225	0.7946	0.7795
	$N=40,000$	-0.9683	0.8612	0.0218	-1.8877	-0.0160	0.7808	0.7779
	$N=80,000$	-0.9739	0.8411	0.0281	-1.8984	-0.0325	0.7716	0.7915
HM2(a)	$N=5,000$	-0.9184	1.0845	0.0267	-1.9215	0.0787	1.0338	0.7004
	$N=40,000$	-0.9804	1.0113	0.0053	-1.9805	0.0099	1.0059	0.8627
	$N=80,000$	-0.9847	1.0064	0.0027	-1.9849	0.0059	1.0028	0.8753
HM2(b)	$N=5,000$	-0.9057	1.0099	0.0565	-1.8913	0.1045	0.9528	0.6727
	$N=40,000$	-0.9434	0.8502	0.0655	-1.9384	0.0152	0.9161	0.8080
	$N=80,000$	-0.9533	0.8475	0.0463	-1.9492	0.0006	0.8946	0.8341

BBL1(a)/BBL1(b): BBL1 with different sets of initial values. BBL2(a): BBL2 with alternative policy functions chosen very carefully through trial and error (which is possible since we know the true parameter values). BBL2(b): BBL2 with alternative policy functions prescribed in BBL (2007), where we add a random noise to each period's policy function. BBL3(a): BBL3 with alternative policy functions chosen very carefully through trial and error. BBL3(b): BBL3 with alternative policy functions prescribed in BBL (2007), where we add a random noise to each period's policy function. HM1(a)/HM1(b): HM1 with different sets of initial values. HM2(a)/HM2(b): HM2 with different sets of initial values. In BBL1, BBL2, HM1 and HM2, we use 250 simulation draws. Note that no simulation is used in BBL3 since projection is used instead of simulation. Refer to Table 5 for comparison among BBL1, BBL2, BBL3, HM1 and HM2.

Table 7: Comparison of Computational Time

This paper	$N=5,000$	87 minutes
	$N=40,000$	619 minutes
	$N=80,000$	1219 minutes
BBL1(a)	$N=5,000$	12160.61minutes
	$N=40,000$	87437.78 minutes
	$N=80,000$	175956.5 minutes
BBL1(b)	$N=5,000$	11115.28 minutes
	$N=40,000$	88833.23 minutes
	$N=80,000$	184819 minutes
BBL2(a)	$N=5,000$	1002.64 minutes
	$N=40,000$	8156.58 minutes
	$N=80,000$	15437.64 minutes
BBL2(b)	$N=5,000$	1124.63 minutes
	$N=40,000$	8774.82 minutes
	$N=80,000$	16098.25 minutes
BBL3(a)	$N=5,000$	507.85 minutes
	$N=40,000$	4313.02 minutes
	$N=80,000$	9036.67 minutes
BBL3(b)	$N=5,000$	520.61 minutes
	$N=40,000$	3213.9 minutes
	$N=80,000$	6458.53 minutes
HM1(a)	$N=5,000$	10126.47 minutes
	$N=40,000$	19618.66 minutes
	$N=80,000$	38433.91 minutes
HM1(b)	$N=5,000$	9490.94 minutes
	$N=40,000$	18484.75minutes
	$N=80,000$	37027.18 minutes
HM2(a)	$N=5,000$	639.6 minutes
	$N=40,000$	4680.26 minutes
	$N=80,000$	8899.53 minutes
HM2(b)	$N=5,000$	600.68 minutes
	$N=40,000$	5143.67 minutes
	$N=80,000$	10270.58 minutes

Table 8: Definitions of Variables

Static variables	Definition
Low Doc	= 1 if the loan was done with no or low documentation, = 0 otherwise.
Multiple Liens	= 1 if the borrower has other, junior mortgages, = 0 otherwise.
FICO	FICO score at loan origination, a credit score developed by Fair Isaac & Co. Scores range between 300 and 850. Higher scores indicate higher credit quality.
Payment	Monthly payment due.
Prepayment Penalty	= 1 if the loan has prepayment penalty, = 0 otherwise.
Income	Borrower's monthly income, imputed from the "front-end debt-to-income" ratio* and the monthly payment due.
MSA	Metropolitan Statistical Area of the loan property.
Time-varying variables	Definition
Housing Value	Current housing value, imputed by adjusting the appraised property value (at loan origination) by a home price index. Specifically, we use CoreLogic's zip code level home prices index, which is available month-by-month and is based on the transaction prices of properties that undergo repeat sales at different points in time.
Net Equity	Current housing value - Outstanding loan balance.
Market Rate	Current market interest rate available to the borrower for refinancing the loan. We impute a borrower-specific rate based on a benchmark rate plus a borrower-specific spread, the latter of which we can compute from the interest rate for the current loan.
Unemployment Rate	Monthly unemployment rate at the county level.

*Defined as the ratio of monthly mortgage-related payments to the borrower's income.

