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

May 15, 2017

Abstract. We carry out an empirical analysis of the Home Equity Conversion Mortgage (HECM) program using a unique and detailed dataset on the behavior of HECM borrowers from 2007–2014 to semiparametrically estimate a structural, dynamic discrete choice model of borrower behavior. Our estimator is based on a new identification result for models with multiple terminating actions where we show that the utility function and discount factor are identified without the need to impose ad hoc identifying restrictions (i.e., assuming that the payoff for one choice is zero). Such restrictions can lead to incorrect counterfactual choice probabilities and welfare calculations. Our estimates, which are not based on such an assumption, provide insights about the factors that influence HECM refinance, default, and termination decisions. We use the results to quantify the trade-offs involved for proposed program modifications through a series of counterfactual simulations. We find that income and credit requirements would indeed be effective in reducing undesirable HECM outcomes, at the expense of excluding some borrowers, and we quantify the relative welfare losses due to restricting access to the program. We also investigate how shocks to housing prices affect HECM outcomes and household welfare.

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. We also thank Peter Arcidiacono, Thomas Davidoff, Juan Carlos Escanciano, Sukjin Han, Ashley Langer, Lance Lochner, David Rivers, Daniel Sacks, Kamila Sommer, Mauricio Varela, Tiemen Woutersen, and Ruli Xiao, attendees of the 2016 Calgary Empirical Microeconomics conference and the 2017 AREUEA-ASSA Annual Meeting, as well as seminar participants at Indiana University, Ohio State University (Glenn College of Public Affairs), the University of Arizona, the University of Maryland, and the University of Texas at Austin for useful comments.

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 have lower incomes, lower property tax to income ratios, or 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 remaining HECM credit, 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. Black households and households with initial withdrawals in excess of 80% of the borrowing limit value the program more, while single male homeowners value the program less.

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 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 NPV 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 home 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 level, and thus to examine how these rates vary with individual borrower characteristics.

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 the functional form of utility is known for one choice. In particular, our model has two distinct, observable terminating actions which allow us to identify the period payoff functions for all choices. Our approach generalizes to other single-agent models with multiple terminating actions under conditions we formalize in Section 3. We estimate the model using a multi-step plug-in semiparametric approach inspired by that of BCNP. 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 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 (CCP) 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), and Mastrobuoni and Rivers (2016) underscores the importance of estimating time preferences. We show that the approach of BCNP for identifying the discount factor, based on nonstationarity of the CCPs, is valid in our model as well, and we estimate the discount factor as part of our analysis.

Our results contribute to a growing collection of known sufficient conditions for identification of models and/or counterfactuals without making a functional form assumption for one choice. For example, BCNP show that for non-stationary models such as ours identification is possible if the final decision period is observed (i.e., the panel is "short" and ends before the final model time period). Arcidiacono and Miller (2015) show that the counterfactual CCPs for temporary policy changes involving only changes to payoffs are identified in the short panel case even when the flow payoffs themselves are not. Chou (2016) demonstrated that no utility normalization is needed if there is an "exclusion restriction": a variable that affects the law of motion of the state variables but not the utility function.

Following negative identification results by Rust (1994) and Magnac and Thesmar (2002), the identification literature moved towards considering semiparametric identification of the utility function under a parametric assumption on the choice-specific errors. We follow the same strategy, but 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.

In contrast to previous work, 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, although such requirements have since been enacted and are among the counterfactual policy changes we consider in this paper.[1] 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. During our sample period, borrowers could choose between fixed- (FRM) and adjustable rate (ARM) HECMs. Fixed-rate HECM borrowers received the entire principal limit in an up-front lump sum payment.[2] 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. HECMs are non-recourse loans, meaning that borrowers will 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.[3]

[1]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).

[2]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).

[3]Like all dynamic discrete choice models, in reality a household's HECM decisions are embedded in a

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. Fixed-rate HECMs require borrowers to withdraw all credit upon loan closing, while borrowers with adjustable-rate HECMs may structure their HECMs as lines of credit and can have access to the credit lines later. 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.[4] 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. 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.

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.

[4] 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).

