A Dynamic Discrete Choice Model of Reverse Mortgage Borrower Behavior

Jason R. Blevins[1], Wei Shi[2], Donald R. Haurin[1], and Stephanie Moulton[3]

[1]*Department of Economics, The Ohio State University*
[2]*Institute for Economic and Social Research, Jinan University*
[3]*John Glenn College of Public Affairs, The Ohio State University*

April 9, 2020*
Forthcoming at *International Economic Review*

Abstract. Using unique data on reverse mortgage borrowers in the Home Equity Conversion Mortgage (HECM) program, we semiparametrically estimate a dynamic discrete choice model of borrower behavior. Our estimator is based on a new identification result we develop for models with multiple terminating actions. We show that the per-period utility functions and discount factor are identified without restrictive, ad hoc identifying restrictions that lead to incorrect counterfactual implications. Our estimates provide insights about factors that influence HECM refinance, default, and termination decisions and allow us to quantify the trade-offs involved for proposed program modifications, such as income and credit requirements.

Keywords: dynamic discrete choice, reverse mortgages, identification, semiparametric estimation, microeconometrics.

JEL Classification: C14, C25, C61, G21, R21.

Acknowledgments: The authors acknowledge funding from The MacArthur Foundation, "Aging in Place: Analyzing the Use of Reverse Mortgages to Preserve Independent Living," 2012–14, Stephanie Moulton, PI, and also from the Department of Housing and Urban Development, "Aging in Place: Managing the Use of Reverse Mortgages to Enable Housing Stability," 2013–2015, Stephanie Moulton, PI. Shi acknowledges financial support from the National Natural Science Foundation of China under grant no. 71803062 and the 111 Project of China under grant no. B18026. We also thank Yonghong An, Peter Arcidiacono, Thomas Davidoff, Juan Carlos Escanciano, Paul Grieco, Phil Haile, Sukjin Han, Ashley Langer, Lance Lochner, Salvador Navarro, Joris Pinkse, David Rivers, Mark Roberts, Joseph Rossetti, Daniel Sacks, Kamila Sommer, Mauricio Varela, Tiemen Woutersen, and Ruli Xiao, attendees of the 2016 Calgary Empirical Microeconomics conference, the 2017 AREUEA-ASSA Annual Meeting, 2017 Midwest Econometrics Group meeting, and the 2018 Asian Meeting of the Econometric Society as well as seminar participants at Indiana University, Ohio State Glenn College of Public Affairs, the Office of the Comptroller of the Currency, Pennsylvania State University, University of Arizona, University of Maryland, University of Texas at Austin, and University of Western Ontario for useful comments.

Disclaimer: The work that provided the basis for this publication was supported by funding under a grant with the U.S. Department of Housing and Urban Development. The substance and findings of the work are dedicated to the public. The author and publisher are solely responsible for the accuracy of the statements and interpretations contained in this publication. Such interpretations do not necessarily reflect the view of the Government.

*Original draft: January 31, 2017.

1. Introduction

Home Equity Conversion Mortgage (HECM) loans are federally-insured reverse mortgages backed by the Federal Housing Administration (FHA). The program is designed to help older homeowners age in place by allowing them to access home equity without making monthly payments, with payment of the loan being deferred until the loan is terminated.

Using a unique dataset of HECM borrowers from 2007–2014, we estimate borrowers' utility functions and investigate the implications of various counterfactual scenarios and policy changes on HECM outcomes and borrower welfare. Based on our estimates of HECM borrowers' preference parameters, on average borrowers value HECMs more when they are younger, have less access to revolving credit, or have less net equity (higher outstanding HECM balances relative to the value of their home). They also tend to value the program more when they have higher income, when interest rates are higher, or when housing prices have recently declined. Variations in these variables over time affect how much HECM borrowers value their HECM loans and their decisions to terminate.

The decisions of HECM borrowers to default, terminate, or refinance are inherently dynamic. Terminations are of particular interest because HECM loans are non-recourse loans insured by the FHA. This insurance provides borrowers with a put option which, along with other dynamic considerations, determines when borrowers choose to terminate the loan. Accurately predicting such terminations is important for evaluating the solvency of the FHA's Mutual Mortgage Insurance Fund (MMIF), which pays lenders when mortgagors default.

Naturally, policymakers are interested in reducing adverse terminations and defaults and have enacted participation constraints in the form of initial credit and income requirements. We simulate our estimated model under these requirements in order to evaluate their effects on both loan outcomes and borrower welfare. Our simulations indicate that these policies would indeed decrease default rates and would also lower the fraction of households with negative net equity. The welfare cost is that households with higher than average valuations for the program would be excluded.

Our results complement other recent attempts of using dynamic models to understand how households value reverse mortgages. Nakajima and Telyukova (2017) calibrate a life-cycle model of retirement and use it to analyze the ex-ante welfare gain from reverse mortgages. Davidoff (2015) simulates the value of the put option minus the initial costs and fees in order to estimate a lower bound on the net present value of HECMs to households. He argues that, contrary to a commonly held belief, "high costs" cannot explain weak HECM demand.

In contrast to these studies, our valuations are estimated from the revealed preferences and observed characteristics of borrowers over time in combination with an econometric model of their dynamic decision making behavior. Our approach is based on methods that have been widely used in economics since the pioneering work of authors such as Miller (1984), Wolpin (1984), Pakes (1986), Rust (1987), Hotz and Miller (1993), and Keane and Wolpin (1994, 1997). In housing economics specifically, structural dynamic discrete choice models have formed the methodological basis of recent studies on forward mortgage default by Bajari, Chu, Nekipelov, and Park (2016) (henceforth BCNP), Ma (2014), and Fang, Kim, and Li (2016) as well as work on neighborhood choice by Bayer, McMillan, Murphy, and Timmins (2016).

Our work is also related to the study of reverse mortgage termination and default. Davidoff and Welke (2007) found that HECM borrowers have a high rate of termination and attribute that to selection on mobility and high sensitivity to house price changes. Given the high rates of termination, accurately predicting terminations is important for the HECM program. In an effort to improve assessments of HECM loan performance, Szymanoski, Enriquez, and DiVenti (2007) estimate HECM termination hazards by age and borrower type.

In addition to termination, HECM borrowers can default for not paying property taxes or homeowner insurance premiums. Moulton, Haurin, and Shi (2015) identify the factors that predict default, including borrower credit characteristics and the amount of the initial withdrawal on the HECM. Our work contributes to this area of the literature in that our model allows us to predict rates of termination, tax and insurance default, and refinancing at the borrower-year level, and thus to examine how these rates vary with individual borrower characteristics. Based on the estimated structural parameters, we obtain values from the ex-ante value function which provides a measure on how much an HECM borrower values the HECM loan, and we identify which borrower characteristics are associated with a higher value.

Motivated by the institutional features of the HECM program on which our model is based, we develop a new semiparametric identification result for the household utility function and discount factor in our model which does not require assuming that the functional form of utility is known for one choice.[1] In particular, our model has two distinct, observable terminating actions which allow us to identify the period payoff

[1] In light of work by Aguirregabiria (2005, 2010), Bajari, Hong, and Nekipelov (2013), Norets and Tang (2014), Aguirregabiria and Suzuki (2014), Arcidiacono and Miller (2015), Chou (2016), and Kalouptsidi, Scott, and Souza-Rodrigues (2016), it is now well known in the literature that using an incorrect functional form for one choice as an identifying restriction on utility (i.e., a zero normalization) can lead to bias in conditional choice probability estimates for counterfactuals and also welfare predictions, except in special cases. By developing a model where the full utility function is identified and estimable, our analysis avoids these pitfalls. Additionally, work by Magnac and Thesmar (2002), Chung, Steenburgh, and Sudhir (2014), Fang and Wang (2015), BCNP, Komarova, Sanches, Silvia Junior, and Srisuma (2016), Mastrobuoni and Rivers (2016), and Abbring and Daljord (2018) underscores the importance of estimating time preferences.

functions for all choices. Our approach generalizes to other single-agent models with multiple terminating actions under conditions we discuss in Section 3 and formalize in Appendix A. We estimate the model using a multi-step plug-in semiparametric approach inspired by that of BCNP.[2] This approach is simple and computationally tractable: it does not require solving a nested dynamic programming problem, forward simulation, backwards induction, or optimization of difficult functions. As such, it can be implemented using built-in commands in most statistical packages.

In contrast to previous work by BCNP, Arcidiacono and Miller (2015), Chou (2016) and others, our identification result is applicable in cases where the utility function itself is also of interest (not only counterfactual implications) and when the utility function may be nonlinear. Furthermore, our approach is valid when the final decision period is not necessarily observed or when an appropriate exclusion restriction may not be available. Full identification of the utility function also implies identification of all types of counterfactuals including non-additive and non-linear changes in utilities and changes in transition probabilities. Yet, these broad classes of counterfactuals are problematic when an ad hoc utility assumption is imposed in order to estimate the model (Kalouptsidi et al., 2016).

2. A Model of HECM Borrower Behavior

We begin with some institutional details of the HECM program and then develop a structural, dynamic discrete choice model for households that have or are considering a HECM.

To obtain a HECM a borrower must be 62 years of age or older. The home must be the borrower's principal residence and must be either a single-family home or part of a 2–4 unit dwelling. Potential borrowers must also complete a mandatory counseling session with a HUD approved counseling agency. During our sample period, there were no income or credit requirements for HECM borrowers, although such requirements have since been enacted and are among the counterfactual policy changes we consider in this paper.[3] The amount one can borrow, known as the *principal limit*, is determined by the age of the youngest borrower, the appraised value of the home up to the FHA mortgage limit, and the interest rate. HECMs are non-recourse loans, meaning that borrowers will

[2] We show that the approach of BCNP for identifying the discount factor, based on nonstationarity of the conditional choice probabilities, is valid in our model as well, and we estimate the discount factor as part of our analysis.

[3] Initial disbursement limits on HECMs were enacted by HUD, effective for all loans originated (with case numbers assigned) on or after September 30, 2013 (Mortgagee Letter 2013-27). HUD's requirement for a financial assessment became effective for all HECM loans originated (with case numbers assigned) after April 27, 2015 (Mortgagee Letter 2015-06).

never owe more than the loan balance or 95% of the current appraised value of the home, whichever is lower. Borrowers cannot be compelled to use assets other than the property to repay the debt.

Our model covers decisions related to both HECM take-up and HECM outcomes.[4] Figure 1 summarizes the decisions households make in our model. Households in the model choose whether to take up HECMs, and if they do, what types of HECMs. During our sample period, borrowers could choose between fixed- (FRM) and adjustable rate (ARM) HECMs. Fixed-rate HECM borrowers receive the entire principal limit in an upfront lump sum payment.[5] On the other hand, borrowers with adjustable-rate HECMs have more payment disbursement options. They may, for example, choose to make only a partial withdrawal initially and later make unscheduled withdrawals or receive payments in scheduled installments. Note that some borrowers with adjustable rate HECMs still utilize a large amount of credit (defined as more than 80% of available credit) upon loan closing.[6] The choices of FRM or ARM and the amount of upfront credit utilization have important implications for later years. The unused portion of the credit line grows at the same rate as being charged on the loan balance which equals the interest rate plus the mortgage insurance premium, and can be tapped to fulfill future cash needs. Several important choices are observed for an HECM household, including termination, refinance into another HECM, default on property tax or home insurance, and continue and keep the loan in good standing.

We index households by i and let $t \in \{0, 1, \ldots, T\}$ denote the number of years since loan closing, with $t = 0$ denoting the take-up period. Each period households choose an action a_{it} from a finite set of alternatives $\mathcal{A}_t$. Households make these decisions taking into account their current state as characterized by a state vector s_{it}. We describe the specific state variables used in Section 4 below, when we discuss our data sources. In the remainder of this section we complete the description of the general structural model, including the payoff functions and value functions which are the main objects of interest in our empirical analysis.

Households in period $t = 0$ have completed the mandatory HECM counseling but have not yet closed on a HECM. Hence, cohorts in our data are defined by the year of counseling.

[4]Like all dynamic discrete choice models, in reality a household's HECM decisions are embedded in a larger utility maximization problem with a budget constraint that fully incorporates capital gains and losses. We do not observe household consumption or savings, and we only observe income in the take-up period, so we cannot estimate this larger model. Hence, the scope of this paper is limited to this "partial optimization" model over HECM decisions.

[5]To reduce potential losses to its insurance fund, HUD issued a moratorium on the fixed rate, full draw HECM on June 18, 2014 (Mortgagee Letter 2014-11).

[6]Our definition of a "large draw" as at least 80% of available credit was motivated by the cutoffs used in the HUD/FHA actuarial reports: 0-80% and 80-100% withdrawals for fixed rate HECMs and 0-40%, 40-80%, and 80-100% for adjustable rate HECMs (IFE, 2015, Exhibit IV-10).

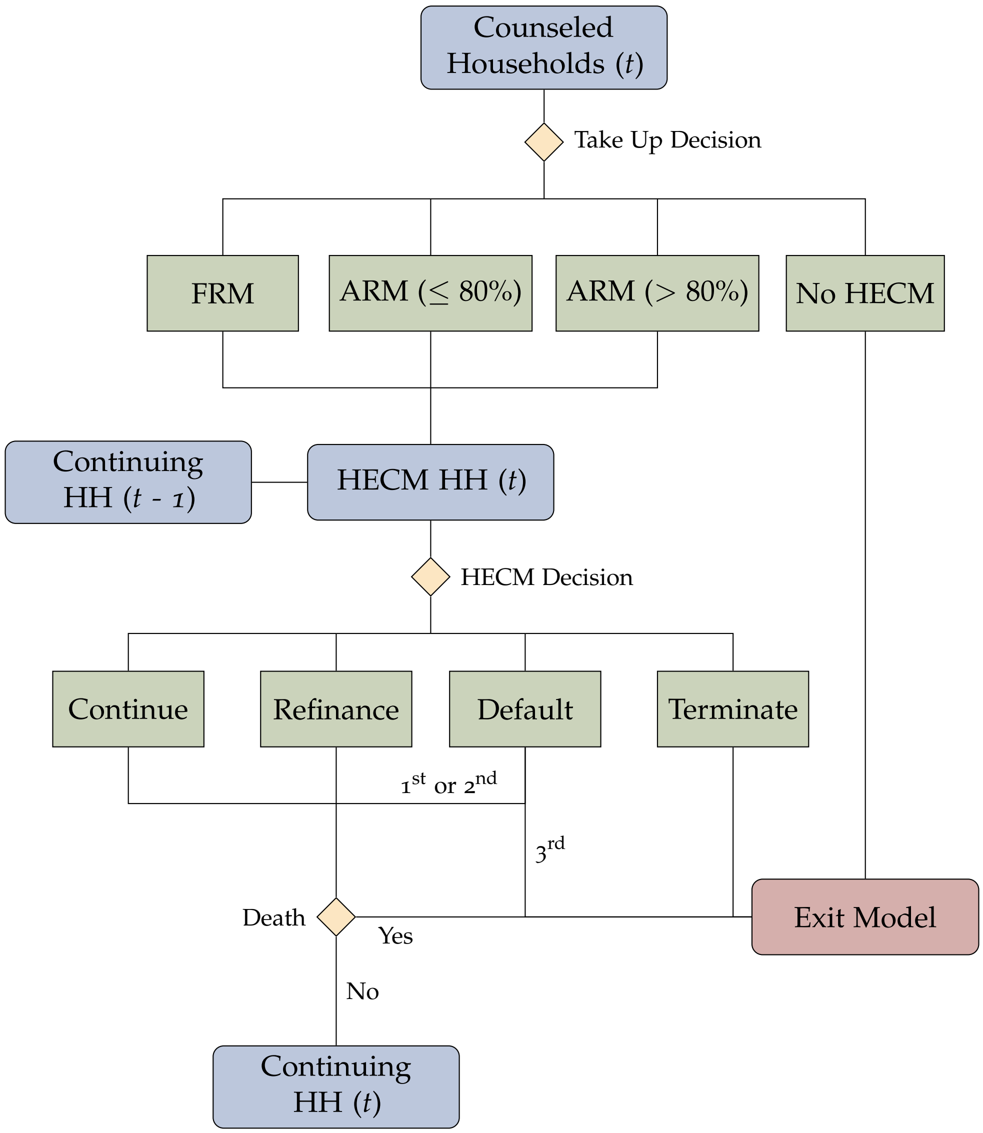

FIGURE 1. Borrower Decision Flow Chart

Households in period $t = 0$ make a take-up decision and, conditional on obtaining a HECM, in periods $t > 0$ they make decisions related to the HECM itself. Our focus is on HECM households ($t > 0$), but we note that accounting for the take-up decisions is important since some of our counterfactuals investigate scenarios where certain households are prohibited from taking-up a HECM. By backwards induction, the continuation values in the take-up problem depend on the decision process for HECM borrowers, so we first discuss the model for HECM households and return to the take-up model for counseled households below.

2.1. HECM Household Decisions

For a HECM household, there are four possible actions in $\mathcal{A}_t$ (corresponding to the decision node in Figure 1). The simplest decision a household can make is to simply continue living in the home and maintaining the reverse mortgage in good standing ($a_{it} = \text{C}$, "continue"). Second, a household could choose to refinance the HECM with another HECM ($a_{it} = \text{R}$, "refinance"). Such households obtain a new HECM with different terms and hence they remain in the pool of HECM households in subsequent periods. Next, households may choose to default ($a_{it} = \text{D}$, "default"). While forward mortgagors default by failing to make the scheduled payments, HECM borrowers are not required to make mortgage payments. Rather, default occurs when the homeowner fails to make scheduled property tax and insurance payments and there are no remaining funds on the HECM credit line (otherwise, the lender could use HECM funds to make the payments on behalf of the homeowner). Default can also occur if the homeowner fails to satisfy other loan obligations such as occupancy and property maintenance (Mortgagee Letter 2015-10). In practice, the HECM is not marked "due and payable" and the foreclosure process do not begin immediately when a household defaults. Some borrowers in our sample remain in default for up to four years without termination of the HECM[7].To account for this, we assume that the loan is not forced to terminate unless a household is in default for three consecutive periods. Finally, a household may terminate the loan ($a_{it} = \text{T}$, "terminate") for events other than defaulting on tax or insurance and refinance, which can happen if the mortgagor(s) sell the home in order to move, downsize or take advantage of house price changes since loan

[7]Our sample ends in 2014. Since 2015, HUD has provided new rules clarifying the process in case of a default. In 2015, HUD clarified that the mortgagee must submit a "due and payable" request once the HECM borrower is in default (Mortgagee Letter 2015-10), and that the mortgagee may delay the initiation of foreclosure proceedings if certain loss mitigation policies can be followed, such as repayment plans, with "goal of keeping HECM borrowers in their homes whenever possible" (Mortgagee Letter 2015-11, 2016-07). Additional delays to the foreclosure process may occur for HECM loans that experience tax or insurance default after being assigned to HUD (GAO 2019). HECM loans in good standing may be assigned to HUD when the balance on the loan reaches 98 percent of the maximum claim amount.

origination. Hence, the set of feasible actions for HECM households is

$$\mathcal{A}_t = \{\text{Continue}, \text{Refinance}, \text{Default}, \text{Terminate}\} = \{\text{C}, \text{R}, \text{D}, \text{T}\}.$$

Households that terminate or terminally default receive the payoff and exit the model immediately. For the remaining households, we account for the possibility that the HECM may terminate exogenously due to death of the borrowers by augmenting the discount factor with a survival probability.[8]

Clearly HECM borrowers make several decisions other than the four simple discrete decisions we focus on in the dynamic discrete choice model. These secondary choices and the factors that influence them are captured by the vector of state variables, introduced below. For example, after choosing to "continue" those households with credit line HECMs also decide how much additional money to withdraw. This choice is incorporated in the state transition process for available HECM credit, with the difference between consecutive periods being driven by the interest rate and the amount withdrawn.