Table 9: Structural Estimates of Per-Period Utility

	Default	Prepay	Pay
FICO	-1.09 (0.126) ***		
MSA dummies	Included		
Housing Value		0.372 (0.144) ***	0.372 (0.144) ***
Monthly Payment		-0.079 (0.028) ***	-0.079 (0.028) ***
Prepayment Penalty		-0.103 (0.036) ***	
Income		-0.009 (0.006)	0.001 (0.001)
Unemployment Rate		-0.298 (0.015) ***	-0.013 (0.003) ***
Low Doc		-0.277 (0.062) ***	-0.026 (0.007) ***
Multiple Liens		-0.259 (0.052) ***	-0.016 (0.011) ***
β (coeff on $\hat{E}[V_{t+1}(s_{i,t+1}) s_{i,t}, a_{i,t}]$)		0.953 (0.016) ***	0.953 (0.016) ***
No. of Obs		478950	478950
R ²		0.9166	0.9986

Table 10: Counterfactual Predictions assuming Myopic Borrowers

Predictions for counterfactual scenario of $\beta = 0$		
	Using parameter estimates obtained when β is also estimated	Using parameter estimates obtained incorrectly imposing $\beta = 0.993$
% Default	3.66% (0.169)	0.05% (0.021)
% Prepay	57.91% (0.422)	58.98% (0.446)
% Pay	38.43% (0.44)	40.97% (0.45)
Predictions for actual scenario (i.e., within-sample predictions)		
	Using parameter estimates obtained when β is also estimated	Using parameter estimates obtained incorrectly imposing $\beta = 0.993$
% Default	23.66% (0.323)	23.96% (0.309)
% Prepay	57.12% (0.379)	56.99% (0.368)
% Pay	19.22% (0.354)	19.05% (0.353)

The first column reports predictions generated using the parameter estimates reported in Table 9. The second column reports predictions generated using parameter estimates obtained under an incorrect assumption that $\beta = 0.993$. The numbers reported inside the parentheses are standard errors. To compute the standard errors, we take the relevant point estimates, generate predictions for 500 samples which differ in the realization of draws for the state transition, and compute standard deviation across the predictions of the 500 samples.

Appendix 1: Optimal Policy Functions

Proposition 1 *Under Assumption 1 there exists a unique decision rule $D_t^*(s_t, \varepsilon_t)$ supported on A for $t = 1, 2, \dots, T$ that solves the maximization problem*

$$\sup_{(D_1, D_2, \dots, D_T) \in A^T} V_{1, \sigma}(s_1).$$

Proof:

Our argument uses backward induction. In the final period (at mortgage maturity) the borrower faces a static optimization problem of choosing among $V_T(0, s_T) + \varepsilon_{0,T}$, $V_T(1, s_T) + \varepsilon_{1,T}$, and $V_T(2, s_T) + \varepsilon_{2,T}$. The optimal decision delivers the highest payoff, yielding the decision rule $D_T^*(s_T, \varepsilon_T) = \arg \max_{k \in A} \{V_T(k, s_T) + \varepsilon_{k,T}\}$. Provided that the payoff shocks are idiosyncratic and have a continuous distribution, the optimal choice probabilities are characterized by continuous functions of $(V_T(k, s_T), k \in A)$. Knowing the optimal decision rule in period T , we can obtain the choice-specific value function in period $T - 1$ as

$$V_{T-1}(k, s_{T-1}) = u(k, s_{T-1}) + \beta E \left[\sum_{k' \in A} \mathbf{1}\{D_T^* = k'\} (V_T(k', s_T) + \varepsilon_{k',T}) \mid s_{T-1}, a_{T-1} = k \right].$$

Provided that the T^{th} period optimal decision has already been derived, the optimal decision problem in $T - 1$ becomes a static choice among three alternatives. Its solution, again, trivially exists and is (almost surely) unique because the distribution of ε_{T-1} is continuous. We iterate this procedure back to $t = 1$.

Appendix 2: Lemma 1

Under our assumptions, the system of equations

$$\sigma_0(z_1, z_2) = \bar{\sigma}_0,$$

$$\sigma_1(z_1, z_2) = \bar{\sigma}_1$$

has a unique solution if and only if $\bar{\sigma}_0 + \bar{\sigma}_1 < 1$.

This result generalizes that in Hotz and Miller (1993) to general full support distributions, and is also proved in Norets and Takahashi (2013). For completeness of exposition, we provide the proof.

Proof:

Consider partial derivatives

$$\begin{aligned}\frac{\partial \sigma_0(z_1, z_2)}{\partial z_1} &= - \int_{-\infty}^{+\infty} \frac{\partial^2 F_\varepsilon}{\partial \varepsilon_0 \partial \varepsilon_1}(\varepsilon_0, \varepsilon_0 - z_1, \varepsilon_0 - z_2) d\varepsilon_0, \\ \frac{\partial \sigma_0(z_1, z_2)}{\partial z_2} &= - \int_{-\infty}^{+\infty} \frac{\partial^2 F_\varepsilon}{\partial \varepsilon_0 \partial \varepsilon_2}(\varepsilon_0, \varepsilon_0 - z_1, \varepsilon_0 - z_2) d\varepsilon_0.\end{aligned}$$

Similarly, we can find that

$$\begin{aligned}\frac{\partial \sigma_1(z_1, z_2)}{\partial z_1} &= \int_{-\infty}^{+\infty} \frac{\partial^2 F_\varepsilon}{\partial \varepsilon_0 \partial \varepsilon_1}(z_1 + \varepsilon_1, \varepsilon_1, z_1 - z_2 + \varepsilon_1) d\varepsilon_1 \\ &\quad + \int_{-\infty}^{+\infty} \frac{\partial^2 F_\varepsilon}{\partial \varepsilon_1 \partial \varepsilon_2}(z_1 + \varepsilon_1, \varepsilon_1, z_1 - z_2 + \varepsilon_1) d\varepsilon_1 \\ &= \int_{-\infty}^{+\infty} \left(\frac{\partial^2 F_\varepsilon}{\partial \varepsilon_0 \partial \varepsilon_1} + \frac{\partial^2 F_\varepsilon}{\partial \varepsilon_1 \partial \varepsilon_2} \right) (\varepsilon_0, \varepsilon_0 - z_1, \varepsilon_0 - z_2) d\varepsilon_0,\end{aligned}$$

and

$$\frac{\partial \sigma_1(z_1, z_2)}{\partial z_2} = - \int_{-\infty}^{+\infty} \frac{\partial^2 F_\varepsilon}{\partial \varepsilon_1 \partial \varepsilon_2}(\varepsilon_0, \varepsilon_0 - z_1, \varepsilon_0 - z_2) d\varepsilon_0.$$