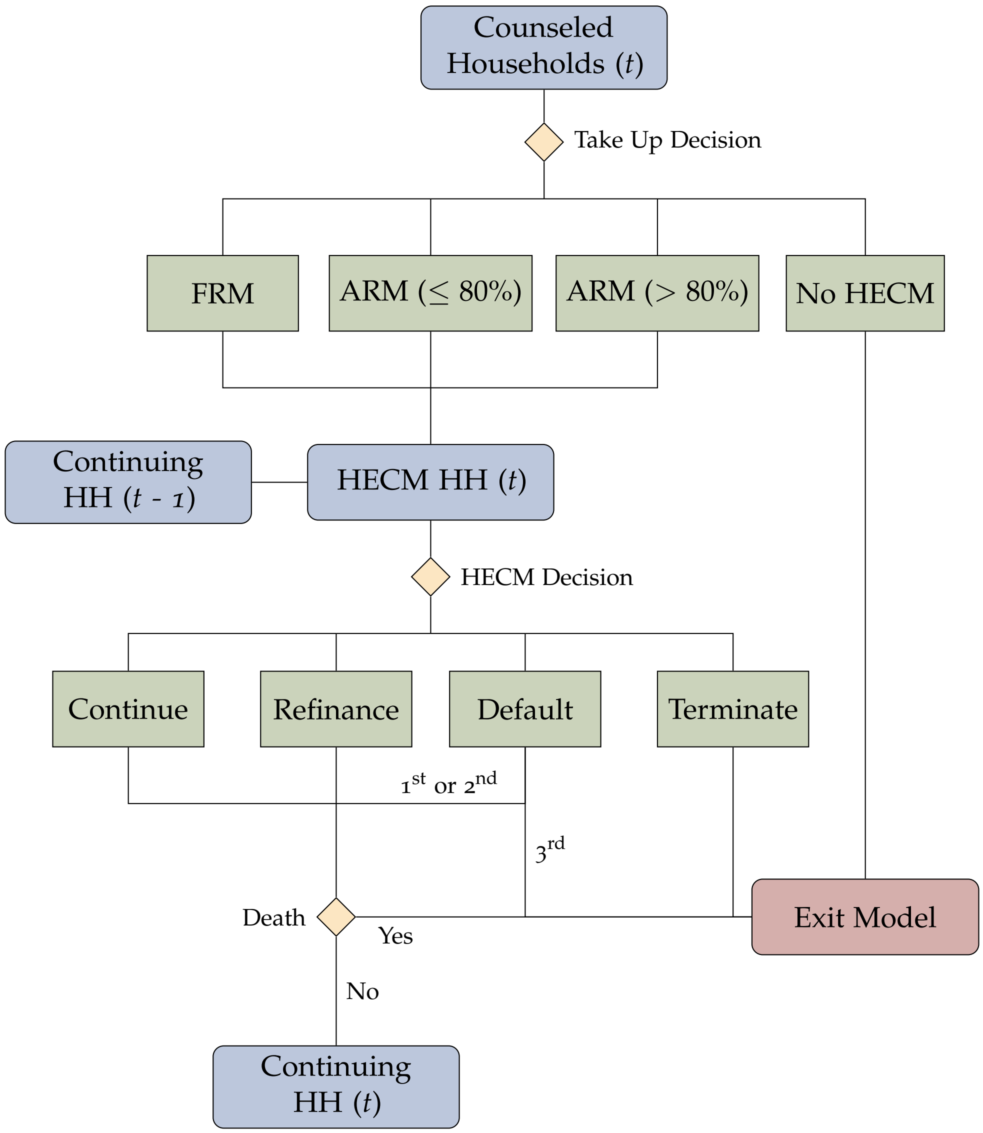

FIGURE 1. Borrower Decision Flow Chart

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 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 the subsequent period. 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). 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.[5] 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 origination, or the HECM becomes "due and payable" when the homeowner fails to live in the home for a consecutive twelve month period due to physical or mental illness (Mortgagee Letter 2015-10). 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.[6] We denote the the survival probability for household i, conditional on age and sex in period t, by $p(s_{it})$.[7]

[5] In 2015, HUD formalized certain loss mitigation policies, such as repayment plans, with "goal of keeping HECM borrowers in their homes whenever possible" (Mortgagee Letter 2015-10; Mortgagee Letter 2015-11). Recently HUD extended the deadline to submit "due and payable" requests through April 2016 (Mortgagee Letter 2016-01), meaning that lenders may wait even longer to take action against delinquent borrowers. This furthers the previous extension granted in October 2015 (Mortgagee Letter 2015-26).

[6] For loans with two borrowers, we use the joint probability that both borrowers die in the same year.

[7] 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.

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}. We also make the standard assumptions that 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 period utility received by a household in state s_{it} that chooses action $a_{it} \in \mathcal{A}_t$ is

$$U_t(s_{it}, a_{it}, \varepsilon_{it}) = u_t(s_{it}, a_{it}) + \varepsilon_{it}(a_{it}), \tag{1}$$

where $u_t(s_{it}, a_{it})$ is the deterministic or mean utility component and $\varepsilon_{it}(a_{it})$ is an idiosyncratic, choice-specific shock.

Households in our model are forward-looking and discount future utility using a discount factor β. 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 function or *policy function* $\sigma_t(s_{it}, a_{it})$.

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,

$$V_t^{\delta}(s_{it}) = \mathrm{E}^{\delta}\left[\sum_{\tau=t}^{T} \beta^{\tau-t} U_t(s_{it}, \delta_t(s_{it}, \varepsilon_{it}), \varepsilon_{it}) I_{it} \,\middle|\, s_{it}\right]. \tag{2}$$

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 $u_t(s_{it}, a_{it})$ 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 a household terminates by choosing $a_{it} = \mathrm{T}$, they receive the mean period utility $u_t(s_{it}, \mathrm{T})$ and the idiosyncratic shock $\varepsilon_{it}(\mathrm{T})$, but no additional utility is received in the future. Hence, $u_t(s_{it}, \mathrm{T})$ 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 ($a_{it} = a_{i,t-1} = a_{i,t-2} = \mathrm{D}$), they receive the mean utility for defaulting $u_t(s_{it}, \mathrm{D})$, the idiosyncratic shock, and because the HECM will be terminated, the termination payoff $u_t(s_{it}, \mathrm{T})$.

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

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

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.[8]

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. In other words, the terminal utility ($u_t(s_{it}, \mathrm{T})$) can be interpreted as present discounted utility for the future values after termination.

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 payoffs 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:

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

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).

[8] Implicitly, $v_t(s_{it}, a_{it})$ as defined here also depends on $a_{i,t-1}$ and $a_{i,t-2}$. For now, we omit this dependence for simplicity: in the general notation, we can simply include the lagged actions in the state vector s_{it}. For the purposes of identification, we formalize this dependence below in Section 3.

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 which leads to choice probabilities of the following form:

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

We state the formal assumptions below, one of which is a conditional independence assumption that limits the persistence of the unobservables (Rust, 1987). In practice, we 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} in the HECM choices in (1). 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.

[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.

3. Semiparametric Identification and Estimation

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.