2.2. *Utility Functions, Dynamic Decisions, and Value Functions*

The dynamic problem faced by HECM households can be thought of as optimal stopping problem since terminal default and termination are irreversible decisions. Hence, these terminating actions are equivalent to choosing a lump sum payoff equal to the present discounted value of the future utility received after leaving the model. Borrowers who continue to pay or refinance receive utility in the period which is a combination of utility from housing services and being able to draw on the line of credit and disutility from making property tax and insurance payments and from maintaining the home. Households who default once or at most twice consecutively also receive utility from housing services but not from the line of credit nor do they incur the disutility of making property tax and insurance payments (and potentially not from maintaining the home).

We will describe the state variables in detail below, but for now we simply assume that all payoff-relevant variables are captured by the observables s_{it} and unobservables ε_{it}. The observable states s_{it} include choice-specific state variables relevant for each individual discrete choice as well as auxiliary state variables that do not appear in the choice-specific payoffs directly, but influence the evolution of the payoff-relevant state variables. For example, two state variables included in the payoffs in our application are net equity and available HECM credit. These evolve in part based on the interest rate and the housing price index, which are therefore included as auxiliary components of s_{it} that do not appear

[8] We assume that each household's beliefs about continuing to the next period are consistent with mortality rates from the United States obtained from the 2011 CDC life tables. For loans with two borrowers, we use the joint probability that both borrowers die in the same year.

in the payoffs directly. State variables such as these are important for capturing the influence of secondary decisions made by HECM households, such as cash withdrawals from credit-line HECMs (which in part determine the evolution of available HECM credit).

In general, the period utility received by a household in state s_{it} that chooses action $a_{it} \in \mathcal{A}_t$ is $U_t(s_{it}, a_{it}, \varepsilon_{it})$, where $\varepsilon_{it}(a_{it})$ is an idiosyncratic, choice-specific error term. Households in our model are forward-looking and discount future utility using a discount factor b. We denote the the survival probability for household i, conditional on age and sex in period t, by $p(s_{it})$. As we show below, the discount factor is identified in our model and we estimate it along with the utility function. A decision rule for a household is a function $\delta_t : (s_{it}, \varepsilon_{it}) \mapsto a_{it}$ mapping states to actions in the choice set $\mathcal{A}_t$. Because we do not observe the idiosyncratic shocks ε_{it}, we will also work with the corresponding conditional choice probability (CCP) function or *policy function* $\sigma_t(s_{it}, a_{it})$.

Before describing the model more fully, we first briefly state three standard assumptions invoked by Rust (1987), Hotz and Miller (1993), and the literature that followed.

Assumption 1 (Basic Assumptions)**.** The primitives of the dynamic discrete choice model have the following properties:

a. The state variables and errors follow a controlled, time-homogeneous, first-order Markov process where the joint transition density can be factored as follows:

$$f(s_{i,t+1}, \varepsilon_{i,t+1} \mid s_{it}, \varepsilon_{it}, a_{it}) = f(s_{i,t+1} \mid s_{it}, a_{it}) f(\varepsilon_{i,t+1}).$$

b. The payoffs are additively separable in the choice-specific errors and the deterministic component is a time invariant function of s_{it} and a_{it}:

$$U_t(s_{it}, a_{it}, \varepsilon_{it}) = u(s_{it}, a_{it}) + \varepsilon_{it}(a_{it}).$$

c. The choice-specific errors follow a known joint CDF $F_\varepsilon(\cdot)$ which is absolutely continuous with respect to Lebesgue measure with strictly positive density on $\mathbb{R}^{|\mathcal{A}_t|}$ and finite first moments.

The first and second parts are conditional independence and additive separability assumptions along the lines of Assumptions AS and CI of Rust (1994). We assume s_{it} follows a first-order Markov process that is conditionally independent from ε_{it} but may depend on a_{it} and that households have rational expectations, hence they know the law of motion of s_{it} and can evaluate the conditional expectation of $s_{i,t+1}$ given s_{it} and a_{it}. The third part requires that the distribution of errors is known and has full support, which allows us to invoke the CCP inversion of Hotz and Miller (1993, Proposition 1).

In our application and in some examples below, we will work under the assumption of type I extreme value errors for analytical convenience. However, in light of results by Norets and Takahashi (2013) and BCNP on the surjectivity of the mapping from CCPs to differences in choice-specific value functions, it is not necessary to assume a specific parametric distribution for our main identification results. It is known in the literature that Assumption 1 alone does not provide sufficient restrictions to identify structure parameters of the model (Rust (1994)). Magnac and Thesmar (2002) show that to identify utility functions in each alternative, in addition to those in Assumption 1, restrictions on the discount rate and the utility of a reference choice are also needed.

Following the literature, we define the *ex ante value function* $V_t^\delta(s_{it})$ as the expected present discounted value received by a household i that behaves according to the sequence of decision rules $\delta = (\delta_0, \delta_1, \ldots, \delta_T)$ in the current period and in the future. Let I_{it} be an indicator variable equal to 0 if household i did not take up a HECM in period $t = 0$ or took up a HECM that is no longer active due to termination, default, or death and equal to 1 otherwise. Then,

$$(1) \qquad V_t^\delta(s_{it}) = \mathrm{E}^\delta \left[\sum_{\tau=t}^{T} b^{\tau-t} U_\tau(s_{i\tau}, \delta_\tau(s_{i\tau}, \varepsilon_{i\tau}), \varepsilon_{i\tau}) I_{i\tau} \,\middle|\, s_{it} \right].$$

Here, E^δ denotes the conditional expectation over future states given the current state and that the household behaves according to the sequence of decision rules δ. The indicator I_{it} ensures that households receive no additional utility after termination, terminal default, death, or initially choosing not to take up a HECM. Since our model is a finite-horizon model, the optimal decision rules can be determined via backwards induction. We assume that households use this sequence of optimal decision rules and therefore we drop the explicit dependence on δ in the remainder.

Importantly, our model has two distinct termination outcomes. As we show below, this property allows us to identify the utility function without a normalization and therefore to make unbiased welfare calculations and counterfactual predictions. For non-terminating actions a_{it}, households receive the mean utility $u(s_{it}, a_{it})$ plus the idiosyncratic shock. Additionally, because they are forward-looking they also expect to receive additional utility in the future. Households discount that utility appropriately and account for uncertainty over future states. This includes periods in which a household chooses to default the first or second time in a row ($a_{it} = \mathrm{D}$). On the other hand, when households terminate by choosing $a_{it} = \mathrm{T}$, they receive the mean period utility $u(s_{it}, \mathrm{T})$ and the idiosyncratic shock $\varepsilon_{it}(\mathrm{T})$, but no additional utility is received in the future. Hence, $u(s_{it}, \mathrm{T})$ can be thought of as a termination payoff that includes any additional discounted expected utility received in the future after leaving the HECM program. Finally, when a household terminates by defaulting for a third time in a row ($a_{it} = a_{i,t-1} = a_{i,t-2} = \text{D}$), they receive the mean utility for defaulting $u(s_{it}, \text{D})$, the idiosyncratic shock, and because the HECM will be terminated, the termination payoff $u(s_{it}, \text{T})$.

In order to calculate conditional choice probabilities, we first introduce the *choice-specific value function* $v_t(s_{it}, a_{it,-1}, a_{it})$ for HECM households in periods $t > 0$. Letting $\beta_{it} = b \times p(s_{it})$ denote the product of the discount factor and survival probability, we have

$$
(2) \quad v_t(s_{it}, a_{it,-1}, a_{it}) = \begin{cases} u(s_{it}, \text{C}) + \beta_{it}\,\text{E}[V_{t+1}(s_{i,t+1}) \mid s_{it}, a_{it} = \text{C}] & a_{it} = \text{C}, \\ u(s_{it}, \text{R}) + \beta_{it}\,\text{E}[V_{t+1}(s_{i,t+1}) \mid s_{it}, a_{it} = \text{R}] & a_{it} = \text{R}, \\ u(s_{it}, \text{D}) + \beta_{it}\,\text{E}[V_{t+1}(s_{i,t+1}) \mid s_{it}, a_{it} = \text{D}] & a_{it} = \text{D}, a_{it} \neq a_{i,t-1} \text{ or } a_{it} \neq a_{i,t-2} \\ u(s_{it}, \text{D}) + u(s_{it}, \text{T}) & a_{i,t-2} = a_{i,t-1} = a_{it} = \text{D}, \\ u(s_{it}, \text{T}) & a_{it} = \text{T}. \end{cases}
$$

The first three cases are standard in dynamic discrete choice models. Households receive period utility and continue to the next period. Importantly, this is also true for the first or second year of default. For a forward mortgage, default is usually considered to be a terminal action (e.g., BCNP), however, in our sample of HECM households, missed property tax or insurance payments (T&I default) were not followed quickly by foreclosure proceedings. In addition, a household could pay off the past due property tax or insurance balance. Therefore, in our model, we allow a household to continue with the HECM after their first or second year of default. The dependence of the choice-specific value function on the past sequence of actions is captured by $a_{it,-1}$, which can vary in different applications. In our setting, $a_{it,-1} = \{a_{i,t-2}, a_{i,t-1}\}$.

The last two cases correspond to the terminating actions: defaulting for three years or direct termination. In our sample, 99.16% of households who default three years in a row continue to default or terminate in the following year. Therefore, it seems reasonable to expect that households who have stayed in default for three years will no longer actively manage their HECM loans. Hence, such households no longer make decisions in our model and instead receive a lump-sum terminal payoff. Similarly, no future utilities are received from the HECM program when the direct termination action is taken.

Although we do not rely on a specific parametric distribution for identification, when estimating the model we assume that the idiosyncratic shocks follow the type I extreme value distribution. The mapping from differences in choice-specific value functions to CCPs is invertible for a very broad class of continuous distributions (Hotz and Miller, 1993; Norets and Takahashi, 2013; Bajari et al., 2016), but it happens that the type I extreme value distribution is also analytically tractable. In this special case the conditional choice probabilities have a convenient closed form in terms of the choice-specific value function:

$$\sigma_t(s_{it}, a_{it,-1}, a_{it}) = \frac{\exp\left(v_t(s_{it}, a_{it,-1}, a_{it})\right)}{\sum_{j\in\mathcal{A}_t}\exp\left(v_t(s_{it}, a_{it,-1}, j)\right)}. \tag{3}$$

We formalize our modeling assumptions below (see Assumption 1), such as additive separability of payoffs and conditional independence of the idiosyncratic errors, which are quite standard in the literature on structural dynamic discrete choice models (Rust, 1994; Aguirregabiria and Mira, 2010).

2.3. HECM Take-Up Decisions

For a counseled household, there are four possible actions in $\mathcal{A}_0$ (corresponding to the take-up decision node in Figure 1). Households can take up an adjustable-rate HECM with either a small ($a_{i0} = \text{A}$) or large ($a_{i0} = \text{AL}$) initial withdrawal, a fixed rate HECM ($a_{i0} = \text{F}$), or they can choose not to take up a HECM at all ($a_{i0} = \text{N}$). For fixed-rate HECMs, households necessarily make a full draw so we do not distinguish between small and large initial withdrawals. Households that choose not to take up a HECM ($a_{i0} = \text{N}$) exit the model.[9] The type of HECM and, in the case of an adjustable-rate HECM, whether the initial withdrawal was large or not, become state variables and therefore affect the household's later decisions. Hence, the set of feasible actions for HECM households is

$$\begin{aligned}\mathcal{A}_0 &= \{\text{Adjustable Rate}, \text{Adjustable Rate (Large Draw)}, \text{Fixed Rate}, \text{No HECM}\}\\ &= \{\text{A}, \text{AL}, \text{F}, \text{N}\}.\end{aligned}$$

As with the HECM model, the utility of the choices associated with HECM take-up are functions of the state variables and are additively separable in the error term as

$$U_0(s_{i0}, a_{i0}, \varepsilon_{i0}) = u_0(s_{i0}, a_{i0}) + \varepsilon_{i0}(a_{i0}). \tag{4}$$

In this case the payoffs $u_0(s_{i0}, a_{i0})$ represent sums of current payoffs and discounted continuation values determined by the HECM household model. As in the model above for HECM households, we will assume that $\varepsilon_{i0}(a_{i0})$ follows type I extreme value distribution

[9] Although HECM counseling is valid for two years, 99% of households in our sample who took up HECMs after counseling did so in the same year. Hence, to construct a parsimonious model of HECM take-up we assume that households either take up in the same year or not at all.

which leads to choice probabilities of the following form:

$$\sigma_0(s_{i0}, a_{i0}) = \begin{cases} \frac{\exp(u_0(s_{i0}, a_{i0}))}{1+\sum_{j\in\{\text{A,AL,F}\}} \exp(u_0(s_{i0}, j))} & a_{i0} \in \{\text{A}, \text{AL}, \text{F}\}, \\ \frac{1}{1+\sum_{j\in\{\text{A,AL,F}\}} \exp(u_0(s_{i0}, j))} & a_{i0} = \text{N}. \end{cases}$$

Recall that Assumption 1 includes a conditional independence assumption that limits the persistence of the unobservables (Rust, 1987). In practice, we further assume that the error terms are independent across time and individuals and in particular, the errors ε_{i0} in the take-up choices (4) are independent from the error terms ε_{it} corresponding to subsequent HECM choices. Although this greatly simplifies the dynamic problem faced by HECM households, which can be separately studied from the HECM take-up choices, it is nonetheless a limitation of our analysis.[10]

In the next section, we show that in models such as ours, with multiple terminating actions and a finite horizon, both the utility function and the discount factor are semiparametrically identified without a utility normalization. Furthermore, we show that although welfare (actual and counterfactual) and counterfactual CCPs are not identified in general, all of these quantities are identified in our model. We then propose an estimator for the model, which is a multi-step plug-in semiparametric procedure based on BCNP.

3. Semiparametric Identification and Estimation

In this section we briefly describe our semiparametric identification results followed by the semiparametric estimator we use. Following most of the literature on identification of dynamic discrete choice models under Assumption 1, we first note that by Hotz and Miller (1993) one can identify differences in conditional value functions. To identify the discount factor and then the utility function itself we exploit the structure of our model, with multiple terminating actions as in (2). An example of this structure from labor economics would be an employee who can either quit immediately (immediate termination) or be fired by failing to meet performance criteria for multiple periods in a row (repeated termination). In addition, we require two more assumptions: completeness of a particular conditional transition distribution and nonstationarity of the CCPs (Assumptions 2 and 3 in the Appendix). See Appendix A for a complete derivation of our identification results and Appendix B for the proofs.

[10] In a reduced-form model, Moulton et al. (2015) allowed the unobservables that determine HECM take-up and default to be correlated. They estimated the correlation between the HECM take-up and default to be −0.38, so there is moderate evidence of negative correlation. Unfortunately, it would be computationally prohibitive at present to allow serial correlation in unobservables in our dynamic structural, which has many continuous state variables (Blevins, 2016).

It is known in the literature that if there is a sequence of choices which makes the future value not dependent on the initial choice, the expected payoff can be represented as a sum of per-period utilities and CCPs for certain time periods, and CCP-based estimators can be used (Arcidiacono and Miller, 2011). Our model builds on this framework. To identify the utility functions without normalization, we observe that with two sequences of choices that can terminate the decision problem and a completeness assumption, each component in the sum of per-period utilities can be identified, rather than the difference relative to a choice.

The identification of the discount factor relies on nonstationary conditional choice probabilities, which has been recognized in the literature (Bajari et al., 2016). In finite-horizon models such as ours, the CCPs are nonstationary, which means that if we fix the current states, the CCPs across two periods can still be different, and the variation comes only from the continuation value which is a product of the discount factor and a future value term that can be shown to be identified from the data. The discount factor, being the only term unknown, is then identified.

Estimation proceeds in multiple steps using a plug-in semiparametric approach. The procedure is based on BCNP, but with some modifications since we do not assume one of the choice-specific payoff functions is known nor do we need to observe the final decision period. In the first step, as in BCNP, we nonparametrically estimate the conditional choice probabilities. Specifically, returning to our empirical model with choice set $\mathcal{A} = \{\mathrm{C}, \mathrm{R}, \mathrm{D}, \mathrm{T}\}$, we use a series representation of the log odds ratio

$$\log \frac{\sigma_t(s, a_{-1}, a)}{\sigma_t(s, a_{-1}, \mathrm{T})} = \sum_{l=1}^{\infty} r_l(t, a) q_l(s, a_{-1}) \approx \sum_{l=1}^{L} r_l(t, a) q_l(s, a_{-1})$$

for choices $a \in \mathcal{A}$ relative to termination ($a = \mathrm{T}$). The functions $\{q_l(\cdot)\}_{l=1}^{\infty}$ are basis functions and $\{r_l(\cdot)\}_{l=1}^{\infty}$ are the coefficients which will be estimated. The conditional choice probabilities depend on the time period in order to capture nonstationary choice behaviors in finite horizon models.

In practice we use a Hermite polynomial sieve as the basis functions because it can approximate multivariate smooth functions of variables with unbounded supports well (Chen, 2007). $q_l(s, a_{-1})$ is a vector from the tensor products of Hermite polynomials of the state variables. We approximate the infinite sum using a finite number of Hermite polynomials of the continuous state variables and interactions among them and with the discrete state variables, and L denotes the highest degree of the polynomials which is selected by 10-fold cross-validation. The sample is randomly partitioned into 10 groups. For each of the 10 groups (testing dataset), we first estimate the parameters in a candidate model leaving this group out and then obtain the log likelihood of the model using the testing dataset. The pool of candidate models includes Hermite polynomials up to degree 3. The model with the highest average log likelihood on the testing datasets is selected. We also select the model using Hermite polynomials up to degree 4, and the results are similar.