We assumed that the joint distribution of errors has a continuous density with a full support on $\mathbb{R}^3$. Provided that $\frac{\partial \sigma_0(z_1, z_2)}{\partial z_1} \frac{\partial \sigma_0(z_1, z_2)}{\partial z_2} > 0$ the mapping $z_1 \mapsto z_2$ implicitly defined by equation $\sigma_0(z_1, z_2) = \bar{\sigma}_0$ is invertible. Moreover, if we denote this mapping $z_2 = m_0(z_1, \bar{\sigma}_0)$, then using the result regarding the derivative of the inverse function, we can conclude that

$$\frac{\partial m_0(z_1, \bar{\sigma}_0)}{\partial z_1} \leq 0.$$

Similarly, we can define a map $z_2 = m_1(z_1, \bar{\sigma}_1)$, then using the result regarding the derivative of the inverse function, we can conclude that

$$\frac{\partial m_1(z_1, \bar{\sigma}_1)}{\partial z_1} \geq 0.$$

We can explore the asymptotic behavior of both maps. Consider m_0 first. Suppose that $z_1 \rightarrow -\infty$. Then $\lim_{z_1 \rightarrow -\infty} m_0(z_1, \bar{\sigma}_0) = z_2^*$, where z_2^* solves $\int \mathbf{1}\{\varepsilon_0 \geq z_2^* + \varepsilon_2\} F_\varepsilon(d\varepsilon) = \bar{\sigma}_0$. Also let z_1^* solve $\int \mathbf{1}\{\varepsilon_0 \geq z_1^* + \varepsilon_1\} F_\varepsilon(d\varepsilon) = \bar{\sigma}_0$. Then $\lim_{z_1 \rightarrow z_1^*} m_0(z_1, \bar{\sigma}_0) = -\infty$.

Next consider m_1 . Suppose that z_2^{**} is the solution of $\int \mathbf{1}\{\varepsilon_1 \geq z_2^{**} + \varepsilon_2\} F_\varepsilon(d\varepsilon) = \bar{\sigma}_1$. Then as $z_1 \rightarrow +\infty$, the map approaches asymptotically to the line: $m_1(z_1, \bar{\sigma}_1) \rightarrow z_1 + z_2^{**}$. Suppose that z_1^{**} is the solution of $\int \mathbf{1}\{z_1^{**} + \varepsilon_1 \geq \varepsilon_0\} F_\varepsilon(d\varepsilon) = \bar{\sigma}_1$. Then $\lim_{z_1 \rightarrow z_1^{**}} m_1(z_1, \bar{\sigma}_1) = -\infty$. Thus m_0 is a continuous strictly decreasing mapping from $(-\infty, z_1^*]$ into $(-\infty, z_2^*]$ and m_1 is a continuous strictly increasing mapping from $[z_1^{**}, +\infty)$ into the real line.

Provided that both curves are continuous and monotone, they intersect if and only if their projections on z_1 and z_2 axes overlap. The projections on the z_2 axis are guaranteed to overlap $((-\infty, z_2^*] \subset \mathbb{R})$. The projections on the z_1 axis will overlap if and only if $z_1^{**} < z_1^*$. Given that function $\sigma(z) = \int \mathbf{1}\{\varepsilon_0 - \varepsilon_1 \leq z\} F_\varepsilon(d\varepsilon)$ is strictly monotone in z , then $z_1^{**} < z_1^*$ if and only if $\bar{\sigma}_0 + \bar{\sigma}_1 < 1$.

This proves the statement of Lemma 1.

Appendix 3: Asymptotic Theory for the Plug-In Estimator

Section 3 outlined the structure of the two-step plug-in estimator for the structural parameters, which include the per-period payoffs and the discount factor. This appendix provides the asymptotic theory for the constructed estimator. We assume a parametric specification for the per-period utility, although our theory allows for an immediate extension to a nonparametric specification of the per-period utility. A key requirement of the plug-in semiparametric procedure is that the first-stage nonparametric estimator of the policy functions converge at a sufficiently fast rate. Our results for the consistency and the convergence rate of the first-stage estimator rely on the results in Wong and Shen (1995), Andrews (1991), and Newey (1997).

To assure consistency and a fast convergence rate for the first-stage estimator, we need the following assumption.

ASSUMPTION 2

- (i) *In addition to the Markov assumption (Assumption 1.iii), for each period t the distribution of states $s_t | s_{t-1}$ is identical across borrowers and over time, and the choice probabilities $\sigma_{k,t}(\cdot)$ are uniformly bounded from 0 and 1 for each $k = 0, 1, 2$. The state space $\mathcal{S}$ is compact.*
- (ii) *The eigenvalues of $E[q^L(s_t)q^{L'}(s_t) | a_t]$ are bounded away from zero uniformly over L , and $|q_l(\cdot)| \leq C$ for all l .*
- (iii) *$\frac{\sigma_{k,t}(s)}{\sigma_{0,t}(s)}$ belongs to a separable functional space with basis $\{q_l(\cdot)\}_{l=1}^\infty$. For each $t \leq T$ and $k \in \{1, 2\}$ the selected series terms provide a uniformly good approximation for the probability ratio*

$$\sup_{s \in \mathcal{S}} \left\| \log \frac{\sigma_{k,t}(s)}{\sigma_{0,t}(s)} - \text{proj} \left(\log \frac{\sigma_{k,t}(s)}{\sigma_{0,t}(s)} \mid q^L(\cdot) \right) \right\| = O(L^{-\alpha})$$

for some $\alpha \geq \frac{1}{2}$.