In our application, we are limited by our data to observing households in some cases for only three years after the initial HECM take-up decision. Although a non-trivial number of borrowers do terminate their HECMs within this timeframe, in most cases this does not cover the terminal period. The results of Arcidiacono and Miller (2015), Chou (2016), and Kalouptsidi et al. (2016), among others, show that using the true utility levels is of utmost importance to avoid biasing the counterfactual policy outcomes of interest.

Before turning to identification, we first formally state the assumptions we have maintained so far, which have been standard assumptions invoked by Rust (1987) 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:

$$f(s_{t+1}, \varepsilon_{t+1} \mid s_t, \varepsilon_t, a_t) = f(s_{t+1} \mid s_t, a_t) f(\varepsilon_{t+1} \mid s_{t+1}).$$

b. The payoffs are additively separable in the choice-specific errors and the deterministic component is a time invariant function of s_t and a_t:

$$U_t(s_t, a_t, \varepsilon_t) = u(s_t, a_t) + \varepsilon_t(a_t).$$

c. The choice-specific errors are independent from s_t and 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). 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 exables 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.

In the next section, we show that in models such as ours, with multiple terminating actions, 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.

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

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

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,

$$V_t(s) = v_t(s,a) - \log\sigma_t(s,a) + \gamma, \tag{5}$$

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:

$$V_t(s) = u(s,1) - \log\sigma_t(s,1) + \gamma. \tag{6}$$

Substituting (6) 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:

$$(7)\qquad v_t(s,0) - v_t(s,1) = [u(s,0) - u(s,1)] + \beta\,\mathrm{E}[u(s',1) \mid s, a=0] \\ - \beta\,\mathrm{E}[\log \sigma_{t+1}(s',1) - \gamma \mid s, a=0]$$

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 (7) 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 proof is reserved for Appendix A.

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).

3.2. Identification of the Utility Function

For simplicity, we consider identification in 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:

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

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 u per se, 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.

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

$$\sigma_t(s,0,0) = \int 1\left\{v_t(s,0,0) + \varepsilon(0) \geq u(s,1) + \varepsilon(1),\ v_t(s,0,0) + \varepsilon(0) \geq v_t(s,0,2) + \varepsilon(2)\right\} F_\varepsilon(d\varepsilon), \tag{8}$$

$$\sigma_t(s,0,2) = \int 1\left\{v_t(s,0,2) + \varepsilon(2) \geq u(s,1) + \varepsilon(1),\ v_t(s,0,2) + \varepsilon(2) \geq v_t(s,0,0) + \varepsilon(0)\right\} F_\varepsilon(d\varepsilon). \tag{9}$$

$$\sigma_t(s,2,0) = \int 1\left\{v_t(s,2,0) + \varepsilon(0) \geq u(s,1) + \varepsilon(1),\ v_t(s,2,0) + \varepsilon(0) \geq u(s,2) + \varepsilon(2)\right\} F_\varepsilon(d\varepsilon), \tag{10}$$

$$\sigma_t(s,2,2) = \int 1\left\{u(s,2) + \varepsilon(2) \geq u(s,1) + \varepsilon(1),\ u(s,2) + \varepsilon(2) \geq v_t(s,2,0) + \varepsilon(0)\right\} F_\varepsilon(d\varepsilon). \tag{11}$$

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 (10) and (11).

Equations (8), (9), (10), and (11) define a mapping Γ from payoff differences to choice probabilities:[10]

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

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 $\Gamma_2^{-1}(\sigma_t(s,2,\cdot)) = u(s,2) - u(s,1)$.

Now, the ex-ante value function in (2) can be written recursively as

$$\begin{aligned} V_t(s,a_{-1}) &= \mathrm{E}\left[\max_a \left\{u(s,a) + \varepsilon(a) + \beta\,\mathrm{E}[V_{t+1}(s',a) \mid s,a]\right\} \,\middle|\, s,a_{-1}\right] \\ &= \sum_{a\in\mathcal{A}} \int 1\{\delta_t(s,a_{-1},\varepsilon) = a\}[v_t(s,a_{-1},a) + \varepsilon(a)]F_\varepsilon(d\varepsilon), \end{aligned}$$

Next, we define the function

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

With this definition, as guaranteed by the Arcidiacono-Miller lemma we arrive at the following CCP representation for an arbitrary error distribution:

$$\begin{aligned} V_t(s,0) &= u(s,1) + w\left(v_t(s,0,0) - u(s,1),\, v_t(s,0,2) - u(s,1)\right), \\ &= u(s,1) + w\left(\Gamma^{-1}(\sigma_t(s,0,\cdot))\right). \end{aligned}$$

This representation generalizes that of (5), for the special case of the type I extreme value distribution, using immediate termination ($a = 1$) as the reference action and continuation ($a_{-1} = 0$) here as the previous action. Similarly for repeated termination,

$$V_t(s,2) - u(s,1) = w\left(\Gamma^{-1}(\sigma_t(s,2,\cdot))\right).$$

Since the choice probabilities $\sigma_t(s,\cdot,\cdot)$ are identified, these ex-ante value functions are

[10] 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 $\sigma_t(s,a_{-1},1) = 1 - \sigma_t(s,a_{-1},0) - \sigma_t(s,a_{-1},2)$.