Let $\hat{\sigma}_t(s, a_{-1}, a)$ denote the estimated choice probabilities, obtained as

$$(5) \quad \hat{\sigma}_t(s, a_{-1}, a) = \frac{\exp\left(\sum_{l=1}^{L} \hat{r}_l(t, a) q_l(s, a_{-1})\right)}{1 + \sum_{j \in \mathrm{A} \setminus \{\mathrm{T}\}} \exp\left(\sum_{l=1}^{L} \hat{r}_l(t, j) q_l(s, a_{-1})\right)}$$

for $a \in \mathcal{A} \setminus \{\mathrm{T}\}$ and

$$\hat{\sigma}_t(s, a_{-1}, \mathrm{T}) = 1 - \hat{\sigma}_t(s, a_{-1}, \mathrm{C}) - \hat{\sigma}_t(s, a_{-1}, \mathrm{R}) - \hat{\sigma}_t(s, a_{-1}, \mathrm{D}).$$

We nonparametrically estimate the take-up probabilities for $t = 0$ in a similar fashion.

As in BCNP, next we must nonparametrically estimate the period-ahead expected ex-ante value function. We assume that the choice-specific errors follow type I extreme value distributions. It follows that the period-ahead expected ex-ante value function is identified directly from the data through the relationship:

$$(6) \quad \mathrm{E}[V_{t+1}(s', a) \mid s, a] = -\,\mathrm{E}[\log \sigma_{t+1}(s', a, \mathrm{T}) \mid s, a] + \mathrm{E}[u(s', \mathrm{T}) \mid s, a] + \gamma.$$

The first conditional expectation on the right hand side can be estimated using a nonparametric regression of $\log \hat{\sigma}_{t+1}(s', a, \mathrm{T})$ on s and a as in BCNP, where the choice probabilities $\hat{\sigma}_{t+1}(s', a, \mathrm{T})$ are estimated in the previous step (5). The advantage of estimating this conditional expectation term using nonparametric regression compared with alternative approaches such as forward simulation is its computational simplicity. We implement this by method of sieves using Hermite polynomials with the number of basis functions selected by cross validation, $\hat{\mathrm{E}}[\log \hat{\sigma}_{t+1}(s', a, \mathrm{T}) \mid s, a] = \sum_{l=1}^{\tilde{L}} \tilde{r}_l(t) q_l(s, a)$. Meanwhile, γ is a known constant. Because we do not assume that the termination utility function is the zero function there is an additional term on the right hand side of (6) relative to BCNP. In their case, the second term on the right hand side of (6) vanishes. This additional term is also a function of s and a and can be estimated given the parametric form for the utility function and the estimated law of motion of the state variables, $\mathrm{E}[u(s', \mathrm{T}) \mid s, a] = \int u(s', \mathrm{T}) f_{s'|s,a}(ds')$.

Although, our procedure involves this additional step of estimating the state transition distribution, it is not new and is part of the first step in other multi-step estimators such as Aguirregabiria and Mira (2002, 2007), Bajari, Benkard, and Levin (2007), and Pesendorfer and Schmidt-Dengler (2007). One could avoid this step by assuming that the termination payoff is the zero function, however, if that assumption was incorrect the estimates would be biased. In our empirical setting, we hypothesized that the payoff to termination would be different based on demographics and household finances and our estimates indeed support that view.

Finally, we estimate the structural parameters via nonlinear least squares. This includes the utility parameters θ and β which is a product of the discount factor and the survival probability. The estimating equations are the log odds ratios for the choices $a \in \mathcal{A}$:

$$\begin{aligned}\log \frac{\sigma_t(s, a_{-1}, a)}{\sigma_t(s, a_{-1}, \mathrm{T})} =& u(s, a; \theta) - u(s, \mathrm{T}; \theta) + \beta \, \mathrm{E}\left[V_{t+1}(s', a) \mid s, a\right] \\ =& u(s, a; \theta) - u(s, \mathrm{T}; \theta) \\ & - \beta \, \mathrm{E}[\log \sigma_{t+1}(s', a, \mathrm{T}) \mid s, a] + \beta \, \mathrm{E}[u(s', \mathrm{T}; \theta) \mid s, a] + \beta\gamma.\end{aligned}$$

for $a \in \{\mathrm{C}, \mathrm{R}, \mathrm{D}\}$. Substituting in estimated quantities from the first step yields

$$\begin{aligned}\log \frac{\hat{\sigma}_t(s, a_{-1}, a)}{\hat{\sigma}_t(s, a_{-1}, \mathrm{T})} =& u(s, a; \theta) - u(s, \mathrm{T}; \theta) \\ & - \beta \sum_{l=1}^{\tilde{L}} \tilde{r}_l(t) q_l(s, a) + \beta \int u(s', \mathrm{T}; \theta) f_{s'|s,a}(ds') + \beta\gamma.\end{aligned}$$

This allows us to estimate the structural parameters θ and β by nonlinear least squares.

This procedure defines a semiparametric plug-in estimator of the kind considered by Ai and Chen (2003). The first step is a series estimator for the conditional choice probabilities for which consistency and a $n^{1/4}$ rate of convergence follow from Wong and Shen (1995), Andrews (1991), and Newey (1997). BCNP provide regularity conditions to establish these properties for the first step estimator, which is the same estimator we use, as well as a proof of asymptotic normality of a closely-related second step estimator. Asymptotic normality of our second-step estimator follows as a straightforward modification of their conditions.

Furthermore, in our application we assume that the period utility for each choice a is linear in the state variables s with coefficients θ^a: $u(s, a; \theta) = s'\theta^a$. Identification of θ in this case was established in Lemma 2. We also parameterize the state transition rule by assuming that stochastic time varying components of s follow a first order VAR process. The state transition rules are estimated by fitting a system of seemingly unrelated regressions using data on s', s and a, thus allowing for contemporaneous correlations between errors associated with the time varying stochastic state variables. The estimating equations are now simplified to:

$$\log \frac{\hat{\sigma}_t(s, a_{-1}, a)}{\hat{\sigma}_t(s, a_{-1}, \mathrm{T})} = s'\theta_a - s'\theta_\mathrm{T} - \beta \sum_{l=1}^{\tilde{L}} \tilde{r}_l(t) q_l(s, a) + \beta \hat{\mathrm{E}}\left[s' \mid s, a\right]' \theta_\mathrm{T} + \beta\gamma,$$

for $a \in \{C, R, D\}$.

4. Data

Our data is drawn from a sample of 20,239 senior households counseled for a reverse mortgage during the years 2007 to 2011, from a single HUD counseling agency. These data include demographic and socio-economic characteristics of the counseled household, as well as credit data at the time of counseling and annually thereafter for at least three years post counseling. Our entire sample spans the years 2007–2014. The credit attributes data includes credit score, outstanding balances and payment histories on revolving and installment debts, and public records information. For those originating a HECM (58.42 percent of counseled households in our sample), counseling data is linked to HUD HECM loan data using confidential personal identifiers. HUD HECM loan data includes information on origination, withdrawals, terminations and tax and insurance defaults.

Our rich dataset allows us to include many state variables in the dynamic discrete choice model that help capture household demographics and financial well-being as well as the economic conditions they face. Household characteristics and the economic climate in turn inform the decisions households make. Although some state variables are fixed over time, many key variables such as the credit score are time-varying.

To control for differences in household demographics, we include age and age squared as state variables along with indicator variables for young borrowers (less than 65 years old), Hispanic and black borrowers, as well as single male and single female borrowers. Additionally, we include many measures of household financial health as state variables. We observe borrowers' credit data annually which allows us to follow the evolution of the credit score, total available revolving credit, and the balances of any revolving and installment credit lines. Each year we also observe several variables related to the borrowers' HECMs including the HECM balance (principal plus accumulated interest) and the balance on defaulted tax and insurance (T&I) payments. Additionally, we observe the value of the property at closing, its zipcode, and the evolution of the housing price index,[11] allowing us to forecast the value of the home over time. From this we calculate borrowers' net equity and two variables we will refer to as "HECM credit" and "excess credit". These variables are further defined below. The remaining financial variables are observed at the time of HECM counseling and are time-invariant. These include monthly income, non-housing assets, and the property tax to income ratio. We also include indicator variables for households with fixed-rate HECMs and households who took large initial

[11] We use the Federal Housing Finance Agency MSA level all-transactions house price index. For households located outside a MSA, we use the state housing price index. These indices are deflated by the consumer price index (CPI).

withdrawals (80% or more).

In the following are several of the financial variables that deserve special attention. These variables are similar in what they measure, but they move over time in distinct ways that allow us to study whether and how households value various features of an HECM loan.

HECM Balance The current HECM balance is calculated based on the amounts a borrower withdraws over time. This balance grows at a rate equal to the interest rate plus a monthly mortgage insurance premium. For FRM borrowers, the entire line of credit is drawn at closing, and so no additional withdrawals can be made. ARM borrowers can choose their initial withdrawal amount and may make subsequent withdraws, as needed or on an installment basis[12].

Net Equity Net equity is defined to be the current value of the home less the current HECM balance. For example, the net equity for a household with a home valued at $200,000 and with a HECM balance of $70,000 would be $130,000. A *ceteris paribus* increase in net equity represents the effect of home equity increasing, controlling for the amount of HECM credit that can still be accessed and the insurance value of the HECM (excess credit). To allow for asymmetric effects of positive and negative net equity, we also include the absolute value of negative net equity as a state variable. This variable is positive only when a household has negative net equity; it is defined to be zero when a household has positive equity.

HECM Credit The current available HECM credit is the amount of money that a borrower can withdraw from HECM line of credit after adjusting for past withdrawals and credit line growth. This variable is zero for FRM borrowers after the first year because FRM HECMs are structured as closed-end mortgages and borrowers are not permitted to make any additional withdrawals after closing. For ARM borrowers, like the HECM balance, this amount also grows at a rate equal to the interest rate plus the mortgage insurance premium. A *ceteris paribus* increase in HECM credit represents the immediate liquidity that can be extracted from the HECM, which is independent of the home value.

Excess Credit We define excess credit to be the difference between the available HECM credit and the current home value when this quantity is positive, or $0 otherwise. In other words, we say a household has excess credit when the available HECM credit exceeds the value of the home. For example, for a household with $170,000 in available HECM

[12] ARM borrowers have more disbursement options, including a lump-sum payment, a tenure payment (a fixed monthly payment for as long as the borrower lives in the home and satisfies other loan obligations), a term payment (a fixed monthly payment for a fixed number of periods) or a line of credit.

credit and a home valued at \$160,000 the excess credit would be \$10,000. If the home were instead valued at \$180,000, excess credit would be \$0 since the home value exceeds the available credit. For most households in our sample, excess credit is \$0. Due to the non-recourse nature of the loan, when excess credit is positive it represents the amount of money the household could save by drawing all funds before terminating the HECM.

To illustrate these three variables, we consider two example households with homes originally valued at \$200,000 and with identical HECMs. Both households had initial principal limits of \$120,000 and initial withdrawals equal to \$70,000. Suppose the first household's home value has held steady at \$200,000 but the second household's home has significantly fallen in value to \$110,000. For simplicity, suppose that the decline happens immediately after closing so that we can abstract away from growth in the HECM balance and HECM credit. For comparison, the values of the net equity, HECM credit, and excess credit variables for these two households are shown in Table 1.

Clearly, net equity is higher for the first household. Since the HECMs and withdrawals are identical, the available HECM credit is the same for both households. However, excess credit is only non-zero for the second household, which has borrowing power (HECM credit) in excess of net equity.

TABLE 1. Example Households: Net Equity, HECM Credit, and Excess Credit

Variable	Household 1	Household 2
Original Home Value	\$200,000	\$200,000
Current Home Value	\$200,000	\$110,000
HECM Credit Limit	\$120,000	\$120,000
HECM Balance	\$70,000	\$70,000
Net Equity	\$130,000	\$40,000
HECM Credit	\$50,000	\$50,000
Excess Credit	\$0	\$10,000

An existing HECM can be refinanced into another HECM. The borrower can increase the amount of borrowing from the new HECM if the interest rate, property value, or the principal limit factor as set forth by HUD moves in its favor[13]. To capture these factors, we have included net equity in the state variables, which should capture the effect of home price movements ("cash-out refi" motives). To further capture "rate refi" motives, for each HECM household and each year in our sample, we compute the *refinance benefit*, as $\Delta\text{IPL}_{\text{refi},it} = \text{IPL}_{\text{refi},it} - \text{HECM Credit Limit}_{it}$, where $\text{IPL}_{\text{refi},it}$ is the estimated

[13]A borrower may choose to refinance due to other reasons such as adding a second borrower to the loan. HUD defines HECM refinance as "a HECM refinance case is the refinance of an existing HECM with a new HECM for the same borrower and same property with different loan specifications." These household specific factors are modeled by choice specific random shocks in our model.

initial principal limit should household i refinances in year t, and its formula is

$$\text{IPL}_{\text{refi},it} = \text{principal limit factor}_{it} \times \min\{\text{est. home value}_{it}, \text{FHA loan limit}_{it}\}.$$

HUD provides a principal limit factor table which shows the value of the factor for each combination of borrower age and interest rate. During our sample period, the table was revised in 2009, 2010, 2013 and 2014. The FHA loan limit has also been revised from time to time and this has been taken into account. For fixed rate borrowers, we also compute the *interest rate differential* which is the difference between the loan rate and current best market rate[14].

Table 2 reports the summary statistics for our HECM sample. The reported means and standard deviations are at the household-year level, meaning that there are multiple observations for each household for each year until the HECM terminates. The first four columns report the mean for each variable conditional on the current household action a_{it}. The last column reports the overall mean and standard deviation for each variable. Recall that households are counted in these statistics for multiple years until termination, which explains why the default action (which can be repeated) is observed much more often than termination (which is immediate).

Comparing across actions, we see relatively few refinance and termination actions relative to default, in part because those households leave the sample while households who default can remain in the sample for multiple years (and they tend to remain in default). Around 40% of our observations are for single female households, 14% are black, and 8% are Hispanic. Average monthly income at time of origination is $2,380. Approximately 56% of observations are for FRMs and 70% of observations correspond to borrowers who took large initial withdrawals. The overall mean age of HECM borrowers across observations in our sample is 73 years. Borrowers who refinance tend to be slightly younger, on average around 72 years old, while the mean age at termination is 76.

For household-year observations where we observe a default, households are more likely to have taken large initial withdrawals and have fixed rate HECMs. They also have lower incomes, lower credit scores, little available credit (HECM and other credit), lower net equity, higher excess credit, and have T&I default balances. The average credit score is 701, however, for borrowers who default it is 590. For refinance observations and to a less extent termination observations, households tend to have higher net equity, more available revolving credit, higher income, and higher property tax/income ratios.

Similarly, Table 3 reports the summary statistics for variables observed in the year of counseling for our take-up sample. Over half of the counseled borrowers do ultimately

[14] We compute average interest rates by state and year from the full HECM loan dataset for adjustable and fixed rate HECMs, and the interest rate used is the lower of the two.

TABLE 2. Summary Statistics for the HECM Sample

	Terminate	Refinance	Default	Continue	All Loans	
	Mean	Mean	Mean	Mean	Mean	SD
Time-Invariant Variables						
Young borrower	0.113	0.168	0.216	0.148	0.151	0.358
Hispanic	0.082	0.074	0.123	0.081	0.084	0.278
Black	0.059	0.174	0.268	0.131	0.139	0.346
Single male	0.200	0.215	0.187	0.150	0.156	0.363
Single female	0.406	0.383	0.479	0.393	0.400	0.490
Monthly income [†]	0.262	0.246	0.194	0.241	0.238	0.164
Property tax/income	0.106	0.114	0.104	0.091	0.093	0.095
Non-housing assets[†]	6.264	2.152	2.734	4.545	4.464	17.618
Fixed rate HECM	0.524	0.523	0.674	0.557	0.564	0.496
Initial withdrawal > 80%	0.604	0.718	0.890	0.682	0.695	0.461
# of borrowers	609[‡]	149[⋆]	996[*]	11,167[◇]	12,906	
Time-Varying Variables						
Age	75.693	72.188	72.856	72.979	73.009	7.561
Credit score	718.187	701.168	590.326	705.896	701.317	93.633
Available revolving credit[†]	2.328	3.131	0.313	2.250	2.175	3.040
Revolving & installment debt[†]	1.108	1.200	0.931	1.265	1.249	2.242
Net equity[†]	13.521	16.345	4.197	10.750	10.540	13.163
Negative net equity[†]	0.010	0.000	0.213	0.056	0.061	0.610
Excess credit[†]	0.032	0.000	0.182	0.072	0.076	0.506
Tax & insurance balance[†]	0.003	0.003	0.162	0.000	0.007	0.090
Available HECM credit[†]	3.003	3.156	0.130	3.164	3.037	6.109
Refinance benefit[†]	0.811	1.394	0.053	0.426	0.419	1.875
Interest rate differential	0.112	0.091	0.135	0.102	0.103	0.192
# of borrower-year observations	609	149	1,778	40,862	43,398	

[†] Monetary variables are measured in units of $10,000. [‡] Borrowers who terminated their HECMs in the sample. [⋆] Borrowers who refinanced at least once in the sample. [*] Borrowers who defaulted at least once in the sample. [◇] Borrowers who did not terminate, refinance, nor default in the sample.

Table 3. Summary Statistics for the Take-Up Sample

	FRM	ARM (≤ 80%)	ARM (> 80%)	No HECM	All HH. in Sample		Senior HH.*	
	Mean	Mean	Mean	Mean	Mean	SD	Mean	SD
Pre-HECM Variables								
Age	70.962	74.142	72.243	70.758	71.511	7.975	70.132	9.223
Young borrower	0.188	0.114	0.131	0.177	0.167	0.373	0.170	0.376
Hispanic	0.065	0.069	0.180	0.100	0.088	0.284	0.043	0.203
Black	0.148	0.079	0.204	0.229	0.174	0.379	0.066	0.249
Single male	0.160	0.150	0.168	0.188	0.170	0.376	0.114	0.318
Single female	0.385	0.450	0.400	0.374	0.393	0.488	0.310	0.462
Monthly income †	0.247	0.221	0.213	0.232	0.234	0.168	0.407	0.316
Property tax/income	0.078	0.118	0.106	0.085	0.090	0.094		
Non-housing assets†	4.507	4.392	2.158	4.492	4.313	17.252	28.373	43,522
Credit score	684.735	726.315	671.323	659.184	680.240	101.496		
Available revolving credit†	2.104	3.216	2.597	1.804	2.204	3.587		
Revolving & installment debt†	1.721	1.288	1.651	1.595	1.590	2.939		
Net equity†	14.952	23.397	15.882	16.557	17.125	21.496	15.649	17.978
Negative net equity†	0.047	0.021	0.014	0.304	0.147	2.122	0.252	4.435
1 year change in house price index	-0.055	-0.064	-0.083	-0.064	-0.062	0.054		
Average interest rate (ARM)	5.256	5.315	5.383	5.262	5.278	0.186		
Average interest rate (FRM)‡	5.280	5.303	5.355	5.277	5.284	0.209		
Initial HECM Variables								
Initial withdrawal > 80%	0.987	0.000	1.000	–	–	–		
Initial principal limit†	12.670	15.772	14.253	–	–	–		
# of households	6,927	3,450	1,446	8,416	20,239			

† Monetary variables are measured in units of $10,000. ‡ Fixed rate HECMs became widely available after April 2009. Therefore average interest rates of fixed rate HECMs are based on observations after April 2009. * Homeowning households with a member age 62 or greater from the 2010 wave of the Health and Retirement Study.

take up a HECM. Those that do take up a HECM tend to be older and in our sample, more households choose FRMs than ARMs. Households that choose small-draw ARMs have the highest average credit scores and those that choose large-draw ARMs have the lowest credit scores. Households with fewer non-housing assets tend to choose large-draw ARMs in particular. Lower income households tend to choose ARMs somewhat more often than FRMs.