Assumption 2 can be verified for particular classes of polynomials and sieves (see Chen, 2007). Assumption 2 implies the following result establishing the consistency and convergence rate of the first-stage

estimator for the policy functions.³³

Theorem 3 *Under Assumptions 1 and 2, the estimator (4) is consistent uniformly over s :*

$$\sup_{s \in \mathcal{S}} \|\hat{\sigma}_{k,t}(s) - \sigma_{k,t}(s)\| = o_P\left(J^{-1/4}\right)$$

provided that $L \rightarrow \infty$ with $\frac{J}{L \log(J)} \rightarrow \infty$ as $J \rightarrow \infty$.

The asymptotics in this theorem is in terms of the number of loans J , reflecting the fact that each loan is observed only once for a given t . We use the estimated first-stage policy functions as inputs for the estimation of the second-stage structural parameters. Our approach is based on applying existing plug-in implementations for estimating the system of equations (5). These techniques involve constructing nonparametric elements based on a statistical model (in our case, the policy functions) that are then plugged into a fully parametric second step. Estimation in the second step is commonly performed by means of a weighted minimum distance procedure, with weights that are chosen optimally to maximize the efficiency of the resulting estimator.

To establish the asymptotic properties of the designed procedure we impose the following assumptions.

ASSUMPTION 3

- (i) *Parameter space Θ is a compact subset of $\mathbb{R}^p$.*
- (ii) *The per-period payoff is Lipschitz-continuous in parameters.*
- (iii) *The variance of the one-period-ahead policy function is bounded ($\sup_{s \in \mathcal{S}} E[\sigma_{k,t+1}(s_{t+1})^2 | s_t = s] < 1$) and strictly positive ($\inf_{s \in \mathcal{S}} E[\sigma_{k,t+1}(s_{t+1})^2 | s_t = s] > 0$) for any $t < T$.*

Under this assumption and the technical assumption described in Appendix 4, which restricts the complexity of the class of functions that is associated with our “nonparametric multinomial logit” estimator, we can use the results regarding semiparametric plug-in estimators in Ai and Chen (2003) and Chen, Linton and van Keilegom (2003), and establish the following result for the estimator for the second-stage structural parameters.

Theorem 4 *Under Assumptions 1, 2 and 3, the estimator (5) is consistent and has asymptotic normal distribution:*

$$\sqrt{JT^*} \left(\left(\hat{\theta}(a), \hat{\beta} \right) - (\theta_0(a), \beta) \right) \xrightarrow{d} N(0, V).$$

where variance V is determined by the functional structure of the model.

³³Proof is in Appendix 4.

The result of this theorem follows from Theorem 3.1 in Ai and Chen (2003). A significant difference between equations (5) used for our estimation and the conditional moment equations implied by infinite-horizon Markov dynamic decision processes is that the one-period-ahead values in our moment equations are estimated separately. As a result, the estimated choice-specific value function and the ex ante value function can be considered to be unrelated nonparametric objects (in contrast to infinite-horizon dynamics, in which the two are connected via a fixed point). This feature facilitates the evaluation of the asymptotic variance.

An explicit expression for the variance can be obtained as follows. We introduce

$$J_k(\sigma_{0,t}, \sigma_{1,t}, \sigma_{0,t+1}, \sigma_{1,t+1}, s) = \left(\frac{\partial F_k}{\partial \sigma_{0,t}}, \frac{\partial F_k}{\partial \sigma_{1,t}}, \frac{\partial F}{\partial \sigma_{0,t+1}}, \frac{\partial F}{\partial \sigma_{1,t+1}} \right)'$$

, $J(s) = (J_1(s), J_2(s))'$, and

$$M(s) = E \left[\begin{pmatrix} \frac{\partial u(s_t; \theta(1))}{\partial \theta(1)} & 0 & -\frac{\partial u(s_t; \theta(0))}{\partial \theta(0)} + \beta \frac{\partial u(s_{t+1}; \theta(0))}{\partial \theta(0)} & F(s_{t+1}) + u(s_{t+1}; \theta(0)) \\ 0 & \frac{\partial u(s_t; \theta(2))}{\partial \theta(2)} & -\frac{\partial u(s_t; \theta(0))}{\partial \theta(0)} + \beta \frac{\partial u(s_{t+1}; \theta(0))}{\partial \theta(0)} & F(s_{t+1}) + u(s_{t+1}; \theta(0)) \end{pmatrix} \middle| s_t = s \right],$$

as well as

$$\Omega(s) = \text{Var} \left((\hat{\sigma}_{0,t}(s_t), \hat{\sigma}_{1,t}(s_t), \hat{\sigma}_{0,t+1}(s_{t+1}), \hat{\sigma}_{1,t+1}(s_{t+1}))' \mid s_t = s \right).$$

Then, the variance of the second-stage estimates is determined by the sampling noise and the error from the first stage estimates:

$$V = E \left[M(s_t)^{-1} (E [Z_t' \text{Var}(\epsilon_t | s_t) Z_t \mid s_t] + J(s_t) \Omega(s_t) J(s_t)') M(s_t)^{-1'} \right].$$

As an alternative to using the asymptotic formula, we can use the subsampling approach to estimate the variance.