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,[11] 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 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 Lemma 3 below), then the utility function u is identified.*

Proof. First, note that $u(s,2) - u(s,1) = \Gamma_2^{-1}(\sigma_t(s,2,\cdot))$ 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

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

As shown previously, $v_t(s,0,2) - u(s,1)$, $u(s,2) - u(s,1)$ and $v_t(s,2,0) - u(s,1)$ are identified from the data. Substituting and solving to obtain an expression for the remaining

[11] 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.

unknown, $u(s',1)$, yields

$$\text{(12)}\quad \mathrm{E}\left[u(s',1) \mid s, a=2\right] = \beta^{-1}\left[v_t(s,0,2) - u(s,1)\right] - \beta^{-1}\left[u(s,2) - u(s,1)\right] \\ - \mathrm{E}\left[w(v_{t+1}(s',2,0) - u(s',1),\, u(s',2) - u(s',1)) \mid s, a=2\right]$$

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 $v_t(s,0,0)$ gives

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

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

$$\text{(13)}\quad u(s,0) = u(s,1) + \left[v_t(s,0,0) - u(s,1)\right] \\ - \beta\,\mathrm{E}\left[u(s',1) + w(v_{t+1}(s',0,0) - u(s,1),\, v_{t+1}(s',0,2) - u(s,1)) \mid s, a=0\right].$$

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

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 A.

Lemma 2. *Suppose Assumption 1 holds and β is identified (e.g., it is known or identified 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. $\mathrm{E}\left[u(s',1;\theta^1) - u(s',1;\theta_0^1) \mid s, a=2\right] = 0$ *for all* s *if and only if* $\theta^1 = \theta_0^1$.

iii. $u(s,2;\theta^2) = u(s,2;\theta_0^2) = 0$ *for all* s *if and only if* $\theta^2 = \theta_0^2$.

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* $\mathrm{E}[s_t s_t^\top]$ *and the conditional autocovariance matrix* $\mathrm{E}[s_{t+1} s_t^\top \mid a_t = 2]$ *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.

3.3. 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 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, $\Pr[\sigma_{t+2}(s,2,\cdot) \neq \sigma_{t+1}(s,2,\cdot)] > 0$.

Remark. The nonstationarity required by Assumption 3 is used only for identifying β, not u. It requires that at least three periods of data are available. Our identification result in Theorem 1 above for u was conditional on β being identified. Assumption 3 is sufficient for that, but our argument for identification of the utility function also extends to stationary models (e.g., infinite-horizon models under commonly used assumptions) in which β is otherwise identified or known.

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

Proof. Consider the expressions for $v_t(s,0,2)$ in adjacent time periods:

$$v_t(s,0,2) = u(s,2) + \beta\,\mathrm{E}\left[w(v_{t+1}(s',2,0) - u(s',1), u(s',2) - u(s',1)) + u(s',1) \mid s, a=2\right]$$
$$v_{t+1}(s,0,2) = u(s,2) + \beta\,\mathrm{E}\left[w(v_{t+2}(s',2,0) - u(s',1), u(s',2) - u(s',1)) + u(s',1) \mid s, a=2\right].$$

Subtracting these equations and solving for β, we find

$$\begin{aligned}\beta = \big\{&\mathrm{E}[w(v_{t+1}(s',2,0) - u(s',1),\, u(s',2) - u(s',1)) \mid s, a = 2] \\ &- \mathrm{E}[w(v_{t+2}(s',2,0) - u(s',1),\, u(s',2) - u(s',1)) \mid s, a = 2]\big\}^{-1} \\ &\times [(v_t(s,0,2) - u(s,1)) - (v_{t+1}(s,0,2) - u(s,1))].\end{aligned}$$

Using the Γ mapping, we can restate the equality as follows:

$$\beta = \frac{\Gamma_2^{-1}(\sigma_t(s,0,\cdot)) - \Gamma_2^{-1}(\sigma_{t+1}(s,0,\cdot))}{\mathrm{E}\left[w\left(\Gamma^{-1}(\sigma_{t+1}(s',2,\cdot))\right) \mid s, a = 2\right] - \mathrm{E}\left[w\left(\Gamma^{-1}(\sigma_{t+2}(s',2,\cdot))\right) \mid s, a = 2\right]}.$$

Assumption 3 guarantees that the denominator is nonzero. All terms on the right hand side can be identified from the CCPs and the transition density of the state variables, and therefore β is identified. ■

3.4. Semiparametric Estimation

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 $\mathcal{A}_t = \{\mathrm{C}, \mathrm{R}, \mathrm{D}, \mathrm{T}\}$ and type I extreme value errors, we use a series representation of the log odds ratio

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

for choices $a \in \mathcal{A}_t$ relative to termination ($a = \mathrm{T}$). The functions q_l are basis functions and $r_l(t,a)$ are the coefficients which will be estimated. In practice we approximate the infinite sum using a finite but large number of basis functions and coefficients, denoted by L. Let $\hat{\sigma}_t(s, a_{-1}, a)$ denote the estimated choice probabilities, obtained as

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

for $a \in \mathcal{A}_t \setminus \{\mathrm{T}\}$ 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, which is identified directly from the data through the relationship

$$\text{(14)} \quad \mathrm{E}[V_{t+1}(s', a) \mid s, a] = -\,\mathrm{E}[\log \sigma_{t+1}(s', a, \mathrm{T}) \mid s, a] + \mathrm{E}[u(s', \mathrm{T}) \mid s, a] + \gamma.$$

The first conditional expectation on the right hand side, which is a function of current s and a, can be estimated nonparametrically using data on the period-ahead choices a_{t+1}. Meanwhile, γ is a known constant. However, because we do not assume the termination utility function is the zero function there is an additional term on the right hand side of (14) relative to BCNP. In their case, the second term on the right hand side of (14) 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 the law of motion of the state variables.

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 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 the discount factor β. The estimating equations are the log odds ratios for the choices $a \in \mathcal{A}_t$:

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

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

$$\log \frac{\hat{\sigma}_t(s, a_{-1}, a)}{\hat{\sigma}_t(s, a_{-1}, \mathrm{T})} = u(s, a; \theta) - u(s, \mathrm{T}; \theta) - \beta \hat{\mathrm{E}}[\log \sigma_{t+1}(s', a, \mathrm{T}) \mid s, a] + \beta \hat{\mathrm{E}}[u(s', \mathrm{T}; \theta) \mid s, a] + \beta\gamma.$$

This allows us to estimate the structural parameters θ and β by nonlinear least squares. The parameters in the take-up model can be similarly estimated.