To gauge to what extent our counseled sample differs from the general population of seniors, we compare our sample with observations from the 2010 wave of the Health and Retirement Study (HRS) who are homeowners with a household member at least 62 years of age. The last two columns of Table 3 report summary statistics of variables that are also available in the HRS data. Compared with the general senior population, households in our counseled sample are slightly older, more likely to come from minority racial groups, more likely to be single, have lower levels of income and non-housing assets, but a higher level of net equity. These observations are in line with the analysis of the characteristics of reverse mortgage borrowers by Davidoff (2014) and Haurin, Moulton, and Shi (2018). As such, our estimates below should be interpreted in terms of the preferences of HECM borrowers only. The socio-economic conditions that drive HECM counseling and take-up decisions may also affect subsequent HECM choices, so our estimates may not be representative of preferences in the overall population of seniors.

5. Estimation Results and Counterfactual Analysis

5.1. Reduced Form Policy Function Estimates

The conditional choice probabilities are estimated by a sieve multinomial logit model using the HECM borrower sample and all of the state variables from the structural model, which include variables in the per-period payoffs and time lags of macroeconomic variables that affect borrowers' expectations on future state variables such as local house price indices and average interest rates of adjustable and fixed-rate mortgages. For finite time horizon models such as ours, decision rules are likely nonstationary, and therefore we estimate the model separately for each HECM loan age (years since loan origination). The basis functions are Hermite polynomial sieves, and the order is selected by cross-validation as discussed in Section 3.

The non-stationarity assumption (Assumption 3 in Appendix A) is also supported in the reverse mortgage data. To test this, we fit a sieve multinomial logit model which is restricted to have same coefficients across different years from loan origination except for indicator variables for each loan age:

$$\log \frac{\sigma_t(s, a_{-1}, a)}{\sigma_t(s, a_{-1}, T)} = x(s, a_{-1})'\theta_a + \sum_{s=2}^{4} 1(s = t)\theta_{s,a}. \tag{7}$$

$x(s, a_{-1})$ has the same set of variables as the one used in the first step conditional choice probability estimation in the paper. $\theta_{s,a}$ can be interpreted as differences in the log odds ratio of choice a relative to the choice of termination between year s and year 1.

If there are 3 periods of data, t, $t+1$, $t+2$, such that $\theta_{t+1,a}$ and $\theta_{t+2,a}$ are significantly different for some a, Assumption 3 will be satisfied. Note that because (7) is a special case of models with more flexible specifications of time effects, the rejection of the hypothesis that the time effects are equal in (7) implies that the time effects are also not equal and that the conditional choice probabilities are not stationary for the more general model. Table 4 reports the estimates for $\theta_{t+2,a} - \theta_{t+1,a}$. Indeed, the differences are significant at the 5% level for the default choice for loan ages 2 vs. 1 and 3 vs. 2 and for the continue choice for loan ages 2 vs. 1 and 3 vs. 4. Wald tests on the joint equality of the time effects for the choices of continue, refinance, and default also reject the null hypothesis of stationarity. Therefore Assumption 3 is satisfied for both $t = 1, 2, 3$ and $t = 2, 3, 4$.[15]

Table 4. Non-stationarity of Choice Probabilities

		Loan Age 2 vs. 1	Loan Age 3 vs. 2	Loan Age 4 vs. 3
Continue/Terminate		-0.622^{***}	0.145	0.348^{***}
		(0.160)	(0.112)	(0.131)
Refinance/Terminate		0.434	-0.310	-0.395
		(0.327)	(0.270)	(0.331)
Default/Terminate		0.666^{***}	0.307^{**}	0.193
		(0.235)	(0.142)	(0.158)
Joint tests	Z statistic	82.42	8.43	15.25
	p-value	0.0000	0.0379	0.0016

This table reports tests on the equality of coefficients. In the upper panel, reported are differences between coefficients of pairs of loan age indicators for each choice in the sieve multinomial logit model. In parentheses are standard errors, which are computed based on the estimated variance-covariance matrix of the coefficients. Other variables are the same as the ones used in the first step conditional choice probability estimation. In the lower panel, Wald tests on the joint equality of pairs of loan age indicators for the choices of continue, refinance, and default are performed. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5 reports the within-sample fit of the HECM policy function estimates. The average predicted choice probabilities are compared with the data using the full sample, as well as sub-samples as defined by HECM characteristics and some key borrower

[15]We use observations with loan ages 2, 3 and 4 in the estimation for the discount factor, because there is more variation in observed choices in year 2 than in year 1.

state variables. Although these are non-targeted moments (i.e., we estimate the model using maximum likelihood rather than GMM based these moments), the predicted choice probabilities still capture the overall patterns in the data well. To further evaluate the fit of the model, we randomly divide the sample into two halves where the first half is used as a training sample to estimate the model parameters, and the model predictions are compared with the actual choices for observations in the second half sample. These results are reported in Table 12 in Appendix C, where we see that the model also has good out-of-sample prediction power.

In addition, we also use a sieve multinomial logit model to estimate the conditional choice probabilities for HECM take-up using the counseled sample. Starting in April 2009, both fixed rate and adjustable rate HECMs are available. Households who choose the fixed rate HECM receive the HECM proceeds as a lump sum, while adjustable rate HECM borrowers can select between different payment plans including a line of credit, tenure, term, and combinations thereof. Large upfront HECM credit utilization has been recognized as a significant risk factor for default, and we model that adjustable rate HECM borrowers are making a choice on whether they make large upfront draws. Large draw is defined as initial HECM credit utilization exceeding 80% of the credit limit. Because fixed rate HECMs were not available before April 2009, the available choices for households counseled before that date are not taking up an HECM, adjustable rate HECM with large upfront draw, and adjustable rate HECM with small upfront draw. Table 6 reports the within sample fit of the HECM take-up policy function estimates and shows that the estimated policy functions fit the data distribution well[16].

5.2. Structural Utility Function Estimates

The total value for a household consists of a choice-specific period utility, a continuation value conditional on the state variables and choice taken this period, and an i.i.d. type I extreme value error. Table 7 contains estimates of the per-period, choice-specific utility coefficients along with 95% bias-corrected bootstrap confidence intervals. The utility coefficients of the choices of continue, refinance and default are relative to those in the termination. Section 3 shows that observing two terminating actions allows us to identify the utility coefficients for every choice, rather than only the difference relative to some reference choice. Therefore in Table 7, the per-period utility coefficients of termination also vary with state variables. For the continue, refinance, and default decisions the reported estimates are differences in the coefficients relative to termination. In other words, the termination coefficients are levels while the continue, refinance, and default are differences.

[16] For simplicity, this table reports estimates only for the observations after April 1, 2009, when all HECM products became available.

Table 5. In-Sample Fit of Reduced Form HECM Policy Function Estimates

Sample	Termination Prediction	Termination Data	Refinance Prediction	Refinance Data	Default Prediction	Default Data
Unconditional						
All	1.43%	1.40%	0.37%	0.34%	4.09%	4.10%
By HECM Type						
Fixed Rate	1.27%	1.26%	0.33%	0.31%	4.79%	4.76%
Adjustable Rate	1.66%	1.61%	0.44%	0.39%	3.10%	3.15%
By Loan Age						
1	0.68%	0.67%	0.29%	0.29%	0.44%	0.44%
2	1.77%	1.75%	0.47%	0.47%	3.24%	3.25%
3	1.94%	1.87%	0.39%	0.35%	5.91%	5.92%
4	1.34%	1.33%	0.30%	0.23%	7.76%	7.77%
By Credit Score						
Q1	1.08%	0.99%	0.39%	0.36%	11.86%	12.09%
Q2	1.37%	1.32%	0.36%	0.33%	3.19%	3.29%
Q3	1.53%	1.63%	0.39%	0.39%	0.82%	0.69%
Q4	1.71%	1.67%	0.36%	0.30%	0.44%	0.26%
By Net Equity						
Q1	1.15%	1.03%	0.16%	0.16%	8.69%	8.68%
Q2	1.39%	1.35%	0.31%	0.20%	4.66%	4.67%
Q3	1.46%	1.67%	0.41%	0.46%	2.22%	2.20%
Q4	1.69%	1.56%	0.57%	0.55%	0.80%	0.83%
By Available HECM Credit						
Q1	1.45%	1.41%	0.40%	0.36%	5.87%	5.88%
Q2	1.38%	1.34%	0.33%	0.39%	0.49%	0.46%
Q3	1.33%	1.49%	0.25%	0.15%	0.19%	0.22%
Q4	1.43%	1.36%	0.37%	0.39%	0.05%	0.02%

This table shows the within-sample fit of the policy function estimates, both unconditionally and conditional on some explanatory variables. Q1–Q4 denote the first through fourth quartiles of the stated variables.

Table 6. In-Sample Fit of Reduced Form HECM Take-Up Policy Function Estimates

	FRM		ARM, Small Draw		ARM, Large Draw	
Sample	Prediction	Data	Prediction	Data	Prediction	Data
			Unconditional			
All	37.95%	37.95%	15.28%	15.28%	2.86%	2.86%
			By Year of Counseling			
2009	36.12%	36.19%	14.86%	14.65%	6.29%	6.34%
2010	37.67%	37.65%	17.23%	17.37%	2.83%	2.80%
2011	38.71%	38.71%	13.30%	13.20%	2.01%	2.03%
			By Age			
Q1	39.62%	38.99%	9.85%	9.90%	2.47%	2.40%
Q2	40.12%	40.87%	12.06%	12.07%	2.66%	2.80%
Q3	38.44%	38.47%	16.72%	17.22%	2.92%	2.71%
Q4	33.06%	32.97%	23.42%	22.76%	3.46%	3.64%
			By Income			
Q1	33.98%	33.87%	15.00%	15.12%	2.55%	2.61%
Q2	36.87%	36.62%	16.59%	16.67%	2.69%	2.61%
Q3	39.08%	39.22%	15.61%	15.47%	2.92%	2.85%
Q4	41.87%	42.09%	13.91%	13.85%	3.28%	3.37%
			By Credit Score			
Q1	33.89%	33.59%	6.68%	6.94%	2.98%	2.85%
Q2	39.28%	39.72%	10.39%	10.00%	3.03%	3.24%
Q3	41.07%	41.27%	18.02%	17.69%	3.01%	2.90%
Q4	37.55%	37.21%	26.13%	26.59%	2.41%	2.45%
			By Net Equity			
Q1	38.14%	38.72%	3.48%	3.37%	2.37%	2.10%
Q2	43.66%	42.80%	10.98%	11.27%	2.84%	3.42%
Q3	39.69%	40.03%	19.34%	19.26%	3.09%	2.89%
Q4	30.26%	30.22%	27.36%	27.26%	3.14%	3.02%

This table shows the within-sample fit of the policy function estimates, both unconditionally and conditional on some explanatory variables. Q1–Q4 denote the first through fourth quartiles of the stated variables. The sample is restricted to households counseled after April 1, 2009.

Table 7. Coefficient Estimates for Per-Period Payoffs

	Continue		Refinance		Default		Terminate	
Constant	6.487	(-7.332, 28.711)	-1.393	(-14.241, 23.637)	5.721	(-7.343, 28.684)	-8.910	(-48.740 , 14.790)
Hispanic	-0.073	(-2.164, 6.332)	0.096	(-2.032, 10.936)	-0.055	(-2.747, 4.423)	-0.276	(-7.906, 2.981)
Black	-0.194	(-2.191, 2.304)	0.361	(-2.047, 3.378)	-0.099	(-2.328, 2.702)	0.446	(-7.446, 3.056)
Single male	-0.721	(-2.072, 1.599)	0.992*	(-0.255, 4.409)	-0.697	(-2.236, 1.817)	0.709	(-3.112, 2.682)
Single female	-1.210	(-2.673, 1.019)	-1.451	(-2.972, 1.350)	-1.334	(-2.766, 1.002)	1.364	(-1.884, 3.270)
Income†	-6.354**	(-15.418, -1.981)	-7.759**	(-17.825, -1.213)	-7.915**	(-18.736, -2.792)	7.216**	(0.463, 18.435)
Property tax/income	-10.778*	(-18.981, 0.377)	-10.235*	(-23.170, 0.373)	-10.287*	(-20.923, 1.488)	12.605	(-4.098, 24.710)
Non-housing assets†	-0.006	(-0.069, 0.033)	-0.027	(-0.131, 0.045)	0.003	(-0.051, 0.049)	0.009	(-0.051, 0.095)
Fixed rate HECM	-0.931	(-3.930, 1.237)	-0.163	(-3.844, 2.211)	-0.157	(-2.620, 2.490)	1.145	(-2.016, 4.964)
First year credit utilization $> 80\%$	2.435*	(-0.032, 5.846)	3.777**	(1.289, 8.353)	3.683**	(0.410, 7.046)	-3.484*	(-8.193, 0.410)
Credit score	-0.005	(-0.035, 0.012)	-0.007	(-0.042, 0.011)	-0.011	(-0.037, 0.005)	0.009	(-0.022, 0.064)
Available revolving credit†	0.368**	(0.260, 0.807)	0.367**	(0.294, 0.765)	0.079	(-0.066, 0.449)	-0.622**	(-1.925, -0.305)
Revolving & installment debt†	-0.244	(-0.861, 0.224)	-0.124	(-0.926, 0.383)	-0.391	(-1.250, 0.189)	0.512	(-0.625, 1.971)
Net equity†	0.130**	(0.045, 0.187)	0.189**	(0.060, 0.310)	0.133**	(0.040, 0.259)	-0.182**	(-0.284, -0.036)
Negative net equity†	1.036	(-2.222, 6.412)	0.379	(-1.250, 21.929)	0.882	(-252.412, 5.901)	-0.063	(-0.324, 0.214)
Excess credit†	-0.220	(-4.468, 1.767)	-0.977	(-3.938, 1.076)	-8.767	(-125.189, 3.239)	-0.107	(-0.487, 0.077)
Available HECM credit†	0.011	(-0.047, 0.059)	-0.103	(-0.288, 0.034)	-0.298**	(-0.404, -0.114)	-0.081	(-0.162, 0.030)
Unpaid T&I balance	-6.323	(-44.796, 188.837)	-0.210	(-346.356, 5.948)			-0.867	(-11.336, 2.451)
$\Delta \text{IPL}_{\text{refi}}$			0.145**	(0.019, 0.628)				
Loan rate - current rate			0.357	(-3.366, 3.932)				
Discount factor b		0.898**				(0.707, 1.000)		

The reported coefficients for the payoffs of continue, refinance and default are relative to those in the termination payoff. 95% bias-corrected bootstrap confidence intervals in parentheses (1,200 replications).
* significance at 10%. ** significance at 5%. † Monetary variables are reported in units of $10,000.

Table 8. Coefficient Estimates for Per-Period Payoffs, with Zero Restrictions on Termination Payoffs

	Continue		Refinance		Default		Terminate
Constant	0.239	(-1.456, 1.191)	-8.151**	(-13.395, -4.584)	-0.775	(-5.822, 2.141)	0
Hispanic	-0.158	(-0.496, 0.354)	-0.542	(-1.910, 0.372)	-0.270	(-1.839, 0.628)	0
Black	0.230	(-0.405, 0.592)	0.347	(-1.055, 1.137)	0.300	(-0.767, 0.728)	0
Single male	0.057	(-0.094, 0.609)	0.510	(-0.240, 1.760)	0.022	(-0.376, 0.477)	0
Single female	0.048	(-0.111, 0.327)	0.153	(-0.531, 1.194)	-0.048	(-0.455, 0.565)	0
Income†	-0.666	(-1.018, 0.094)	-1.562	(-7.336, 2.843)	-2.164**	(-4.209, -0.810)	0
Property tax/income	-0.569	(-1.413, 1.713)	-0.313	(-4.483, 4.348)	0.205	(-3.358, 2.278)	0
Non-housing assets†	-0.001	(-0.002, 0.002)	-0.002	(-0.007, 0.012)	0.003	(-0.005, 0.010)	0
Fixed rate HECM	0.113	(-0.273, 0.465)	-0.151	(-1.195, 0.779)	0.856**	(0.185, 1.946)	0
First year credit utilization $>$ 80%	0.250	(-0.141, 0.502)	1.238**	(0.295, 2.108)	1.026	(-0.793, 1.905)	0
Credit score	0.000	(-0.001, 0.001)	-0.001	(-0.005, 0.005)	-0.005	(-0.008, 0.002)	0
Available revolving credit†	0.011	(-0.020, 0.039)	0.043	(-0.036, 0.136)	-0.204**	(-0.311, -0.101)	0
Revolving & installment debt†	0.037**	(0.007, 0.060)	0.041	(-0.040, 0.107)	-0.048	(-0.084, 0.027)	0
Net equity†	0.004	(-0.007, 0.011)	0.038*	(-0.001, 0.070)	-0.029	(-0.055, 0.027)	0
Negative net equity†	0.846	(-2.295, 7.090)	0.283	(-1.436, 21.835)	0.520	(-246.963, 4.724)	0
Excess credit†	-0.067	(-4.735, 1.771)	-0.810	(-3.532, 1.342)	-8.685	(-122.881, 3.767)	0
Available HECM credit†	0.044**	(0.009, 0.063)	0.036	(-0.054, 0.095)	-0.152**	(-0.203, -0.051)	0
Unpaid T&I balance	-6.583	(-41.179, 209.661)	-0.325	(-360.621, 7.167)			0
ΔIPL_{refi}			0.054	(-0.079, 0.107)			0
Loan rate - current rate			1.306	(-2.138, 2.933)			0
Discount factor b			0.898**			(0.707, 1.000)	

95% bias-corrected bootstrap confidence intervals in parentheses (1,200 replications).
* significance at 10%. ** significance at 5%. † Monetary variables are reported in units of $10,000.

We do this to mirror the situation when termination is chosen as the baseline choice with coefficients fixed at zero. To show that coefficient estimates depend on the normalization restriction, Table 8 shows the results when the termination payoff is indeed restricted to be zero everywhere.

The utility functions represent the per-period sum of benefits and losses of taking actions on the HECM, as functions of the state variables. Because of heterogeneity, households naturally have higher or lower utilities than others for each choice. For borrowers who continue, default, or refinance, the utility function represents the net current period benefits of taking each action. For a borrower who terminates the HECM, the period utility for termination represents a one-time payoff since no utility is received in future periods. The higher the termination value relative to the payoff from other choices, the more likely that the HECM will be terminated.