Appendix 4: Proof of Theorem 3

In this proof by n we denote the sample size corresponding to the borrowers observed with t periods from mortgage origination. We introduce the notation for the trinomial logit function $\ell(z_1, z_2) = \frac{\exp(z_1)}{1 + \exp(z_1) + \exp(z_2)}$. By $\tilde{\sigma}_{k,t}^L(s)$ we denote the choice probability

$$\tilde{\sigma}_{k,t}^L(s) = \ell(\tilde{r}^{L'}(t, k)q^L(s), \tilde{r}^{L'}(t, j)q^L(s)), \quad j \neq k$$

where $\tilde{r}^L(t, k)$ are the coefficients of the projection of the probability ratio $\log \frac{\sigma_{k,t}(s)}{\sigma_{0,t}(s)}$ on L first orthogonal polynomials. We also denote $\tilde{\sigma}_{0,t}^L(s) = 1 - \tilde{\sigma}_{1,t}^L(s) - \tilde{\sigma}_{2,t}^L(s)$. We note that $\frac{\partial \ell}{\partial z_1}, \frac{\partial \ell}{\partial z_2} \leq \frac{1}{2}$. Thus, $\frac{1}{2}$ is a uniform Lipschitz constant and

$$\begin{aligned} \sup_{s \in \mathcal{S}} |\tilde{\sigma}_{k,t}^L(s) - \sigma_{k,t}(s)| &= \sup_{s \in \mathcal{S}} \left| \ell \left(\log \frac{\tilde{\sigma}_{k,t}^L(s)}{\tilde{\sigma}_{0,t}^L(s)}, \log \frac{\tilde{\sigma}_{j,t}^L(s)}{\tilde{\sigma}_{0,t}^L(s)} \right) - \ell \left(\log \frac{\sigma_{k,t}(s)}{\sigma_{0,t}(s)}, \log \frac{\sigma_{j,t}(s)}{\sigma_{0,t}(s)} \right) \right| \\ &\leq \frac{1}{2} \sqrt{\sup_{s \in \mathcal{S}} \left| \log \frac{\tilde{\sigma}_{1,t}^L(s)}{\tilde{\sigma}_{0,t}^L(s)} - \log \frac{\sigma_{1,t}(s)}{\sigma_{0,t}(s)} \right|^2 + \sup_{s \in \mathcal{S}} \left| \log \frac{\tilde{\sigma}_{2,t}^L(s)}{\tilde{\sigma}_{0,t}^L(s)} - \log \frac{\sigma_{2,t}(s)}{\sigma_{0,t}(s)} \right|^2} = O(L^{-\alpha}) \end{aligned}$$

for some $\alpha \geq \frac{1}{2}$. This guarantees the quality of approximation of the choice probability using a logit transformation of the series expansion.

Now we omit index t in the variables (whenever the period of time under consideration is known), use r_k^L in place of $r^L(t, k)$, and construct function

$$\begin{aligned} \rho(a, s; r_1^L, r_2^L) &= (\mathbf{1}\{a = 1\} - \mathbf{1}\{a = 0\}) \ell(r_1^{L'} q^L(s), r_2^{L'} q^L(s)) \\ &\quad + (\mathbf{1}\{a = 2\} - \mathbf{1}\{a = 0\}) \ell(r_2^{L'} q^L(s), r_1^{L'} q^L(s)). \end{aligned}$$

Then we can express the sample quasi-likelihood as

$$\widehat{Q}(r_1^L, r_2^L) = E_n [\rho(a, s; r_1^L, r_2^L)] + E_n [\mathbf{1}\{a = 0\}],$$

where we adopted the notation from the empirical process theory where $E_n[\cdot] = \frac{1}{n} \sum_{i=1}^n$. Also introduce the population likelihood with the series expansion

$$Q(r_1^L, r_2^L) = E [\rho(a, s; r_1^L, r_2^L)] + E [\mathbf{1}\{a = 0\}].$$

Consider function

$$f(a, s; r_1^L, r_2^L, \tilde{r}_1^L, \tilde{r}_2^L) = \rho(a, s; r_1^L, r_2^L) - \rho(a, s; \tilde{r}_1^L, \tilde{r}_2^L) - E[\rho(a, s; r_1^L, r_2^L)] + E[\rho(a, s; \tilde{r}_1^L, \tilde{r}_2^L)].$$

Provided that we established that function $\ell(\cdot, \cdot)$ is Lipschitz, we can evaluate

$$\text{Var}(f(a, s; r_1^L, r_2^L, \tilde{r}_1^L, \tilde{r}_2^L)) = O \left(L \sup_{k=1,2, l \leq L} \|r_{k,l} - \tilde{r}_{k,l}\| \right) = O(L).$$

where $r_k^L = (r_{k,1}, \dots, r_{k,L})$. Next we impose a technical assumption that allows us to establish consistency

of estimator (4).