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. This form further simplifies the problem, yielding estimating equations of the following form for $a \in \{C, R, D\}$:

$$\log \frac{\hat{\sigma}_t(s, a_{-1}, a)}{\hat{\sigma}_t(s, a_{-1}, T)} = s'\theta_a - s'\theta_T - \beta \hat{E}\left[\log \hat{\sigma}_{t+1}(s', a, T) \mid s, a\right] + \beta \hat{E}\left[s' \mid s, a\right]' \theta_T + \beta\gamma.$$

4. Data

4.1. State Variables

Our data is drawn from a sample of 21,564 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 report attributes 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 (61 percent of counseled households in our sample), counseling data is linked to HUD loan data using confidential personal identifiers. HUD HECM loan data includes details 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, others 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 reports annually which allows us to follow the evolution of the FICO 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 and the evolution of the housing price index,[12] 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 withdrawals (80% or more).

Three of the financial variables deserve special attention: net equity, HECM credit, and excess credit. 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 the insurance component of the HECM program.

HECM Balance The current HECM balance is calculated based on the amounts a borrower withdraws over time. This balance grows 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 choose their initial withdrawal amount and may make subsequent withdraws, as needed or on an installment basis.

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,

[12]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).

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 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 aspect 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

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

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	0.046	0.109	0.054	0.083	0.081	0.272
Hispanic	0.090	0.075	0.123	0.082	0.084	0.277
Black	0.058	0.177	0.270	0.132	0.138	0.344
Single male	0.197	0.211	0.194	0.147	0.150	0.357
Single female	0.402	0.381	0.465	0.386	0.390	0.488
Monthly Income†	0.263	0.246	0.201	0.244	0.242	0.165
Property tax/income	0.105	0.114	0.102	0.091	0.091	0.093
Non-housing assets†	6.361	2.181	2.563	4.383	4.314	16.994
Fixed rate HECM	0.529	0.524	0.704	0.598	0.602	0.489
Initial withdrawal $>$ 80%	0.606	0.714	0.904	0.717	0.724	0.447
Time-Varying Variables						
Age	75.784	72.245	73.217	73.094	73.134	7.545
FICO	717.957	701.381	594.682	706.950	701.602	93.369
Available revolving credit†	2.341	3.149	0.343	2.244	2.156	3.005
Revolving & installment debt†	1.081	1.201	0.966	1.268	1.250	2.272
Net equity†	13.259	16.296	4.237	10.391	10.148	12.976
Negative net equity†	0.025	0.000	0.221	0.057	0.064	0.609
Excess credit†	0.062	0.000	0.187	0.076	0.081	0.527
Tax & insurance balance†	0.005	0.003	0.178	0.000	0.009	0.097
Available HECM credit†	4.000	3.779	0.230	3.979	3.795	6.619
Household-year observations	624	147	2,250	43,078	44,697	46,099

† Monetary variables are measured in units of $10,000.

Table 3. Summary Statistics for the Take-Up Sample

	FRM	ARM (≤ 80%)	ARM (> 80%)	No HECM	All Households	
	Mean	Mean	Mean	Mean	Mean	SD
Pre-HECM Variables						
Age	70.948	74.127	72.264	70.758	71.503	7.968
Young borrower	0.188	0.114	0.130	0.177	0.167	0.373
Hispanic	0.065	0.068	0.180	0.100	0.088	0.284
Black	0.148	0.079	0.204	0.229	0.174	0.379
Single male	0.159	0.150	0.167	0.188	0.170	0.376
Single female	0.383	0.446	0.398	0.374	0.391	0.488
Monthly Income†	0.247	0.221	0.213	0.232	0.234	0.168
Property tax/income	0.078	0.118	0.106	0.085	0.090	0.094
Non-housing assets†	4.529	4.403	2.156	4.492	4.322	17.285
FICO	684.774	726.452	671.308	659.174	680.191	101.522
Available revolving credit†	2.098	3.225	2.596	1.803	2.202	3.588
Revolving & installment debt†	1.721	1.287	1.655	1.595	1.590	2.942
Change in housing price index	-0.055	-0.065	-0.084	-0.064	-0.062	0.054
Average interest rate (ARM)	5.281	5.318	5.387	5.286	5.297	0.183
Average interest rate (FRM)	5.282	4.402	2.315	5.081	4.837	1.496
Initial HECM Variables						
Initial withdrawal > 80%	1	0	1	–	–	–
Net equity†	14.914	23.331	15.844	–	–	–
Negative net equity†	0.047	0.022	0.014	–	–	–
Excess credit†	0	0.001	0.005	–	–	–
Tax & insurance balance†	0	0	0	–	–	–
Available HECM credit†	7.785	13.409	8.421	–	–	–
Household observations	6,871	3,419	1,441	8,415	20,146	20,146

† Monetary variables are measured in units of $10,000.

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 $2,420. Approximately 60% of observations are for FRMs and 72% 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 FICO scores, little available credit (HECM and other credit), lower net equity, higher excess credit, and have T&I default balances. The average FICO score is 702, however, for borrowers who default it is 594. For refinance observations, households tend to be younger, have higher net equity, more available revolving credit, higher income, and higher property tax/income ratios.

Similarly, Table 3 reports the summary statistics for our take-up sample. These values are averages over household-year observations, as in Table 2. Over half of the counseled borrowers do ultimately 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 FICO scores and those that choose large-draw ARMs have the lowest FICO 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.