At 5% significance level, borrowers who receive more value from termination are those with higher income but lower net equity or unused revolving credit. At HECM loan termination, the loan balance has to be repaid, which can be satisfied by selling the home, a deed in lieu, a short sale, foreclosure, or using the personal funds of the borrowers or their heirs. Although we do not have data on how HECMs are repaid, anecdotal evidence indicates that the most common way of repayment is through selling the home, and consistent consistent with this, borrowers with more net equity find it less attractive to move out and have lower value from termination. At the same time, because HECM credit limits do not change with declines in house prices, HECM borrowers are insured against house price declines to the extent of their HECM credit limits, and the insurance value is greater the more the home price drops below the HECM credit limit. Borrowers with more unused revolving credit may have more financial resources to cope with unexpected events such as increased medical spending, and are better able to stay in their homes and age in place. Note that we include both net equity (the level, whether positive or negative, say NE_{it}) and negative net equity (the absolute value of the negative part, NNE_{it}) as state variables. Hence, the total effect of net equity on choice-specific utility for a household is $\rho_{\mathrm{NE}}\mathrm{NE}_{it} - 1\{\mathrm{NE}_{it} < 0\}\rho_{\mathrm{NNE}}\,|\mathrm{NE}_{it}|$. For households with positive net equity, an increase in its value increases the period payoff of the choices of continue, refinance and default relative to immediate termination, hence makes it more attractive to continue with HECMs. At the same time, if net equity drops below zero, a further decrease will increase the period payoff of continue and default relative to immediate termination. Unlike forward mortgages, with no monthly payment on HECM loan balance, negative equity does not make it more costly to continue with HECM loans. The variable of excess credit measures the extent to which the home value drops below the HECM credit limit, and its point estimate indicates that a higher excess credit decreases the period payoffs of default, and to a lesser extent, refinance and continue, relative to immediate termination. However, because the number of observations with negative equity or nonzero excess credit is rather limited, the estimates of the coefficients of negative net equity and excess credit are quite noisy and are not statistically significant, indicating that there is no strong evidence to support that the choice payoffs vary nonlinearly with home value as it crosses the thresholds of HECM credit limit and loan balance, and therefore households may not strategically terminate their HECM loans. Davidoff and Wetzel (2014) also conclude that HECM borrowers do not behave strategically when HECMs terminate, as their HECM credit utilization behavior do not vary when home values drop below HECM credit limits.

Borrowers with lower income or who have used more of their HECM credit lines receive higher value from default relative to other choices, which means that if the continuation value is fixed, these households are more likely to default this period. HECM borrowers are more likely to refinance if doing so can increase the amount of borrowing. A higher amount of unused revolving credit increases the period payoff of continue and refinance relative to other choices, indicating that having access to additional financial resources are complementary to continuing with the HECMs.

The utility estimates depend on the normalization used, as restricting the termination payoff to zero significantly changes many coefficient estimates (Table 8). For example, according to the unrestricted estimates (Table 7), higher net equity increases the period payoff of default relative to termination and is significant at 5%, but with the restricted model (Table 8), the coefficient becomes negative and insignificant. Therefore the structural parameter estimates can be biased if the termination payoff is restricted to zero but in reality it varies with the state variables.

In addition to the utility coefficients, we also estimate the discount factor β which is a product of the factor of time discounting b and the survival probability. Our estimate for b is 0.898 and is within the range of previous results in the literature on discount factors for the elderly, which range from 0.67–0.90 (Trostel and Taylor, 2001; Harrison, Lau, and Williams, 2002; Read and Read, 2004). Factoring in the survival probability, the effective discount rate will be lower for older or single borrowers.

5.3. Ex-Ante Value Function Estimates

The ex-ante value function in (1) measures the expected discounted future value from HECM loans for existing HECM borrowers, can be taken as a measure of the value households place on the HECM program. Note that this does not include the one-time payoff that arises from taking up the loan, such as paying off some other debt. To include this, the HECM take up choice should be part of a household consumption, saving and portfolio choice model such as (Nakajima and Telyukova, 2017), and is beyond the scope of our model and data. The ex-ante value function is nonlinear, but to summarize how this value varies across households of different types, we report in Table 9 the results of a linear regression of $V_t(s_{it})$ on state variables. This allows us to examine how households' valuations for the HECM loans vary with household and loan characteristics and economic conditions.

At 5% significance level, the net benefits of HECMs are larger if the borrower has a lower available revolving credit amount, a lower net equity amount or experiences a recent house price decline in the borrower's MSA. The ex-ante value is also larger for younger borrowers, while there are no statistically significant nonlinear effects for borrowers who just become eligible for HECMs (age between 62 and 64). If the adjustable rate HECM is structured as a credit line, the unused portion of the credit line grows at the compounding rate that depends on a market benchmark. Not surprisingly, the ex-ante value is higher for a higher average interest rate for adjustable rate HECMs. Notice that a higher income is associated with a higher ex-ante value with the effect significant at 10%. The utility coefficient estimates indicate that at 5% significance level, a higher income increases the payoff of termination, while decreases the period payoffs of the choices of continue, refinance and default, and as a result of these opposite effects, the overall effect of income on the ex-ante value is only weakly positively significant.

To provide some external validation of the welfare measure, Figure 2 shows that there is a negative relationship between state average ex-ante values and consumer dissatisfaction with reverse mortgages. Note that because the utility from termination is not restricted to zero, the ex-ante value may be negative which arises when the expected present value from the HECM loan is outweighed by the cost of termination. We compute average ex-ante values for households in different states using the estimated model parameters. Consumer dissatisfaction is measured by the number of complaints with reverse mortgages that have been received by the Consumer Financial Protection Bureau[17], divided by the number of HECM loans in 2012. The solid line shows the fitted values from a linear regression[18]. Its slope is -0.007 with p-value 0.026.

5.4. Welfare Implications of Borrower Eligibility Requirements

In our policy experiments we study the effects of imposing certain underwriting criteria on borrower behavior and welfare. A significant program change in recent years is the introduction of the financial assessment requirements effective April 27, 2015 which are designed, among other things, to improve the financial position of the MMIF through decreasing the rate of property tax & insurance defaults (Mortgagee Letter 2015-06).

[17] The first complaint in the database was received in December 1, 2011.

[18] Hawaii is excluded as its value is an outlier in the figure.

Table 9. Regression of Ex-Ante Values on Borrower Characteristics

Dependent Variable	$V_t(s_{it}, a_{it,-1})$	
	Coeff.	95% CI
Hispanic	-0.455	(-12.464 , 2.828)
Black	1.309	(-11.306 , 3.900)
Single male	0.354	(-3.469 , 2.317)
Single female	1.167	(-2.081 , 3.147)
Income†	6.306*	(-0.497 , 17.229)
Property tax/income	11.855	(-4.985 , 23.864)
Non-housing assets†	0.010	(-0.050 , 0.096)
Fixed rate HECM	1.393	(-1.727 , 5.588)
First year credit utilization > 80%	-3.184*	(-8.048 , 0.663)
Credit score	0.008	(-0.021 , 0.066)
Available revolving credit†	-0.624**	(-1.935 , -0.296)
Revolving & installment debt†	0.524	(-0.595 , 1.973)
Net equity†	-0.196**	(-0.296 , -0.056)
Negative net equity†	0.300	(-0.103 , 2.129)
Excess credit†	0.032	(-0.532 , 0.541)
Available HECM credit†	-0.049	(-0.128 , 0.061)
Unpaid T&I balance	1.329	(-4.621 , 8.388)
$\Delta\text{IPL}_{\text{refi}}$	0.000	(-0.044 , 0.036)
Young borrower	0.005	(-0.516 , 0.538)
Age	-0.092**	(-0.149 , -0.025)
Age^2	0.000**	(0.000 , 0.001)
HPI change	-3.551**	(-6.014 , -0.470)
HPI change, 1 year lag	-1.052	(-3.476 , 2.336)
HPI change, 2 year lag	2.197*	(-0.517 , 3.947)
Average interest rate (ARM)	0.705**	(0.343 , 0.994)
Average interest rate (FRM)	-0.010	(-0.189 , 0.165)
Defaulted in T&I last year	-0.223	(-0.905 , 0.616)
In default for two years	1.319**	(0.412 , 6.173)
Loan age		
2	-0.769**	(-1.154 , -0.219)
3	-0.515*	(-0.980 , 0.023)
4	-0.218	(-0.855 , 0.319)

The reported coefficients are for a linear regression of $V_t(s_{it}, a_{it,-1})$ on s_{it} and other variables. Since V_t is not a linear function, these estimates reflect average relationships rather than marginal effects. 95% bias-corrected bootstrap confidence intervals in parentheses (1,200 replications).
* significance at 10%. ** significance at 5%. † Monetary variables are reported in units of $10,000.

Figure 2. Average Ex-Ante Values and Borrower Satisfaction

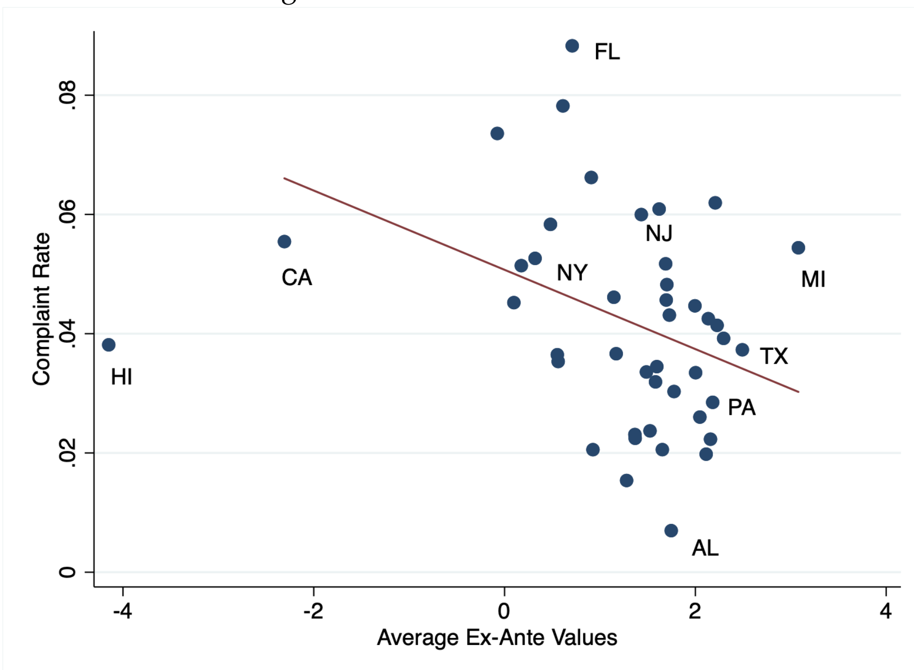

Previous studies have examined the effects of imposing underwriting criteria on default rates (Moulton et al., 2015), but the welfare cost of limiting program participation is not yet fully understood.

Specifically, we simulate the effects of imposing borrower eligibility requirements on credit scores, income and age. The results are summarized in Table 10. The first column indicates the levels of the cutoff values for the credit requirement, initial income requirement, and initial age requirement respectively. For example, under the initial credit requirement of 490, borrowers whose credit scores are below 490 become ineligible for HECMs. The statistics reported are based on our sample of HECM borrowers that are still eligible for HECMs. The next three columns report the annual rates of termination, refinance, and default decisions in the model, averaged over all households and all four years of our sample. For both the credit and income requirements, the default rate declines considerably. The average termination rate becomes slightly higher while the average refinance rate is largely unaffected by the policies. Surprisingly, the initial income requirement also significantly reduces the fraction of households with negative net equity (fifth column) as well as the amount of their negative equity (sixth column, in $1 million units), but the opposite happens for the initial age requirement.

The cost of these policies is, of course, a decline in HECM volume due to households being excluded and a decrease in total borrower welfare. We report the *total* of the ex-ante values $V_t(s_{it})$ over borrowers in the sample, averaged over the four years of our data, under each scenario in the seventh column of Table 10 and the percentage change in this welfare measure in the eighth column. The next two columns report the ex-ante values that are obtained when we restrict the termination payoff to zero. The final two columns measure the reduction in HECM volume (in number of households and the percentage changes) due to these participation constraints. With a more stringent requirement on credit scores, income or age, more households become ineligible for HECMs, and the average borrower welfare as measured by the ex-ante value drops, as percentage decreases in ex-ante values are greater than the decreases in HECM volumes.

Compared with the initial credit requirement, imposing an initial income requirement would reduce the default rate less for a similar reduction in HECM volume, and its welfare cost is greater. To see that the credit requirement is more effective, in terms of welfare, at reducing defaults, we can compare the implications of a credit score requirement of 550—a threshold that excludes 10.81% households in our sample—with those of an income requirement at 2 times the Federal Poverty Level. The baseline default rate before the restrictions are imposed is 3.97%. The income requirement decreases this to 3.08% yet it would exclude 43.38% of borrowers in our sample, thus reducing total ex ante value from 87,718 to 40,033, a 54.36% reduction. In contrast, the credit requirement reduces the default rate slightly more, to 2.94%, and it does so by excluding fewer borrowers, only 10.81%. Furthermore, the drop in ex ante value is also less, at 17.25%, so the welfare cost is lower. Imposing an initial age requirement is even less effective in reducing defaults. On the other hand, the income requirement is better in terms of reducing negative net equity.

The normalized ex-ante values which restrict termination payoffs to zero mask the heterogeneous impacts of the borrower eligibility requirements in different regions. Figures 3 and 4 show how the initial credit (minimum credit score of 550) and income (2X federal poverty line) requirements impact different regions differently in terms of the loss in borrower welfare as measured by the ex-ante values. We report the relative changes in the number of eligible loans, total ex-ante values, and total ex-ante values where the termination utility is restricted to zero, for states where the sample has at least 30 observations. The predicted decreases in the normalized ex-ante values are similar to the decreases in the number of eligible loans, while if the termination utility is not normalized, states of Colorado and New York would experience much larger average percentage decreases in ex-ante values, suggesting that borrowers in these two states would be hit the hardest under the initial credit and income eligibility constraints.

Besides borrowers, lenders and the MMIF are also important participants in this market. Abstracting away from investors in the market for HECM mortgage-backed securities, HECM lenders earn interest income and origination fees on HECM loans. Aside from the