ASSUMPTION 4

Consider the class of functions indexed by n

$$\mathcal{F}_n = \left\{ f(\cdot, \cdot; r_1^{L_n}, r_2^{L_n}, \tilde{r}_1^{L_n}, \tilde{r}_2^{L_n}) - E[f(\cdot, \cdot; r_1^{L_n}, r_2^{L_n}, \tilde{r}_1^{L_n}, \tilde{r}_2^{L_n})], r_{k,l} \in \Theta, l \leq L_n, k = 1, 2 \right\},$$

where Θ is the compact subset of $\mathbb{R}$ and $\tilde{r}_1^{L_n}$ and $\tilde{r}_2^{L_n}$ are the coefficients of projections of population probability ratios on L_n series terms. Then for each $L_n \rightarrow \infty$ such that $n/(L_n \log n) \rightarrow \infty$ the $\mathbf{L}_1$ covering number for class $\mathcal{F}_n$, N , has the following bound

$$\log N(\delta, \mathcal{F}_n, \mathbf{L}_1) \leq An^{r_0} \log \frac{1}{\delta},$$

where $0 < r_0 \leq \frac{3}{4}$ and $r_0 \downarrow 0$ is assumed to correspond to the factor $\log n$.

This is the condition restricting the complexity of the functions created by logit transformations of series expansions. By construction any $f \in \mathcal{F}_n$ is bounded $|f| < 1 < \infty$. We established that $\text{Var}(f) = O(L_n)$ for $f \in \mathcal{F}_n$. The symmetrization inequality (30) in Pollard (1984) holds if $\varepsilon_n/(16n\mu_n^2) \leq \frac{1}{2}$. This will occur if $\frac{n\varepsilon_n}{N^2} \rightarrow \infty$. Provided that the symmetrization inequality holds, we can follow the steps of Theorem 37 in Pollard (1984) to establish the tail bound on the deviations of the sample average of f via a combination of the Hoeffding inequality and the covering number for the class $\mathcal{F}_n$. As a result, we obtain that

$$\begin{aligned} & P\left(\sup_{f \in \mathcal{F}_n} \frac{1}{n} \left\| E_n[f(\cdot)] \right\| > 8\mu_n\right) \\ & \leq 2 \exp\left(An^{r_0} \log \frac{1}{\mu_n}\right) \exp\left(-\frac{1}{128} \frac{n\mu_n^2}{L_n}\right) + P\left(\sup_{f \in \mathcal{F}_n} \frac{1}{n} \left\| E_n[f(\cdot)]^2 \right\| > 64L_n\right). \end{aligned}$$

The second term can be evaluated with the aid of Lemma 33 in Pollard (1984):

$$P\left(\sup_{f \in \mathcal{F}_n} \frac{1}{n} \left\| E_n[f(\cdot)]^2 \right\| > 64L_n\right) \leq 4 \exp\left(An^{2r_0} \log \frac{1}{L_n}\right) \exp(-nL_n).$$

As a result, we find that

$$\begin{aligned} & P\left(\sup_{f \in \mathcal{F}_n} \frac{1}{n} \left\| E_n[f(\cdot)] \right\| > 8\mu_n\right) \\ & \leq 2 \exp\left(An^{r_0} \log \frac{1}{\mu_n}\right) \exp\left(-\frac{1}{128} \frac{n\mu_n^2}{L_n}\right) + 4 \exp\left(An^{r_0} \log \frac{1}{L_n} - nL_n\right). \end{aligned}$$

We start the analysis with the first term. Consider the case with $r_0 > 0$. Then the log of the first term takes the form

$$An^{r_0} \log(1/\mu_n) - \frac{1}{128} \frac{n\mu_n^2}{L_n}.$$

Then one needs that $\frac{n\mu_n^2}{L_n n^{r_0} \log n} \rightarrow \infty$ if $r_0 > 0$ and $\frac{n\mu_n^2}{L_n \log^2 n} \rightarrow \infty$ if $r_0 \downarrow 0$. Hence the first term is of $o(1)$. This condition also guarantees that the second term vanishes. We note also that the CLT applies to the term $E_n[\mathbf{1}\{a = 0\}] = E[\mathbf{1}\{a = 0\}] + O_p(\frac{1}{\sqrt{n}})$. Now for some slowly diverging sequence $\delta_n \rightarrow \infty$ such that $\mu_n = \delta_n \sqrt{\frac{L_n n^{r_0} \log n}{n}} \rightarrow 0$, we establish that

$$\sup_{(r_1^{L_n}, r_2^{L_n}) \in \Theta^{L_n} \times \Theta^{L_n}} \left\| \widehat{Q}(r_1^L, r_2^L) - Q(r_1^L, r_2^L) + \widehat{Q}(\tilde{r}_1^L, \tilde{r}_2^L) - Q(\tilde{r}_1^L, \tilde{r}_2^L) \right\| = O_p\left(\mu_n + \frac{1}{\sqrt{n}}\right) = o_p(1).$$

Thus, the sample quasi-likelihood converges uniformly to the population quasi-likelihood and the estimated choice probabilities are uniformly consistent over $\mathcal{S}$. To establish the rate for the estimated choice probabilities, we consider a neighborhood of the population projections defined by $\sup_{k=1,2, l \leq L_n} \|r_{k,l} - \tilde{r}_{k,l}\| \leq \varepsilon$. Using Lemma 2.3.1 from van der Vaart and Wellner (1996), we can find that

$$E \left[\sup_{f \in \mathcal{F}_n} \sqrt{n} E_n[f(\cdot)] \right] \leq C n^{r_0/2} \sqrt{L_n} \varepsilon \log \frac{1}{\sqrt{L_n} \varepsilon},$$

for some constant C . Using Theorem 3.4.1 from van der Vaart and Wellner (1996) and the derived inequality, we can express the convergence rate for the estimated parameters of the approximated choice probabilities as $\rho_n^2 n^{r_0/2} \sqrt{L_n} \frac{1}{\rho_n} \log \frac{\rho_n}{\sqrt{L_n}} \leq \sqrt{n}$. Then

$$\sup_{s \in \mathcal{S}} \|\widehat{\sigma}_{k,t}(s) - \tilde{\sigma}_{k,t}(s)\| = O_p\left(\frac{L_n}{\rho_n}\right).$$