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. In this model, we included all of the state variables from the structural model as well as HECM loan age (years since origination) and interactions of the loan age with selected state variables. The loan age is included because the model has a finite time horizon and decision rules may vary as the loan age changes. The specification we use was selected by minimizing the Akaike information criterion (AIC).[13]

Table 4 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 borrower state variables. Overall the predicted choice probabilities capture the patterns in the data reasonably well. Note that the data is censored, because choices are observed only for households who survive, i.e. who are not forced to exit due to death, etc, and this may contribute to the discrepancies between the observed and predicted choice frequencies.

In addition, we also use a sieve multinomial logit model to estimate the conditional choice probabilities for HECM take-up using the counselee 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 the 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 5 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.

5.2. Structural Utility Function Estimates

The total value for a household consists of a choice-specific period payoff, a continuation value conditional on the state variables and choice taken this period, and an i.i.d. type 1 extreme value error. 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. Table 6 contains estimates of the per-period, choice-specific utility coefficients along with 95% bootstrap confidence intervals.

The higher the termination value relative to the payoff from other choices, the more likely that the HECM will be terminated. Borrowers that receive more value from ter-

[13] We considered many alternative specifications, some of which included cubic splines of the state variables with 3 to 5 equally spaced knots and/or additional interaction terms. The final specification was selected by choosing the one with the minimum AIC value.

Table 4. In-Sample Fit of Reduced Form HECM Policy Function Estimates

	Termination		Refinance		Default	
Sample	Prediction	Data	Prediction	Data	Prediction	Data
Unconditional						
All	1.36%	1.35%	0.32%	0.32%	4.88%	4.88%
By HECM Type						
Fixed Rate	1.19%	1.19%	0.28%	0.28%	5.69%	5.65%
Adjustable Rate	1.63%	1.61%	0.38%	0.38%	3.61%	3.68%
By Loan Age						
1	0.90%	0.66%	0.38%	0.29%	0.71%	0.48%
2	1.52%	1.75%	0.36%	0.47%	3.49%	3.57%
3	1.72%	1.84%	0.32%	0.35%	6.42%	6.78%
4	1.40%	1.33%	0.23%	0.24%	8.93%	8.86%
5	1.17%	0.82%	0.18%	0.00%	8.69%	8.34%
6	0.87%	0.92%	0.13%	0.00%	9.00%	7.69%
By Credit Score						
Q1	1.01%	0.95%	0.32%	0.34%	13.85%	14.02%
Q2	1.30%	1.34%	0.30%	0.29%	4.04%	4.09%
Q3	1.48%	1.53%	0.34%	0.37%	0.98%	0.88%
Q4	1.66%	1.60%	0.31%	0.28%	0.49%	0.39%
By Net Equity						
Q1	1.14%	1.00%	0.21%	0.14%	10.06%	10.26%
Q2	1.26%	1.27%	0.30%	0.18%	5.39%	5.36%
Q3	1.39%	1.58%	0.35%	0.43%	2.97%	2.80%
Q4	1.64%	1.57%	0.41%	0.53%	1.09%	1.09%
By Available HECM Credit						
Q1	1.32%	1.38%	0.31%	0.35%	7.57%	7.69%
Q2	1.42%	1.23%	0.38%	0.31%	1.19%	0.51%
Q3	1.43%	1.23%	0.35%	0.14%	0.32%	0.34%
Q4	1.43%	1.48%	0.29%	0.36%	0.03%	0.12%

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 5. 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	37.84%	37.84%	15.17%	15.17%	2.83%	2.83%
			By Year of Counseling			
2009	36.10%	36.10%	14.57%	14.57%	6.32%	6.32%
2010	37.53%	37.53%	17.30%	17.30%	2.77%	2.77%
2011	38.61%	38.61%	13.04%	13.04%	2.00%	2.00%
			By Age			
Q1	39.62%	38.87%	9.82%	9.84%	2.43%	2.37%
Q2	39.43%	40.71%	12.36%	12.00%	2.62%	2.77%
Q3	38.58%	38.40%	16.30%	17.08%	2.89%	2.72%
Q4	33.13%	32.87%	23.17%	22.61%	3.46%	3.58%
			By Income			
Q1	34.44%	33.71%	15.54%	15.00%	2.64%	2.56%
Q2	36.77%	36.52%	16.32%	16.45%	2.79%	2.58%
Q3	38.68%	39.11%	15.09%	15.39%	2.83%	2.85%
Q4	41.45%	42.03%	13.69%	13.81%	3.07%	3.34%
			By Credit Score			
Q1	34.43%	33.44%	6.18%	6.81%	3.03%	2.82%
Q2	38.47%	39.55%	10.99%	9.92%	2.87%	3.21%
Q3	39.87%	41.23%	18.11%	17.59%	2.90%	2.87%
Q4	38.58%	37.16%	25.49%	26.47%	2.52%	2.42%
			By Net Equity			
Q1	40.35%	38.78%	8.93%	3.35%	2.57%	2.09%
Q2	40.31%	42.82%	12.35%	11.19%	2.59%	3.37%
Q3	38.32%	39.61%	15.52%	19.02%	2.86%	2.88%
Q4	32.37%	30.13%	23.88%	27.11%	3.30%	2.98%

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 6. Coefficient Estimates for Per-Period Payoffs