Figure 3. Initial Credit Requirements

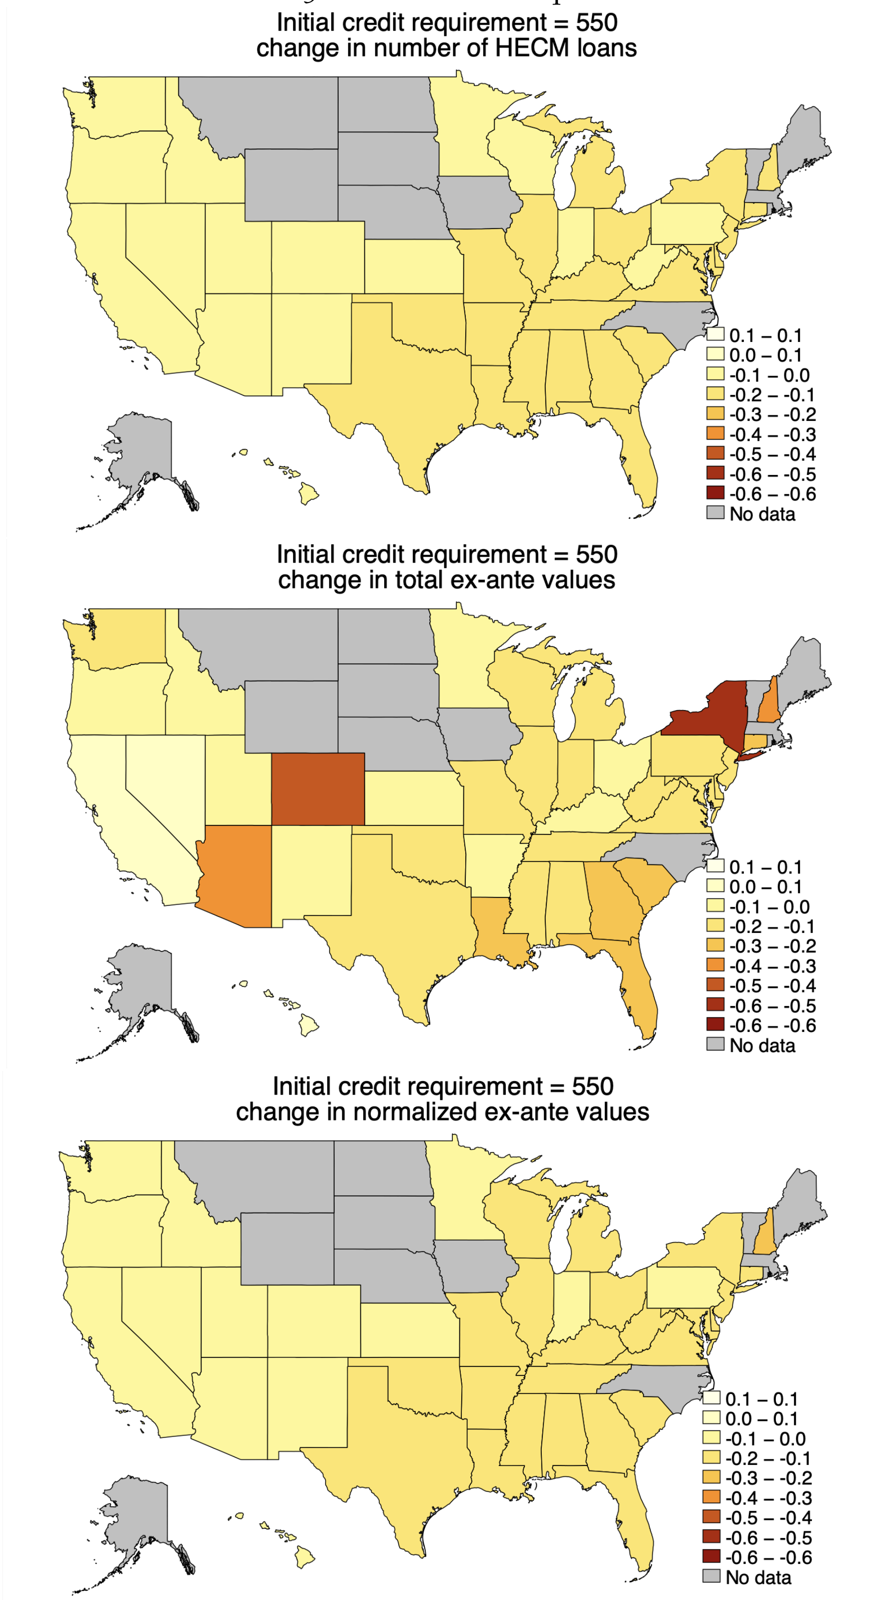

Figure 4. Initial Income Requirements

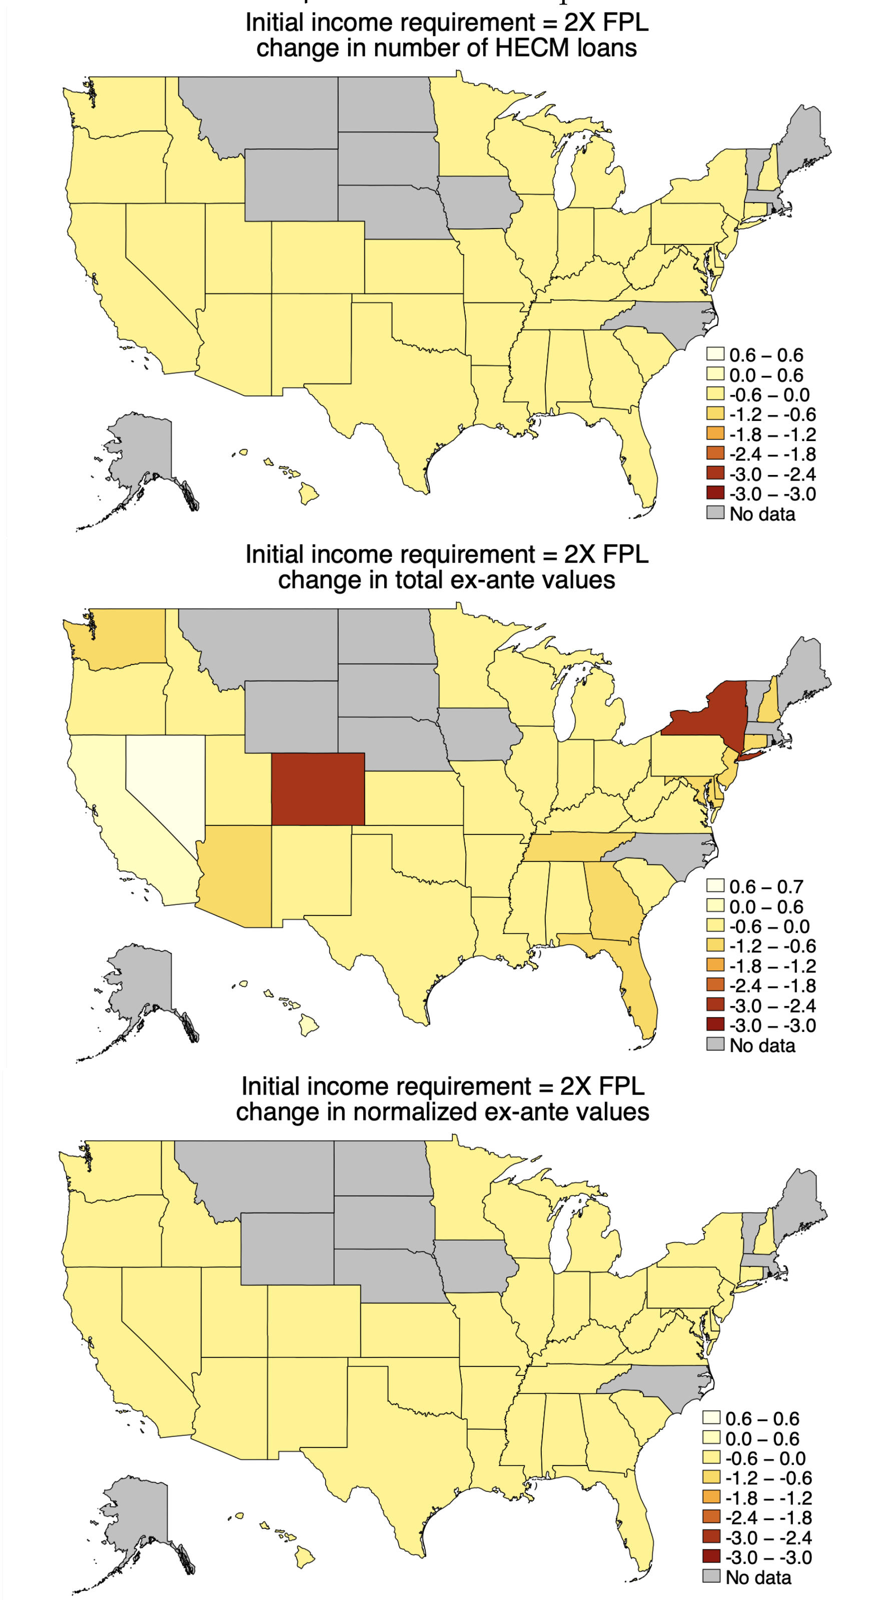

Table 10. Counterfactual Imposition of Borrower Eligibility Requirements

	Termination	Refinance	Default	Negative Equity		Ex-Ante Value		Ex-Ante Value Excluding Termination		HECM Volume	
	%	%	%	% HH	Total $1M	Total	%Δ	Total	%Δ	# HH	%Δ
	1.41%	0.34%	3.97%	2.36%	-19.28	87,718	–	215,259	–	11,656	–
Initial credit requirement											
490	1.43%	0.35%	3.63%	2.34%	-18.63	83,335	-5.00%	207,559	-3.58%	11,355	-2.58%
520	1.44%	0.35%	3.36%	2.36%	-18.26	78,452	-10.56%	198,872	-7.61%	10,978	-5.82%
550	1.45%	0.35%	2.94%	2.35%	-17.25	72,589	-17.25%	187,104	-13.08%	10,396	-10.81%
580	1.47%	0.35%	2.54%	2.37%	-15.63	64,950	-25.96%	174,016	-19.16%	9,825	-15.71%
610	1.50%	0.35%	2.14%	2.35%	-13.89	57,275	-34.71%	160,430	-25.47%	9,165	-21.37%
Initial income requirement											
1 × FPL	1.42%	0.34%	3.69%	2.14%	-16.21	76,300	-13.02%	197,496	-8.25%	10,769	-7.61%
1.25 × FPL	1.43%	0.34%	3.51%	2.03%	-14.28	67,695	-22.83%	180,625	-16.09%	9,897	-15.09%
1.5 × FPL	1.44%	0.34%	3.33%	1.88%	-11.73	56,209	-35.92%	157,634	-26.77%	8,729	-25.11%
1.75 × FPL	1.46%	0.34%	3.20%	1.76%	-9.47	47,134	-46.27%	136,493	-36.59%	7,609	-34.72%
2 × FPL	1.47%	0.34%	3.08%	1.71%	-8.15	40,033	-54.36%	117,816	-45.27%	6,600	-43.38%
Initial age requirement											
No non-borrowing spouse below 62	1.43%	0.34%	3.93%	2.37%	-18.97	82,374	-6.09%	206,462	-4.09%	11,351	-2.62%
63	1.45%	0.34%	3.86%	2.42%	-18.67	79,300	-9.60%	198,603	-7.74%	10,963	-5.95%
64	1.49%	0.34%	3.79%	2.56%	-17.96	72,503	-17.35%	182,766	-15.09%	10,153	-12.89%
65	1.53%	0.33%	3.67%	2.70%	-17.78	64,240	-26.77%	167,473	-22.20%	9,478	-18.69%
66	1.56%	0.33%	3.64%	2.80%	-17.24	59,799	-31.83%	157,054	-27.04%	8,910	-23.56%
67	1.59%	0.33%	3.60%	2.86%	-16.51	55,980	-36.18%	146,879	-31.77%	8,371	-28.18%

The baseline sample (11,656) is slightly smaller than the full sample because the sample is restricted to observations with credit scores available at HECM loan closing. Initial income requirement is measured in terms of the Federal Poverty Level (FPL). Initial age requirement is based on the minimum age of the borrower(s). Reported rates and valuations are four-year averages. Negative net equity values reported are the percentage of households with negative equity (in any amount) and the total amount of household net equity, in $1 million units, for households with negative equity. Ex-ante value is the total ex-ante value of all HECM households measured in utils. HECM volume is measured in terms of the number of counseled households who choose to take-up a HECM in the baseline and are still eligible for HECMs with the eligibility requirement imposed.

risk associated with the uncertainty on when the loan is terminated and paid off, lenders may experience losses on loans that are terminated in adverse circumstances. Based on the HECM loans purchased by Fannie Mae, Begley, Fout, LaCour-Little, and Mota (2019) find that the average loss is 6% of the loan balance at the time of loan termination. For the MMIF, it earns insurance premiums on active loans but may need to pay claims if the loan balance exceeds the value of the home when the loan terminates. The policies considered in Table 10 reduce the occurrence of adverse terminations but at the same time reduce the loan volume.

Whether the lender and the MMIF find alternative policies advantageous depends on the tradeoff between loan volumes and adverse terminations over the life of a loan. Although this is beyond the scope of our model, we report in Table 11 the ratios of the predicted relative changes in loan volumes and total ex-ante values over reductions in default rates and the total amount of negative equity, for the policies considered in Table 10. A lower ratio indicates a lower cost for a given improvement in loan performance. Table 11 shows that to reduce default rates, eligibility requirements based on credit scores are most cost effective, while to decrease the severity of negative equity, income requirements are most cost effective.

6. Conclusion

The contributions of this paper are twofold. We show that both the utility function and the discount factor in a dynamic structural discrete choice model can be fully identified when distinct terminating actions exist. With this result, welfare and counterfactual analysis are more robust as there is no need to impose an ad hoc identifying assumption or "normalization." We then carry out an empirical analysis of the HECM program. Our estimates quantify the effects of factors that influence major HECM decisions, including refinance, default, and termination. Factors associated with higher HECM values include having less access to revolving credit, a lower net equity amount, and younger age. We show how various factors influence household welfare and illustrate the welfare cost of policies that restrict program eligibility, with the aim of reducing defaults and adverse terminations.

References

Abbring, J. and O. Daljord (2018). Identifying the discount factor in dynamic discrete choice models. Working paper, Tilburg University.

Aguirregabiria, V. (2005). Nonparametric identification of behavioral responses to counter-

Table 11. Evaluations of Borrower Eligibility Requirements

	%Δ Volume ÷Δ Default Rates	%Δ Value ÷Δ Default Rates	%Δ Volume ÷%Δ NE	%Δ Value ÷%Δ NE
Initial credit requirement				
490	**7.57**	**14.65**	0.77	1.49
520	**9.59**	**17.41**	1.10	2.00
550	**10.51**	**16.76**	1.03	1.64
580	**10.98**	**18.14**	0.83	1.37
610	**11.68**	**18.97**	0.76	1.24
Initial income requirement				
$1 \times$ FPL	26.91	46.03	**0.48**	**0.82**
$1.25 \times$ FPL	32.41	49.03	**0.58**	**0.88**
$1.5 \times$ FPL	38.91	55.66	**0.64**	**0.92**
$1.75 \times$ FPL	44.99	59.95	**0.68**	**0.91**
$2 \times$ FPL	48.76	61.10	**0.75**	**0.94**
Initial age requirement				
No non-borrowing spouse below 62	58.84	137.00	1.64	3.81
63	54.02	87.20	1.89	3.05
64	69.48	93.46	1.89	2.54
65	62.11	88.97	2.40	3.43
66	70.37	95.07	2.23	3.01
67	75.45	96.86	1.96	2.52

%Δ Volume: percentage change in the number of eligible HECM loans relative to the number of HECM loans in the observed sample. Δ Default Rates: reduction in average default rates. %Δ Value: percentage change in the total ex-ante values of eligible HECM borrowers. %Δ NE: percentage change in the total negative equity amount relative to that of the observed sample.

factual policy interventions in dynamic discrete choice processes. *Economics Letters 87*, 393–398.

Aguirregabiria, V. (2010). Another look at the identification of dynamic discrete decision processes: An application to retirement behavior. *Journal of Business and Economic Statistics 28*, 201–218.

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

Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica 75*, 1–53.

Aguirregabiria, V. and P. Mira (2010). Dynamic discrete choice structural models: a survey. *Journal of Econometrics 156*, 38–67.

Aguirregabiria, V. and J. Suzuki (2014). Identification and counterfactuals in dynamic models of market entry and exit. *Quantitative Marketing and Economics 12*, 267–304.

Ai, C. and X. Chen (2003). Efficient estimation of models with conditional moment restrictions containing unknown functions. *Econometrica 71*, 1795–1843.

Andrews, D. (1991). Asymptotic normality of series estimators for nonparametric and semiparametric regression models. *Econometrica 59*, 307–345.

Arcidiacono, P. and R. A. Miller (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica 79*, 1823–1867.

Arcidiacono, P. and R. A. Miller (2015). Identifying dynamic discrete choice models off short panels. Working paper, Duke University.

Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica 75*, 1331–1370.

Bajari, P., C. S. Chu, D. Nekipelov, and M. Park (2016). Identification and semiparametric estimation of a finite horizon dynamic discrete choice model with a terminating action. *Quantitative Marketing and Economics 14*, 271–323.

Bajari, P., H. Hong, and D. Nekipelov (2013). Game theory and econometrics: A survey of some recent research. In D. Acemoglu, M. Arellano, and E. Dekel (Eds.), *Advances in Econometrics: Tenth World Congress*, Volume 3, Chapter 1, pp. 3–52. Cambridge: Cambridge University Press.

Bayer, P., R. McMillan, A. Murphy, and C. Timmins (2016). A dynamic model of demand for houses and neighborhoods. *Econometrica 84*, 893–942.

Begley, J., H. Fout, M. LaCour-Little, and N. Mota (2019). Home equity conversion mortgages: The secondary market investor experience. *Journal of Housing Economics*.

Blevins, J. R. (2016). Sequential Monte Carlo methods for estimating dynamic microeconomic models. *Journal of Applied Economics 31*, 773–804.

Blundell, R., X. Chen, and D. Kristensen (2007). Semi-nonparametric IV estimation of shape-invariant Engel curves. *Econometrica 75*(6), 1613–1669.

Buchholz, N., M. Shum, and H. Xu (2016). Semiparametric estimation of dynamic discrete choice models. Working paper, Princeton University.

Canay, I. A., A. Santos, and A. M. Shaikh (2013). On the testability of identification in some non-parametric models with endogeneity. *Econometrica 81*, 2535–2559.

Chen, X. (2007). Large sample sieve estimation of semi-nonparametric models. In J. J. Heckman and E. E. Leamer (Eds.), *Handbook of Econometrics, Volume 6B*, pp. 5549–5632. Amsterdam: North-Holland.

Chen, X., V. Chernozhukov, S. Lee, and W. Newey (2014). Local identification of nonparametric and semiparametric models. *Econometrica 82*, 785–809.

Chou, C. (2016). Identification and linear estimation of general dynamic programming discrete choice models. Working paper, University of Southern California.

Chung, D., T. Steenburgh, and K. Sudhir (2014). Do bonuses enhance sales productivity? A dynamic structural analysis of bonus-based compensation plans. *Marketing Science 33*, 165–187.

Darolles, S., Y. Fan, J.-P. Florens, and E. Renault (2011). Nonparametric instrumental regression. *Econometrica 79*, 1541–1565.

Davidoff, T. (2014). Reverse mortgage demographics and collateral performance. Technical report, University of British Columbia.

Davidoff, T. (2015). Can "high costs" justify weak demand for the home equity conversion mortgage? *Review of Financial Studies 28*, 2364–2398.

Davidoff, T. and G. M. Welke (2007). Selection and moral hazard in the reverse mortgage market. Working paper, UC Berkeley, Hass.

Davidoff, T. and J. Wetzel (2014). Do reverse mortgage borrowers use credit ruthlessly. Working paper, University of British Columbia.

Fang, H., Y. S. Kim, and W. Li (2016). The dynamics of subprime adjustable-rate mortgage default: A structural estimation. Working paper, University of Pennsylvania.

Fang, H. and Y. Wang (2015). Estimating dynamic discrete choice models with hyperbolic discounting, with an application to mammography decisions. *International Economic Review 56*, 565–596.

Harrison, G. W., M. I. Lau, and M. B. Williams (2002). Estimating individual discount rates in Denmark: A field experiment. *American Economic Review 92*(5), 1606–1617.

Haurin, D., S. Moulton, and W. Shi (2018). The accuracy of senior households' estimates of home values: Application to the reverse mortgage decision. *Real Estate Economics 46*(3), 655–697.

Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies 60*, 497–529.

Integrated Financial Engineering (2015). Actuarial review of the Federal Housing Administration Mutual Mortgage Insurance Fund HECM loans for fiscal year 2015. Technical report, U.S. Department of Housing and Urban Development.

Kalouptsidi, M., P. T. Scott, and E. Souza-Rodrigues (2016). Identification of counterfactuals in dynamic discrete choice models. Working paper, Princeton University.

Keane, M. and K. I. Wolpin (1994). The solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte Carlo evidence. *Review of Economics and Statistics 76*, 648–672.

Keane, M. and K. I. Wolpin (1997). The career decisions of young men. *Journal of Political Economy 105*, 473–522.

Komarova, T., F. Sanches, D. Silvia Junior, and S. Srisuma (2016). Joint analysis of the discount factor and payoff analysis factor payoff parameters in dynamic discrete choice games games. Working paper, London School of Economics.

Ma, C. (2014). Estimating a dynamic discrete choice model with partial observability for household mortgage default and prepayment behaviors. Working paper, Ohio State University.

Magnac, T. and D. Thesmar (2002). Identifying dynamic discrete decision processes. *Econometrica 70*, 801–816.

Mastrobuoni, G. and D. A. Rivers (2016). Criminal discount factors and deterrence. Working paper, University of Western Ontario.

Miller, R. A. (1984). Job matching and occupational choice. *Journal of Political Economy 92*, 1086–1120.

Moulton, S., D. R. Haurin, and W. Shi (2015). An analysis of default risk in the home equity conversion mortgage (HECM) program. *Journal of Urban Economics 90*, 17–34.

Nakajima, M. and I. A. Telyukova (2017). Reverse mortgage loans: A quantitative analysis. *Journal of Finance 72*, 911–950.

Newey, W. (1997). Convergence rates and asymptotic normality for series estimators. *Journal of Econometrics 79*, 147–168.

Newey, W. K. and J. L. Powell (2003). Instrumental variable estimation of nonparametric models. *Econometrica 71*, 1565–1578.

Norets, A. and S. Takahashi (2013). On the surjectivity of the mapping between utilities and choice probabilities. *Quantitative Economics 4*, 149–155.

Norets, A. and X. Tang (2014). Semiparametric inference in dynamic binary choice models. *Review of Economic Studies 81*, 1229–1262.

Pakes, A. (1986). Patents as options: Some estimates of the value of holding European patent stocks. *Econometrica 54*, 755–784.

Pesendorfer, M. and P. Schmidt-Dengler (2007). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies 75*, 901–928.

Read, D. and N. L. Read (2004). Time discounting over the lifespan. *Organizational Behavior and Human Decision Processes 94*(1), 22–32.

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

Rust, J. (1994). Structural estimation of Markov decision processes. In R. F. Engle and D. L. McFadden (Eds.), *Handbook of Econometrics*, Volume 4, Amsterdam, pp. 3081–3143. North Holland.