To attain the rate $o_p(n^{-1/4})$ we need to assure that $\frac{L_n}{\rho_n} = o(n^{-1/4})$. To assure $n^{-1/4}$ we choose $\delta_n \rightarrow 0$ and set $L_n = \delta_n n^{-1/4} \rho_n$. Then the rate constraint can be re-written as

$$\rho_n^2 n^{-3/8+r_0/2} \frac{\log \frac{n^{1/4} \sqrt{\rho_n}}{\sqrt{\delta_n}}}{\frac{n^{1/4} \sqrt{\rho_n}}{\sqrt{\delta_n}}} \leq \sqrt{n}.$$

Provided that $\lim_{x \rightarrow \infty} \log x / x = 0$, we conclude that $\rho_n = O(n^{7/8-r_0/2})$, meaning that $L_n = o(n^{5/8-r_0/2})$.

We note that the slowest rate for the choice of L_n has to satisfy

$$\frac{n^{1-r_0} \mu_n^2}{L_n \log n} \rightarrow \infty,$$

for $\mu_n \rightarrow 0$. Thus, the estimator with rate $o(n^{-1/4})$ is plausible if $r_0 < 3/4$. Using the triangle inequality and our previous result, we find that

$$\sup_{s \in \mathcal{S}} \|\widehat{\sigma}_{k,t}(s) - \sigma_{k,t}(s)\| = O_p\left(\frac{L_n}{\rho_n} + L_n^{-\alpha}\right) = o_p\left(n^{-1/4}\right),$$

if $\alpha \geq 1$.

Online Appendix: Tables for Monte Carlo Exercises

Table A1: Comparison of Mean Bias

		θ_1	α_1^1	α_1^2	θ_2	α_2^1	α_2^2	β
This paper	$N=5,000$	5.7123	5.2231	2.0974	5.5658	4.8489	2.6234	-13.672
	$N=40,000$	3.1747	1.5297	0.9533	3.1542	1.3907	1.0178	-5.4287
	$N=80,000$	3.2628	1.2910	0.1332	3.3087	1.1319	0.0574	-3.3584
BBL1(a)	$N=5,000$	3.3247	5.9672	5.6439	5.4135	8.5817	3.5415	-4.2469
	$N=40,000$	4.7362	13.6958	4.7977	7.2334	14.3583	3.2132	-6.2645
	$N=80,000$	6.7146	20.2232	6.4603	8.5880	19.3171	5.4794	-8.2311
BBL1(b)	$N=5,000$	13.7438	3.7704	0.9994	16.4871	6.2838	-1.8439	-15.9367
	$N=40,000$	49.7180	17.4173	-13.8947	55.7520	14.5511	-16.1525	-57.6659
	$N=80,000$	69.4362	21.0092	-22.8156	77.3846	15.2185	-25.6514	-78.8891
BBL2(a)	$N=5,000$	-5.2047	0.5186	1.1318	0.3778	6.9057	-4.4351	-26.5861
	$N=40,000$	-3.3558	0.7880	4.9877	2.5082	6.6755	-0.3417	-25.8874
	$N=80,000$	-0.3934	2.9986	4.1463	5.1014	7.8644	-0.8551	-22.8384
BBL2(b)	$N=5,000$	-32.5547	-3.7195	0.4073	-26.5160	3.0902	-5.6922	-62.4828
	$N=40,000$	15.0168	19.3482	10.6168	23.6887	28.6986	1.7645	-178.7200
	$N=80,000$	47.4134	29.7298	18.8310	57.1967	40.2921	8.5133	-253.2178
BBL3(a)	$N=5,000$	-2.8593	4.3519	-0.8784	3.4580	10.5634	-6.9455	-30.1626
	$N=40,000$	-1.9735	4.6534	3.3207	4.4901	11.1902	-2.4893	-22.3120
	$N=80,000$	-1.6668	2.5248	7.0449	3.7844	6.7554	2.0464	-22.5950
BBL3(b)	$N=5,000$	77.6410	-5.7795	38.4520	-4.8360	0.1639	-33.5603	-47.4055
	$N=40,000$	202.9911	-41.0756	50.8638	-130.8136	0.6153	-99.2623	-130.1522
	$N=80,000$	287.6685	-60.6680	64.3482	-179.7309	0.9400	-150.1427	-182.4968
HM1(a)	$N=5,000$	0.7673	-0.2665	-0.3653	0.8276	0.4383	-0.0972	-0.2985
	$N=40,000$	1.8873	-1.9989	-0.8245	2.2081	-0.1935	-0.2341	-0.7217
	$N=80,000$	2.6138	-3.3833	-1.5720	3.0180	-0.6636	-0.9218	-0.8861
HM1(b)	$N=5,000$	2.6356	-10.0098	2.5081	7.4878	1.5900	-14.5250	-8.5222
	$N=40,000$	6.3431	-27.7619	4.3553	22.4689	-3.1927	-43.8472	-24.4259
	$N=80,000$	7.3727	-44.9299	7.9406	28.7389	-9.1993	-64.5924	-30.6784
HM2(a)	$N=5,000$	5.7677	5.9750	1.8848	5.5527	5.5629	2.3921	-14.1129
	$N=40,000$	3.9275	2.2609	1.0510	3.9023	1.9719	1.1809	-7.4517
	$N=80,000$	4.3147	1.8137	0.7558	4.2795	1.6610	0.7971	-6.9760
HM2(b)	$N=5,000$	6.6683	0.7017	3.9979	7.6851	7.3898	-3.3351	-16.0742
	$N=40,000$	11.3281	-29.9700	13.0902	12.3150	3.0302	-16.7757	-18.3927
	$N=80,000$	13.2031	-43.1293	13.0823	14.3737	0.1563	-29.8175	-18.6369