	Continue		Refinance		Default		Terminate	
Constant	3.890	(-3.079, 20.458)	-1.188	(-9.609, 15.757)	5.281	(-2.777, 20.980)	13.932	(-0.494, 43.062)
Hispanic	0.551	(-0.146, 2.577)	-1.416	(-4.021, 0.605)	0.780	(-0.020, 2.347)	3.898	(1.241, 11.407)
Black	0.562	(-0.102, 1.931)	0.922	(-0.460, 2.412)	0.812	(0.005, 2.176)	2.415	(-0.091, 9.506)
Single male	-0.286	(-1.446, 0.149)	0.292	(-1.657, 1.702)	-0.224	(-1.485, 0.296)	-1.865	(-7.924, 0.769)
Single female	-0.346	(-1.845, 0.114)	-0.563	(-2.218, 0.715)	-0.413	(-2.145, 0.075)	-2.240	(-9.272, 0.349)
Income†	0.430	(-0.962, 4.094)	-0.940	(-5.183, 2.626)	0.727	(-0.557, 6.004)	1.007	(-7.404, 3.988)
Property tax/income	0.736	(-0.796, 8.447)	1.624	(-6.022, 11.513)	1.437	(-2.961, 8.805)	6.579	(0.147, 22.647)
Non-housing assets†	-0.003	(-0.045, 0.001)	-0.018	(-0.079, 0.109)	-0.006	(-0.038, 0.025)	-0.015	(-0.087, 0.009)
Fixed rate HECM	-0.033	(-0.831, 0.720)	1.139	(-0.470, 2.731)	-0.047	(-0.984, 0.475)	-0.105	(-3.390, 1.561)
First year credit utilization $>$ 80%	-0.023	(-0.898, 2.643)	-0.405	(-2.732, 2.468)	0.526	(-0.576, 3.149)	2.320	(-1.201, 10.857)
FICO	-0.003	(-0.020, 0.002)	-0.005	(-0.024, 0.002)	-0.011	(-0.032, -0.002)	-0.016	(-0.052, -0.000)
Available revolving credit†	0.023	(-0.399, 1.102)	0.090	(-0.467, 1.022)	-0.018	(-0.521, 0.906)	0.007	(-0.759, 1.458)
Revolving & installment debt†	-0.362	(-1.119, -0.019)	-0.444	(-1.306, -0.046)	-0.465	(-1.151, -0.135)	-0.771	(-2.433, -0.051)
Available HECM credit†	-0.090	(-0.388, -0.007)	-0.241	(-0.477, 0.008)	-0.409	(-0.675, -0.218)	-0.116	(-1.090, -0.033)
Net equity†	-0.004	(-0.075, 0.111)	0.074	(-0.018, 0.178)	-0.007	(-0.134, 0.053)	-0.037	(-0.102, 0.183)
Negative net equity†	0.313	(-0.011, 0.669)			0.163	(-0.973, 0.642)	-0.017	(-0.337, 0.218)
Excess credit†	0.210	(-0.072, 1.835)			-0.191	(-1.105, 1.274)	0.199	(-0.045, 3.283)
Discount factor	0.829	(0.230, 0.999)						

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

mination are Hispanic households and households with higher property tax to income ratios, lower credit scores, or lower revolving and installment debt. For ARM households, the termination values are higher with larger HECM credit utilization (i.e., less available HECM credit). Because the HECM credit limit does not change with a decline in house prices, ARM households are insured against house price declines to the extent of their HECM credit limit, and the insurance value is greater the more the home price drops below the HECM credit limit. The excess credit variable is significant at the 10% level in the per period termination payoff, suggesting that households may strategically terminate their HECM loans. At the same time, in our counterfactuals we do find that households derive more utility from the HECM program when housing prices are falling, suggesting that the insurance embedded in an HECM is valued by households. Black households and households with lower credit scores, less revolving and installment debt, and higher HECM utilization (less available HECM credit) receive higher per-period value from default, which means that if the continuation value is fixed, these households are more likely to default this period.

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 $\rho_{\text{NE}}\text{NE}_{it} + 1\{\text{NE}_{it} < 0\}\rho_{\text{NNE}}\,|\text{NE}_{it}|$.

5.3. Ex-Ante Value Function Estimates

We define the normalized ex-ante value function as $\overline{V}_t(s_{it}) = V_t(s_{it}) - u_t(s_{it}, \text{T})$, which is the expected discounted present value over and above the state-specific termination payoff. This represents the value households place on the HECM program relative to the outside option of terminating the loan. This is a nonlinear function, but to summarize how this value varies across households of different types, we report in Table 7 the results of a linear regression of $\overline{V}_t(s_{it})$ on state variables. This allows us to examine how households' valuations for remaining in the HECM program vary with household and loan characteristics and economic conditions. The higher the normalized ex-ante value, the more likely that the household will keep their HECMs. At 5% significance level, households value HECMs more if they have high initial credit utilization, have more available HECM credit, or are black. Households with high property tax to income ratios or high incomes value HECMs less, as do single male homeowners. The value is also higher when the net equity is lower (especially negative) and following a recent house price decline in the borrower's MSA.