Srisuma, S. and O. Linton (2012). Semiparametric estimation of Markov decision processes with continuous state space. *Journal of Econometrics 166*, 320–341.

Szymanoski, E. J., J. C. Enriquez, and T. R. DiVenti (2007). Home equity conversion mortgage terminations: Information to enhance the developing secondary market. *Cityscape*.

Taber, C. R. (2000). Semiparametric identification and heterogeneity in discrete choice dynamic programming models. *Journal of Econometrics 96*, 201–229.

Trostel, P. and G. Taylor (2001). A theory of time preference. *Economic Inquiry 39*(3), 379–395.

U.S. Department of Housing and Urban Development. Mortgagee letter 2013-27.

U.S. Department of Housing and Urban Development. Mortgagee letter 2014-11.

U.S. Department of Housing and Urban Development. Mortgagee letter 2015-06.

U.S. Department of Housing and Urban Development. Mortgagee letter 2015-10.

U.S. Department of Housing and Urban Development. Mortgagee letter 2015-11.

U.S. Department of Housing and Urban Development. Mortgagee letter 2016-07.

U.S. Government Accountability Office (2019, September). Reverse mortgages: FHA needs to improve monitoring and oversight of loan outcomes and servicing.

Wolpin, K. I. (1984). An estimable dynamic stochastic model of fertility and child mortality. *Journal of Political Economy 92*, 852–874.

Wong, W. H. and X. Shen (1995). Probability inequalities for likelihood ratios and convergence rates of sieve MLEs. *Annals of Statistics 23*, 339–362.

A. Semiparametric Identification with Multiple Terminating Actions

In this section, we consider semiparametric identification and estimation of a finite-horizon dynamic discrete choice model with multiple terminating actions. We show that the presence of distinct terminating actions has substantial identifying power and leads directly to identification of the entire utility function without the need to impose an ad hoc "normalization". The model is directly motivated by the empirical setting we consider, where households may terminate directly or be forced to terminate by remaining in default for three consecutive periods. To our knowledge, this paper is the first to consider semiparametric identification of such models.

First we show that the main model primitives of interest—the period utility functions and the discount factor—are identified from functions that are potentially observable in the data. Potentially observable functions are those which can be consistently, nonparametrically estimated in a first step and include the conditional choice probability function, the laws of motion for the state variables, and certain other conditional expectations. Second, we describe how we estimate the model following the semiparametric plug-in procedure of BCNP, modified appropriately to account for features of our model.

Aguirregabiria (2005, 2010), Norets and Tang (2014), Arcidiacono and Miller (2015), Chou (2016), and Kalouptsidi et al. (2016) all discuss identification of counterfactual choice probabilities in dynamic discrete choice models such as ours. They emphasize that arbitrarily normalizing one of the choice-specific utility functions to zero across all states is not innocuous for analyzing counterfactuals. This is contrary to the common practice in applied work, a practice we avoid in this paper. In this section, we characterize a class of models in which semiparametric identification of the utility function is possible without such a normalization.

Arcidiacono and Miller (2015) consider identification in the case of short panel data, such as ours, where the sampling period ends before the model time horizon. They show that the counterfactual CCPs for temporary policy changes involving only changes to payoffs are identified even when the flow payoffs are not. They do not, however, consider identification of the payoff function itself without a normalization. We show that this is possible in cases with multiple terminating actions.

Chou (2016) demonstrates that normalizations affect counterfactual policy predictions and shows that no normalization is needed if there are variables that affect the state transition law but not the per period utilities. Our identification results effectively impose a different form of exclusion restriction, based on the presence of multiple terminating actions. One of the actions is repeated and may affect the available sequences of choices but not current payoffs directly.

Finally, we note that in models with continuous state variables alternative approaches are possible. For example Taber (2000), considered the case where the distribution of errors is instead left unspecified, under certain exclusion restrictions. In more recent work, Buchholz, Shum, and Xu (2016) build on Srisuma and Linton (2012)'s framework for continuous-state models to develop a single-index representation and a corresponding closed-form semiparametric estimator for dynamic binary choice models with linear utility functions and unspecified error distributions. Komarova et al. (2016) also consider models with linear payoffs, but with a focus on identification of the discount factor, payoff coefficients, and switching cost parameters when the distribution of errors is known.

A.1. Non-Identification in a Dynamic Binary Choice Model

To motivate our stated goal of avoiding an ad hoc location assumption or "normalization" on the utility function, we first consider a simple dynamic binary response model. It is known in the literature that Assumption 1 alone does not provide sufficient restrictions to identify the structural primitives of the model (Rust, 1994). Magnac and Thesmar (2002) show that to identify utility functions for each alternative, in addition to those in Assumption 1, restrictions on the discount rate and the utility of a reference choice are also needed. Before introducing the additional assumptions needed for our new identification result, after defining some notation we first revisit this non-identification result since it gives insights into the nature of the problem.

Let $v_t(s,a)$ denote the choice-specific value function for choice a,

$$v_t(s,a) = u(s,a) + \beta \operatorname{E}\left[V_{t+1}(s') \mid s, a\right].$$

As is well-known, the choice probabilities depend only on differences in the choice-specific value function at particular states. For example, in the logistic case the choice probability for $a = 1$ in state s is

$$\sigma_t(s,1) = \frac{\exp(v_t(s,1) - v_t(s,0))}{1 + \exp(v_t(s,1) - v_t(s,0))}.$$

We will make use of Lemma 1 of Arcidiacono and Miller (2011), which extends results of Hotz and Miller (1993) to establish that the ex-ante value function can be written in terms of choice probabilities and the choice-specific value function for an arbitrary reference choice a. When specialized to the case of type I extreme value errors, their result is that for any state s and choice a,

$$V_t(s) = v_t(s,a) - \log \sigma_t(s,a) + \gamma, \tag{8}$$

where γ is Euler's constant. Intuitively, this representation of the ex-ante value has three components: the value from the reference choice ($v_t(s,a)$), a non-negative adjustment term ($-\log \sigma_t(s,a)$) in case the reference choice is not optimal, and the mean of the type I extreme value distribution (γ). Suppose that $a = 1$ is a terminating action after which the agent receives no additional utility: $v_t(s,1) = u(s,1)$. Using termination as the reference choice, we can express the ex-ante function simply in terms of within-period quantities:

$$V_t(s) = u(s,1) - \log \sigma_t(s,1) + \gamma. \tag{9}$$

Substituting (9) at period $t+1$ into the definition of the choice-specific value function for the continuation choice $a = 0$ yields

$$v_t(s,0) = u(s,0) + \beta \, \mathrm{E}\left[u(s',1) - \log \sigma_{t+1}(s',1) \mid s, a = 0\right] + \beta\gamma.$$

Differencing this function across choices (since this difference appears in the choice probabilities) gives an expression involving three differences:

$$\begin{aligned} v_t(s,0) - v_t(s,1) = [u(s,0) - u(s,1)] + \beta \, \mathrm{E}[u(s',1) \mid s, a = 0] \\ - \beta \, \mathrm{E}[\log \sigma_{t+1}(s',1) - \gamma \mid s, a = 0] \end{aligned} \tag{10}$$

This representation highlights two important points. First, following Rust (1994), we can see that under the maintained assumptions the utility function is not identified, even when the error distribution is known. We can construct an observationally equivalent utility function $\tilde{u}$ that yields the same difference $v_t(\cdot,0) - v_t(\cdot,1)$ and therefore the same choice probabilities.

Second, it is clear from (10) that the transition probabilities play an important role here. If we were to assume—incorrectly—that $u(\cdot,1)$ is a constant function (e.g., equal to zero), then the second term is constant and the transition density does not interact with the utility function. If in reality the termination payoff varies with the state variables, then using the choice specific value function based on the incorrectly normalized utility function would yield incorrect welfare measures. We summarize these points in the following lemma. The proofs of this result and others are given in Appendix B.

Lemma 1. *Under Assumption 1, neither the utility function, u, nor the ex-ante value function, V_t, are identified.*

On a positive note, for simple counterfactuals where $\tilde{u}$ is an additive transformation of u, the counterfactual choice probabilities are known to be identified even if u itself is only identified up to differences (Aguirregabiria, 2010; Arcidiacono and Miller, 2015). Unfortunately, many interesting counterfactuals involve either affine or nonlinear transformations of u or changes in the state transition density, in which cases the counterfactual choice probabilities are not identified when u is only known up to differences (Kalouptsidi et al., 2016).

A.2. Identification of the Utility Function

Here, we introduce a simple model in which to consider identification with multiple terminating actions. Specifically, consider a three-choice model with one continuation action ($a = 0$) and two terminating actions ($a = 1$ and $a = 2$). The first terminating action ($a = 1$) results in immediate termination while the second ($a = 2$) must be chosen twice consecutively to result in termination. The arguments can be extended readily to models with more choices and with more complex terminating circumstances, such as our empirical model. This simpler framework contains the essential elements needed for our identification result and is motivated directly by the case of HECM households, which can continue, terminate immediately, or be forced to terminate by defaulting for multiple consecutive periods. An example of this structure from labor economics would be an employee who can either quit immediately (immediate termination) or be fired by failing to meet performance criteria for multiple periods in a row (repeated termination). Letting a_{-1} denote the choice in the previous period, the choice-specific value function can be expressed as follows:

$$v_t(s, a_{-1}, a) = \begin{cases} u(s,0) + \beta \, \mathrm{E}[V_{t+1}(s', 0) \mid s, a = 0] & a = 0, \\ u(s,1) & a = 1, \\ u(s,2) + \beta \, \mathrm{E}[V_{t+1}(s', 2) \mid s, a = 2] & a_{-1} \neq a = 2, \\ u(s,2) & a_{-1} = a = 2. \end{cases} \tag{11}$$

Remark. There is effectively a payoff exclusion restriction on a_{-1} in the representation above. Although the lagged choice a_{-1} can affect payoffs by limiting the available sequences of choices, it does not appear in the payoff function u directly. For example, having defaulted in the last period does not affect period utility u conditional on current state s and action a, but for a household already in default, choosing to default again will terminate the model and no future payoffs will be received. Furthermore, because agents are forward looking, they internalize the expected increase in the probability of termination through foreclosure in future periods. This is similar in spirit to the type of exclusion restrictions considered by Magnac and Thesmar (2002) and Abbring and Daljord (2018), but we also note that those papers assume the utility for one choice is either zero or a known function.

The ex-ante value function can be written recursively as

$$\begin{aligned}V_t(s,a_{-1}=0) = &\mathrm{E}\left[\max\left\{u(s,0)+\varepsilon(0)+\beta\,\mathrm{E}[V_{t+1}(s',0)\mid s,a=0],\right.\right.\\&u(s,1)+\varepsilon(1),\\&\left.\left.u(s,2)+\varepsilon(2)+\beta\,\mathrm{E}[V_{t+1}(s',2)\mid s,a=2]\right\}\mid s,a_{-1}=0\right],\\V_t(s,a_{-1}=2) = &\mathrm{E}\left[\max\left\{u(s,0)+\varepsilon(0)+\beta\,\mathrm{E}[V_{t+1}(s',0)\mid s,a=0],\right.\right.\\&u(s,1)+\varepsilon(1),\\&\left.\left.u(s,2)+\varepsilon(2)\right\}\mid s,a_{-1}=2\right].\end{aligned}$$

Using the choice-specific value functions in (11) and Assumption 1, the ex-ante value function can be written compactly as

$$V_t(s,a_{-1}) = \sum_{a\in\mathcal{A}}\int 1\{\delta_t(s,a_{-1},\varepsilon)=a\}[v_t(s,a_{-1},a)+\varepsilon(a)]F_\varepsilon(d\varepsilon).$$

From Theorem 1 of Arcidiacono and Miller (2011), $V_t(s,a_{-1})$ can be expressed as functions of period payoffs, conditional choice probabilities and state transition probabilities for an arbitrary sequence of choices. Only a finite number of period payoffs is needed in cases where there is a terminal choice such as the mortgage termination choice in our application, a renewal action such as the engine replacement choice in Rust (1987), or in general the distribution of future states following a specific sequence of choices does not depend on initial choices (Arcidiacono and Miller (2011)). Specialized to our application, with the presence of a terminal choice, $V_t(s,a_{-1})$ can be written as period payoff of immediate termination plus functions of the conditional choice probabilities. To see this, define the function

$$\begin{aligned}w(z_a,z_b) = \int &[z_a\,1\{z_a+\varepsilon(0)\geq\varepsilon(1),\,z_a+\varepsilon(0)\geq z_b+\varepsilon(2)\}\\&+z_b\,1\{z_b+\varepsilon(2)\geq\varepsilon(1),\,z_b+\varepsilon(2)\geq z_a+\varepsilon(0)\}]\,F_\varepsilon(d\varepsilon).\end{aligned}\tag{12}$$

With this definition, $V_t(s,a_{-1})$ can be written as

$$\begin{aligned}V_t(s,a_{-1}) =&u(s,1)+w\left(v_t(s,a_{-1},0)-u(s,1),\,v_t(s,a_{-1},2)-u(s,1)\right)\\&+\sum_{a\in\mathcal{A}}\int 1\{\delta_t(s,a_{-1},\varepsilon)=a\}\varepsilon(a)F_\varepsilon(d\varepsilon).\end{aligned}\tag{13}$$

The conditional choice probabilities in this setting are $\sigma_t(s,a_{-1},a)$. We will focus on the following six conditional probabilities: $\sigma_t(s,0,0)$, $\sigma_t(s,0,1)$, $\sigma_t(s,0,2)$, $\sigma_t(s,2,0)$, $\sigma_t(s,2,1)$

and $\sigma_t(s,2,2)$. They can be written in terms of the error distribution F_ε as

$$\sigma_t(s,0,0) = \int 1\{v_t(s,0,0)+\varepsilon(0) \geq u(s,1)+\varepsilon(1),$$
$$v_t(s,0,0)+\varepsilon(0) \geq v_t(s,0,2)+\varepsilon(2)\}\, F_\varepsilon(d\varepsilon), \tag{14}$$
$$\sigma_t(s,0,2) = \int 1\{v_t(s,0,2)+\varepsilon(2) \geq u(s,1)+\varepsilon(1),$$
$$v_t(s,0,2)+\varepsilon(2) \geq v_t(s,0,0)+\varepsilon(0)\}\, F_\varepsilon(d\varepsilon). \tag{15}$$
$$\sigma_t(s,2,0) = \int 1\{v_t(s,2,0)+\varepsilon(0) \geq u(s,1)+\varepsilon(1),$$
$$v_t(s,2,0)+\varepsilon(0) \geq u(s,2)+\varepsilon(2)\}\, F_\varepsilon(d\varepsilon), \tag{16}$$
$$\sigma_t(s,2,2) = \int 1\{u(s,2)+\varepsilon(2) \geq u(s,1)+\varepsilon(1),$$
$$u(s,2)+\varepsilon(2) \geq v_t(s,2,0)+\varepsilon(0)\}\, F_\varepsilon(d\varepsilon). \tag{17}$$

Note here that once the continuation action is taken ($a_{-1}=0$), given s the action in the last period no longer affects the choices going forward. As a result, $v_t(s,0,0) = v_t(s,2,0)$ in (16) and (17).

Equations (14), (15), (16), and (17) define a mapping Γ from payoff differences to choice probabilities:[19]

$$\Gamma : [v_t(s,a_{-1},0)-u(s,1),\ v_t(s,a_{-1},2)-u(s,1)] \mapsto [\sigma_t(s,a_{-1},0),\ \sigma_t(s,a_{-1},2)]\,. \tag{18}$$

Under the full support assumption (Assumption 1.c), Γ is invertible by Proposition 1 of Hotz and Miller (1993) and surjective by Norets and Takahashi (2013) and Lemma 1 of BCNP. Therefore, given any choice probabilities the payoff differences following continuation can be solved uniquely and we will denote the components of the inverse mapping simply as $\Gamma_1^{-1}(\sigma_t(s,0,\cdot)) = v_t(s,0,0) - u(s,1)$ and $\Gamma_2^{-1}(\sigma_t(s,0,\cdot)) = v_t(s,0,2) - u(s,1)$. Similarly, following repeated termination ($a_{-1}=2$) the differences are identified as $\Gamma_1^{-1}(\sigma_t(s,2,\cdot)) = v_t(s,2,0)-u(s,1)$ and $\Gamma_2^{-1}(\sigma_t(s,2,\cdot)) = u(s,2)-u(s,1)$. Hotz and Miller (1993) also prove that there exists a mapping such that $\psi^*(\sigma_t(s,a_{-1},\cdot)) = \sum_{a\in\mathcal{A}} \int 1\{\delta_t(s,a_{-1},\varepsilon)=a\}\varepsilon(a)F_\varepsilon(d\varepsilon)$. Therefore from (13), $V_t(s,a_{-1})$ can be written as

$$V_t(s,a_{-1}) = u(s,1) + \psi(\sigma_t(s,a_{-1},\cdot)), \tag{19}$$

where $\psi(\sigma_t(s,a_{-1},\cdot)) = w\left(\Gamma^{-1}(\sigma_t(s,a_{-1},\cdot))\right) + \psi^*(\sigma_t(s,a_{-1},\cdot))$, and a_{-1} is the previous

[19] Recall that there are three choices, but we note that only two payoff differences and two choice probabilities are relevant for the Γ mapping. The remaining difference $v_t(s,a_{-1},0) - v_t(s,a_{-1},2)$ is determined by the subtracting the two differences appearing as functional arguments, and the remaining choice probability is determined as $\sigma_t(s,a_{-1},1) = 1 - \sigma_t(s,a_{-1},0) - \sigma_t(s,a_{-1},2)$.

action which can be 0 (continuation) or 2 (default). This representation generalizes that of (8), for the special case of the type I extreme value distribution, using immediate termination ($a = 1$) as the reference action. Since the choice probabilities $\sigma_t(s, \cdot, \cdot)$ are identified, these ex-ante value functions are identified up to $u(s, 1)$. The remaining ex-ante value $V_t(s, 1)$ is zero since $a = 1$ is a terminal choice.

The presence of two terminating actions allows us to identify $u(s, 1)$ and therefore the full payoff function u. To show this, we make the following additional completeness assumption, which guarantees that there is sufficient variation in the state transition density, as the following theorem shows. Broadly, completeness is similar to a full rank condition for finite dimensional models,[20] and it has been used as an identifying assumption for nonparametric instrumental variable models (Newey and Powell, 2003; Blundell, Chen, and Kristensen, 2007; Darolles, Fan, Florens, and Renault, 2011; Chen, Chernozhukov, Lee, and Newey, 2014), however Canay, Santos, and Shaikh (2013) show that in some cases it is not testable.

Assumption 2 (Completeness)**.** The conditional distributions $f_{s'|s,a=2}$ is complete for s. In other words, for for all integrable functions h we have $\int h(s') f_{s'|s,a=2}(s')\, ds' = 0$ for all s if and only if $h = 0$.