The mean biases reported in this table are defined as $\sqrt{N}$ · mean of $(\theta_k - \theta_0)$ over Monte Carlo simulations $k = 1, 2, \dots, 500$. We multiply the mean bias by $\sqrt{N}$, where N is the sample size, and report the resulting product. This adjustment makes it easier to make comparisons across different sample sizes, because the unadjusted statistic collapses at a rate of $\sqrt{N}$.

Table A2: Comparison of Standard Error

		θ_1	α_1^1	α_1^2	θ_2	α_2^1	α_2^2	β
This paper	$N=5,000$	2.9934	4.4992	4.3495	3.0771	4.4331	4.6899	7.8599
	$N=40,000$	2.0899	4.0606	3.3104	2.1849	3.9075	3.5386	6.5530
	$N=80,000$	1.8444	4.0907	3.1197	1.8940	4.1848	3.1947	6.3077
BBL1(a)	$N=5,000$	6.9315	20.0076	17.7303	8.4466	20.3803	17.8651	8.4827
	$N=40,000$	3.2431	17.1204	16.8672	5.2253	18.1894	17.2230	4.1949
	$N=80,000$	3.4598	18.2092	17.7524	5.9360	20.2625	18.4132	4.5786
BBL1(b)	$N=5,000$	25.9901	26.6159	25.9271	27.3423	27.4646	25.7705	23.0431
	$N=40,000$	78.6046	49.8145	50.9073	84.3849	56.7870	52.8406	77.2953
	$N=80,000$	116.1301	72.9106	71.8719	125.3301	86.9449	75.5910	110.2835
BBL2(a)	$N=5,000$	27.4869	28.7982	25.3800	27.9265	29.6981	26.7321	44.9152
	$N=40,000$	35.8209	35.9279	30.1506	36.9839	36.0391	30.4802	65.0968
	$N=80,000$	38.3888	35.3436	30.9947	39.4026	35.6309	30.8483	72.8766
BBL2(b)	$N=5,000$	49.1523	35.7521	27.5511	49.8783	36.6442	28.4547	1.0542
	$N=40,000$	65.8194	47.8936	32.6117	67.0132	49.9498	32.4332	1.1813
	$N=80,000$	59.5477	45.6975	32.1448	60.7212	48.4547	31.5765	1.0578
BBL3(a)	$N=5,000$	28.4845	28.9296	26.6105	29.4448	30.3713	27.8174	42.2590
	$N=40,000$	36.4058	37.2834	27.3399	37.6668	38.2676	27.7458	62.1977
	$N=80,000$	41.0776	38.9539	32.0551	41.9036	39.9243	32.1995	75.0387
BBL3(b)	$N=5,000$	35.8480	31.2841	41.9894	160.5501	0.4929	52.3559	20.3130
	$N=40,000$	23.2295	21.4013	23.4473	168.3990	0.9400	96.9919	16.5323
	$N=80,000$	24.1200	20.5093	23.3198	225.8098	1.1898	119.7305	17.1297
HM1(a)	$N=5,000$	0.8850	3.7010	3.2449	1.0047	3.8363	3.6101	1.2905
	$N=40,000$	1.7708	7.7385	6.6008	2.0621	7.8729	7.0390	2.5095
	$N=80,000$	1.7696	7.5472	6.1297	1.9500	7.8639	6.5910	2.5985
HM1(b)	$N=5,000$	9.3842	15.6511	14.5354	10.0093	16.7779	17.6536	29.1480
	$N=40,000$	25.9509	49.1380	39.6466	29.0439	46.2442	51.1704	82.1106
	$N=80,000$	33.3866	65.3914	54.0216	35.9447	63.8271	67.3285	111.9115
HM2(a)	$N=5,000$	2.8498	4.7309	4.0209	2.8726	4.8100	4.3537	7.4089
	$N=40,000$	2.2062	4.0578	3.2891	2.2742	4.1127	3.4403	6.7965
	$N=80,000$	2.0686	3.8624	3.3742	2.1213	4.0462	3.5335	6.7775
HM2(b)	$N=5,000$	10.4843	22.1065	16.8111	10.2904	19.0239	18.8785	20.7287
	$N=40,000$	31.9716	81.6671	55.4017	31.0625	65.4898	65.0559	69.6227
	$N=80,000$	48.5300	102.1482	80.8873	47.2243	96.7825	94.7463	105.2789

The standard errors reported in this paper are defined as $\sqrt{N} \cdot \sqrt{(\text{mean of } (\theta_k - \text{mean}(\theta_k))^2)}$ over Monte Carlo simulations $k = 1, 2, \dots, 500$. We multiply the standard error by $\sqrt{N}$, where N is the sample size, and report the resulting product. This adjustment makes it easier to make comparisons across different sample sizes, because the unadjusted statistic collapses at a rate of $\sqrt{N}$.