We will maintain this high-level assumption for now, but immediately below in Lemma 2 we establish weaker alternative assumptions for common special cases. For example, in our application identification will follow from the parametric, linear specification for u without appealing to Assumption 2. To focus on identification of the utility function, we also assume that the discount factor β is known for now. Although this is a common practice in applied work, in Lemma 3 in below we give a condition under which β is separately identified and we appeal to this result in our application.

Theorem 1. *If Assumptions 1 and 2 hold and β is identified (e.g., it is known or identified by Lemma 3 below), then the utility function u is identified.*

We note that unlike BCNP, our identification result does not require that we observe the final decision period T. This "short panel" setting is common in empirical work and is the subject of a recent study by Arcidiacono and Miller (2015). However, in contrast to their findings for more general models, in our setting the period utility function and discount factor are identified without assuming the utility function is known for one choice.

We conclude our discussion of identification by considering sufficient conditions for the completeness required by Assumption 2 in some common special cases. The proofs appear in Appendix B.

[20]In the finite-dimensional setting, if a square matrix A has *full rank* then $Ax = 0$ implies $x = 0$. In an infinite-dimensional setting, if the distribution of Y is *complete* for X, then $\int g(y) f(y \mid x)\, dy = 0$ for all x implies $g(y) = 0$ for all y.

Lemma 2. *Suppose Assumption 1 holds and β is identified (e.g., it is known or identified by Lemma 3 below).*

a. Constant termination payoffs: *If the termination payoffs are unknown, but constant, then u is identified without additional assumptions.*

b. Parametric utility: *If for each choice a, $u(s, a) = u(s, a; \theta^a)$ with $\theta = (\theta^0, \theta^1, \theta^2)$, then u is identified if the following parametric identification conditions hold:*

i. $u(s, 0; \theta^0) = u(s, 0; \theta^0_0) = 0$ *for all s if and only if* $\theta^0 = \theta^0_0$.

ii. $\mathrm{E}\left[u(s', 1; \theta^1) - u(s', 1; \theta^1_0) \mid s, a = 2\right] = 0$ *for all s if and only if* $\theta^1 = \theta^1_0$.

iii. $u(s, 2; \theta^2) = u(s, 2; \theta^2_0) = 0$ *for all s if and only if* $\theta^2 = \theta^2_0$.

c. Linear utility: *If the payoffs have linear representations of the form $u(s, a) = u(s, a; \theta^a) = s^\top \theta^a$, where θ^a are choice-specific linear coefficients for each a, then u is identified if the covariance matrix $\mathrm{E}[s_t s_t^\top]$ and the conditional covariance matrix $\mathrm{E}[s_{t+1} s_{t+1}^\top \mid s_t, a_t = 2]$ have full rank.*

d. Finite state space: *Suppose that $s \in \mathcal{S}$ with $|\mathcal{S}| < \infty$. Then u is identified if the $|\mathcal{S}| \times |\mathcal{S}|$ choice-specific Markov transition matrix $\Pi_2 = [\Pr(s' \mid s, a = 2)]$ has full rank.*

Therefore, in each case there are weaker alternatives to the completeness assumption we invoke for the general nonparametric u case.

A.3. Identification of the Discount Factor β

As Chung et al. (2014) noted, in finite-horizon models the period utility function is identified by the terminal period leaving the discount factor to be identified by intertemporal variation in observed behavior. BCNP showed that the discount factor is identified when there is variation in the CCPs over time, which is natural in finite-horizon models. Under the same assumption, stated below, we verify that the discount factor β is identified in our model with multiple terminating actions. This does not require that the terminal period is observed or that the termination payoffs are known.

Assumption 3 (Nonstationary Choice Probabilities)**.** For some period t, s and $a_{-1} \in \{0, 2\}$, $\sigma_{t+1}(s, a_{-1}, \cdot) \neq \sigma_{t+2}(s, a_{-1}, \cdot)$.

Remark. The nonstationarity required by Assumption 3 is used only for identifying β, not u, as shown in the following Lemma 3. This assumption requires at least three periods of data to be available. In finite-horizon models, it is natural that optimal decisions depend on the number of time periods remaining; hence the choice probabilities are expected to be nonstationary. In stationary models (e.g., infinite-horizon models under commonly used assumptions), Assumption 3 will not hold and therefore β needs to be identified from other sources. Also note that β is a product of the factor of time discounting b and the survival probability $p(s)$. Because the survival probability is a known function of borrower demographics, identification of β implies identification of the time discounting factor b. Finally, we note that the identification result in Theorem 1 above for the utility function is conditional on β being identified or known.

Lemma 3. *If Assumptions 1 and 3 hold, β is identified.*

B. Proofs

Proof of Lemma 1

Suppose the true utility function is u where the period utility of termination $u(s,1) \neq 0$. First, following Rust (1994) we can find an observationally equivalent utility function $\tilde{u}$ that yields the same observable CCPs σ while satisfying an identifying restriction such as a "zero normalization". For each state s and choice a, define $\tilde{u}(s,1) = 0$ and $\tilde{u}(s,0) = u(s,0) - u(s,1) + \beta\,\mathrm{E}[u(s',1) \mid s, a=0]$. Then, by substituting u and $\tilde{u}$ into (10) above, we can verify that both utility functions yield the same differences in choice-specific value functions and hence the same observable CCPs. Therefore the utility function is not identified with Assumption 1 alone.

Next, using (8) from the Arcidiacono-Miller Lemma, with termination ($a=1$) as the reference choice, we can state the ex-ante value function as in (9). For the true utility function we have $V_t(s) = u(s,1) - \log \sigma_t(s,1) + \gamma$ and for the alternative utility function $\tilde{u}$ we have $\tilde{V}_t(s) = \tilde{u}(s,1) - \log \tilde{\sigma}_t(s,1) + \gamma$. But $\tilde{\sigma}_t = \sigma_t$ and so the value functions are only equal everywhere if $u = \tilde{u}$, which is the case when the utility function is identified.

Proof of Theorem 1

First, note that

$$\int \big(u(s',2) - u(s',1)\big)\, f_{s'|s,a=2}(s')\, ds' = \int \Gamma_2^{-1}(\sigma_{t+1}(s',2,\cdot)) f_{s'|s,a=2}(s')\, ds',$$

where the equality follows from (18). The right hand side is a functional of the observed data. By the completeness assumption on the conditional distribution of $f_{s'|s,a=2}$ (Assumption 2), there is a unique solution for $u(s,2) - u(s,1)$ and hence it is identified.

Next, subtracting $u(s,1)$ from both sides of $v_t(s,0,2)$ and substituting for $V_{t+1}(s',2)$

we have

$$
\begin{aligned}
v_t(s,0,2) - u(s,1) &= u(s,2) - u(s,1) + \beta \, \mathrm{E}\left[V_{t+1}(s',2) \mid s, a=2\right] \\
&= u(s,2) - u(s,1) + \beta \, \mathrm{E}\left[u(s',1) \mid s, a=2\right] \\
&\quad + \beta \, \mathrm{E}\left[\psi(\sigma_{t+1}(s',2,\cdot)) \mid s, a=2\right].
\end{aligned}
$$

As the mapping Γ in (18) is invertible, $v_t(s,0,2) - u(s,1)$ is also identified from the data. Substituting and solving to obtain an expression for the remaining unknown, $u(s',1)$, yields

$$
\begin{aligned}
&\mathrm{E}\left[u(s',1) \mid s, a=2\right] \\
&= \beta^{-1}\left[v_t(s,0,2) - u(s,1)\right] - \beta^{-1}\left[u(s,2) - u(s,1)\right] - \mathrm{E}\left[\psi(\sigma_{t+1}(s',2,\cdot)) \mid s, a=2\right]. \qquad (20)
\end{aligned}
$$

The period payoff $u(s,1)$ is then identified under Assumption 2. Once $u(s,1)$ and the difference $u(s,2) - u(s,1)$ are identified, so is $u(s,2)$. Finally, subtracting $u(s,1)$ from both sides of the expression for the remaining choice-specific payoff $v_t(s,a_{-1},0)$ gives

$$
\begin{aligned}
v_t(s,a_{-1},0) - u(s,1) &= u(s,0) - u(s,1) + \beta \, \mathrm{E}\left[V_{t+1}(s',0) \mid s, a=0\right] \\
&= u(s,0) - u(s,1) + \beta \, \mathrm{E}\left[u(s',1) \mid s, a=0\right] \\
&\quad + \beta \, \mathrm{E}\left[\psi(\sigma_{t+1}(s',0,\cdot)) \mid s, a=0\right].
\end{aligned}
$$

where $a_{-1} \in \{0,2\}$. The left-hand side is identified, and so are all quantities on the right-hand side of the second equality except for $u(s,0)$. Solving for $u(s,0)$ yields

$$
u(s,0) = u(s,1) + \left[v_t(s,a_{-1},0) - u(s,1)\right] - \beta \, \mathrm{E}\left[u(s',1) + \psi(\sigma_{t+1}(s',0,\cdot)) \mid s, a=0\right]. \qquad (21)
$$

Therefore, $u(s,a)$ is identified for all choices $a = 0,1,2$.

Proof of Lemma 2

We consider each case in turn below.

a. *Constant termination payoffs:* Suppose that the termination payoffs are constant: $u(\cdot,1) = c_1$ and $u(\cdot,2) = c_2$. Then the difference is identified immediately as $c_2 - c_1 = u(\cdot,1) - u(\cdot,2) = \Gamma_2^{-1}(\sigma_t(\cdot,2,\cdot))$, where the second equality follows from the proof of Theorem 1. Next, c_1 is separately identified by (20) since $\mathrm{E}\left[u(s',1) \mid s, a=2\right] = c_1$, and then c_2 is also identified. Finally, $u(s,0)$ is identified from (21) as before, in the proof of Theorem 1.

b. *Parametric utility:* Define $\Delta^{1,2}u(s;\theta^{1,2}) \equiv u(s,2;\theta^2) - u(s,1;\theta^1)$, where $\theta^{1,2} = (\theta^1,\theta^2)$. Recall that $\Delta^{1,2}u(\cdot;\theta^{1,2}) = \Gamma_2^{-1}(\sigma_t(\cdot,2,\cdot))$ is identified. Therefore, the set $\Theta^{1,2}$ of parameter values $\theta^{1,2}$ which yield the above identified difference is also identified. This set may not be a singleton, but we note that the vector of true parameters (θ_0^1, θ_0^2) is an element. Importantly, for all elements $\theta^{1,2} \in \Theta^{1,2}$, including the true parameters, the right-hand-side of (20) is constant because it depends only on the difference $\Delta^{1,2}u(\cdot;\theta^{1,2})$ and other identified quantities. The parameter θ_0^1, which appears on the left-hand side of (20) through $u(s,1;\theta_0^1)$ is then separately identified under part ii of the maintained assumption. This, in turn, identifies the function $u(\cdot,2;\theta_0^2)$, and under part iii of the assumption, the parameter θ_0^2 is identified. Finally, as before, the function $u(\cdot,0;\theta_0^0)$ is identified from (21). Under part i of the assumption, the parameter θ_0^0 is identified.

c. *Linear utility:* Suppose that $u(s,a) = s^\top \theta^a$ for all s and a. Define $\Delta\theta^{1,2} \equiv \theta^2 - \theta^1$. Then the condition identifying the utility difference $u(s,2) - u(s,1)$ becomes

$$u(s_{t+1},2) - u(s_{t+1},1) = s_{t+1}^\top \Delta\theta^{1,2} = \Gamma_2^{-1}(\sigma_{t+1}(s_{t+1}, a_t = 2, \cdot)).$$

Premultiplying both sides by s_{t+1}, taking expectations, and using the rank assumption on the conditional covariance matrix allows us to solve for $\Delta\theta^{1,2}$:

$$\Delta\theta^{1,2} = \mathrm{E}[s_{t+1}s_{t+1}^\top \mid s_t, a_t = 2]^{-1}\,\mathrm{E}[s_{t+1}\Gamma_2^{-1}(\sigma_{t+1}(s_{t+1},2,\cdot)) \mid s_t, a_t = 2].$$

Turning to (20), we can premultiply by s_{t+1}, substitute to obtain expressions in terms of conditional choice probabilities, and solve to find

$$\theta^1 = \mathrm{E}[s_{t+1}s_{t+1}^\top \mid s_t, a_t = 2]^{-1}\Big\{\beta^{-1}\,\mathrm{E}[s_{t+1}\Gamma_2^{-1}(\sigma_t(s_t,0,\cdot)) - s_{t+1}s_t^\top\Delta\theta^{1,2} \mid s_t] \\ - \mathrm{E}\left[s_t\psi\left(\sigma_{t+1}(s_{t+1},2,\cdot)\right) \mid s_t, a_t = 2\right]\Big\}$$

This separately identifies θ^1 and therefore θ^2. In line with previous arguments, θ^0 is then identified from (21) under the full rank assumption on $\mathrm{E}[s_t s_t^\top]$.

d. *Finite state space:* In this case, there are a finite number of unknowns represented by choice-specific vectors u_0, u_1, and u_2, each of length $|\mathcal{S}|$. Similarly, let $\sigma_{t,a_{-1},a}$, $w_{t,a_{-1}}$, and $\gamma_{t,a_{-1},j}$ denote, respectively, denote the vectors of values $\sigma_t(s,a_{-1},a)$, $\psi(\sigma_t(s,a_{-1},\cdot))$, and $\Gamma_j^{-1}(\sigma_t(s,a_{-1},\cdot))$ stacked across s. First, the differences in termination payoffs are identified as $u_2 - u_1 = \gamma_{t,2,2}$. Then, stacking (20) yields another matrix equation for u_1: $\Pi_2 u_1 = \beta^{-1}(\gamma_{t,0,2} - \gamma_{t,2,2}) - \Pi_2 w_{t,2}$. Since Π_2 has full rank, this equation identifies u_1, and hence u_2 separately. As in previous cases, u_0 is then identified directly from the vectorized counterpart of (21).

Proof of Lemma 3

Given Assumption 3, the conditional choice probabilities are nonstationary and $\sigma_{t+1}(s, a_{-1}, a) \neq \sigma_{t+2}(s, a_{-1}, a)$ for some s, a_{-1} and a. Because

$$\begin{aligned}\sigma_{t+1}(s, a_{-1}, a) = \int 1 \big(& u(s, a) + \varepsilon(a) + \beta \, \mathrm{E}[V_{t+2}(s', a) \mid s, a_{-1}, a] \\ & \geq u(s, \tilde{a}) + \varepsilon(\tilde{a}) + \beta \, \mathrm{E}[V_{t+2}(s', \tilde{a}) \mid s, a_{-1}, \tilde{a}] \; \forall \tilde{a} \in \mathcal{A}\big) \, F_{\varepsilon}(d\varepsilon),\end{aligned}$$

in order for $\sigma_{t+1}(s, a_{-1}, a) \neq \sigma_{t+2}(s, a_{-1}, a)$, the conditional expectation of future values must be time varying for some $\tilde{a} \in \mathcal{A}$,

$$\text{(22)} \quad \mathrm{E}[V_{t+2}(s', \tilde{a}) \mid s, a_{-1}, \tilde{a}] \neq \mathrm{E}[V_{t+3}(s', \tilde{a}) \mid s, a_{-1}, \tilde{a}].$$

Fix $a = \tilde{a}$, and consider the expressions for $v_t(s, a_{-1}, \tilde{a})$ in two adjacent time periods:

$$\begin{aligned} v_{t+1}(s, a_{-1}, \tilde{a}) &= u(s, \tilde{a}) + \beta \, \mathrm{E}\left[V_{t+2}(s', \tilde{a}) \mid s, a_{-1}, \tilde{a}\right] \\ v_{t+2}(s, a_{-1}, \tilde{a}) &= u(s, \tilde{a}) + \beta \, \mathrm{E}\left[V_{t+3}(s', \tilde{a}) \mid s, a_{-1}, \tilde{a}\right]. \end{aligned}$$

Subtracting these equations and solving for β, we find

$$\text{(23)} \quad \beta = \frac{(v_{t+1}(s, a_{-1}, \tilde{a}) - u(s, 1)) - (v_{t+2}(s, a_{-1}, \tilde{a}) - u(s, 1))}{\mathrm{E}[V_{t+2}(s', \tilde{a}) - V_{t+3}(s', \tilde{a}) \mid s, a_{-1}, \tilde{a}]}.$$

The denominator is nonzero given (22). All terms on the right hand side of (23) can be identified from the CCPs (see (18) and (19)) and the transition density of the state variables, and therefore β is identified.

C. Out-of-Sample HECM Policy Function Fit

Previously in Table 5 we reported the within-sample fit of the HECM policy function estimates. Here, to provide an additional way to evaluate the fit of the first-stage model we randomly divide the sample into two halves. The first half is used as a training sample to estimate the model parameters and then the model predictions are compared with the actual choices for observations in the second half of the sample. These results are reported in Table 12. In addition to good within-sample fit, these results indicate that the model also has good out-of-sample prediction power.

Table 12. Out-of-Sample Prediction of Reduced Form HECM Policy Function Estimates

	Termination		Refinance		Default	
Prediction Sample	Prediction	Data	Prediction	Data	Prediction	Data
			Unconditional			
All	1.35%	1.51%	0.41%	0.32%	4.09%	4.12%
			By HECM Type			
Fixed Rate	1.19%	1.35%	0.29%	0.34%	4.87%	4.70%
Adjustable Rate	1.59%	1.73%	0.57%	0.29%	2.99%	3.32%
			By Loan Age			
1	0.56%	0.79%	0.31%	0.26%	0.41%	0.43%
2	1.76%	1.84%	0.48%	0.45%	3.19%	3.23%
3	1.81%	1.99%	0.39%	0.35%	5.84%	6.02%
4	1.29%	1.42%	0.46%	0.18%	7.95%	7.83%
			By Credit Score			
Q1	1.09%	0.96%	0.44%	0.31%	11.82%	12.08%
Q2	1.27%	1.51%	0.33%	0.35%	3.18%	3.35%
Q3	1.37%	1.80%	0.39%	0.41%	0.75%	0.59%
Q4	1.68%	1.78%	0.47%	0.21%	0.50%	0.37%
			By Net Equity			
Q1	1.10%	1.02%	0.19%	0.17%	8.85%	8.74%
Q2	1.34%	1.60%	0.37%	0.26%	4.52%	4.64%
Q3	1.40%	1.73%	0.48%	0.35%	2.23%	2.36%
Q4	1.57%	1.69%	0.59%	0.50%	0.74%	0.76%
			By Available HECM Credit			
Q1	1.35%	1.53%	0.42%	0.37%	5.89%	5.95%
Q2	1.46%	1.35%	0.44%	0.22%	0.49%	0.39%
Q3	1.25%	1.48%	0.32%	0.13%	0.17%	0.22%
Q4	1.34%	1.53%	0.39%	0.26%	0.01%	0.04%

This table shows the out-of-sample prediction fit of the policy function estimates, both unconditionally and conditional on some explanatory variables. Q1–Q4 denote the first through fourth quartiles of the stated variables. The data is randomly divided into a training sample and a prediction sample of equal sizes. The policy function is estimated using the training sample.