Table 7. Regression of Normalized Ex-Ante Values on Borrower Characteristics

Dependent Variable: $\overline{V}_t(s_{it})$	Coeff.	95% CI
Young borrower	0.393	(0.011, 0.842)
Hispanic	-0.221	(-0.497, 0.078)
Black	0.763	(0.456, 1.105)
Single male	-0.392	(-0.593, -0.138)
Single female	-0.136	(-0.348, 0.075)
Fixed rate HECM	0.276	(-0.054, 0.565)
First year credit utilization $> 80\%$	0.326	(0.037, 0.649)
Defaulted in T&I last year	0.588	(-0.079, 1.345)
In default for two years	0.187	(-1.132, 2.980)
Age	0.375	(0.286, 0.467)
Age^2	-0.002	(-0.003, -0.002)
Income†	-0.839	(-1.311, -0.357)
Property tax/income	-0.976	(-1.850, -0.126)
FICO /100	-0.033	(-0.001, 0.000)
Available revolving credit†	0.001	(-0.025, 0.030)
Available HECM credit†	0.044	(0.026, 0.063)
Net equity†	-0.012	(-0.017, -0.006)
Negative net equity†	0.363	(0.080, 0.841)
Excess credit†	-0.077	(-0.263, 0.172)
Revolving & installment debt†	0.025	(-0.031, 0.080)
Non-housing assets†	-0.001	(-0.003, 0.002)
HPI change	-4.395	(-6.364, -2.320)
HPI change, 1 year lag	0.005	(-2.466, 2.141)
HPI change, 2 year lag	2.244	(0.546, 3.902)
Average interest rate (ARM)	0.811	(0.580, 1.046)
Average interest rate (FRM)	-0.020	(-0.162, 0.097)
Loan age		
2	0.145	(0.049, 0.236)
3	0.295	(0.115, 0.476)
4	0.439	(0.170, 0.707)
5	0.574	(0.217, 0.925)
6	0.705	(0.273, 1.169)
Constant	-12.444	(-16.279, -8.740)

The reported coefficients are for a linear regression of $\overline{V}_t(s_{it})$ 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 ($1,000$ replications).
† Monetary variables are reported in units of $10,000.

5.4. Counterfactual Simulations and Welfare Implications

Our counterfactual simulations have two objectives. In the first set of 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 property tax & insurance defaults (Mortgagee Letter 2015-06). 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. In the second part, we examine how borrowers' behavior and welfare will vary with improvements or downturns in housing market conditions as modeled by a one time unexpected change to house prices during the first year of the HECM.

We first simulated the effects of imposing borrower eligibility requirements on FICO scores and income. The results are summarized in Table 8. The first column indicates the levels of the cutoff values for the credit requirement and initial income requirement, respectively. The next three columns report the 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 and fraction of households with negative equity declined considerably while refinance and termination rates were largely unaffected. Surprisingly, both policies also reduce 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). 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 8 and the percentage change in this welfare in the eighth column. The final two columns measure the reduction in HECM volume (in number of households and the percentage change in households) due to these participation constraints. With a more stringent initial credit or income requirement, more households with relatively low credit scores or income are ineligible for HECMs, and the average borrower welfare as measured by the ex-ante value drops.

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 and negative net equity, we can compare the implications of a FICO requirement of 490—a very low threshold that excludes only 300 households in our sample—with those of an income requirement at between 1–1.25 times the Federal poverty level. The baseline default rate before the restrictions are imposed is 4.60%. The income requirement decreases this only slightly to 4.16–3.97% yet it would exclude 8–16% of borrowers in our sample, thus reducing total ex ante value from 11,551 to 10,675–9,810. In contrast, the credit requirement reduces the default rate slightly more, to 4.10%, and it does so by excluding fewer borrowers, only 3%. Furthermore, the drop in ex ante value is also less, staying at 11,250, so the welfare cost is lower. On the other hand, the income requirement is better in terms of reducing negative net equity.

Next, we simulate changes in house prices. In the counterfactual, HECM borrowers observe a one time change in their home values one year after HECM closing. Specifically, we simulate percentage changes in household home values and local housing price indices. Accordingly, we also adjust household net equity, excess credit, and the relevant HPI lags. Crucially, households' expectations for house prices remain the same in the counterfactual, as the change in housing prices is unexpected, and after this one time change, housing prices are assumed to follow the actual path as observed in the data. The home values in the counterfactual are within the empirical support of the home values in the data. As a result, households will not change their decision rules in the counterfactual, and the decision rules estimated using the actual data can be used to estimate borrowers' behavior and welfare under the alternative housing price scenarios.[14] Similar strategies are used by BCNP in their counterfactual simulations.

The results of a counterfactual decline in housing prices, net of the value of the home and the associated decrease in assets,[15] are summarized in Table 9. When housing prices fall by 8%, the rates of termination and refinance (second and third columns) fall and the rate of default (fourth column) increases slightly (and hence, the rate of continuation increases). The welfare of HECM households, as measured by the sum of the ex-ante values of all households (seventh column), actually *increases* by 2% when housing prices fall by 8%. This potentially surprising result is due to many factors, as we now explain. One factor is the direct change in housing prices. As we saw in Section 5.3, households who live in areas that have experienced recent house price declines tend to value the HECM program more on average. So do households with less net equity (especially negative net equity), and so when prices fall both the fraction of households with negative net equity (fifth column) and the household average dollar amount of that negative net equity (sixth column) increase. When house prices decrease households also experience an increase in excess credit, which is related to the insurance feature of HECM loans, but this does not seem to significantly affect either the period utility or ex-ante values.

[14]Our model is not a general equilibrium model, and therefore it cannot account for all possible effects of changing housing prices, such as the cost of alternative housing.

[15]As mentioned before, we focus only on household utility related to the HECM. Changes in housing prices for most seniors are essentially capital gains or losses, but the change in utility related to the HECM is independent and may even move in a different direction, as we see in our simulations.

Table 8. Counterfactual Imposition of Borrower Eligibility Requirements

	Termination	Refinance	Default	Negative Equity		Ex-Ante Value		HECM Volume	
	%	%	%	% HH	Total \$1M	Total	%Δ	% HH	%Δ
Initial credit requirement									
None	1.38	0.33	4.60	3.00	-244	53,131	–	11,551	–
490	1.41	0.33	4.10	2.39	-229	51,366	-3	11,250	-3
520	1.42	0.33	3.78	2.42	-222	49,543	-7	10,876	-6
550	1.44	0.33	3.32	2.41	-210	46,804	-12	10,304	-11
580	1.46	0.33	2.85	2.43	-188	44,088	-17	9,740	-16
610	1.48	0.33	2.42	2.41	-169	40,940	-23	9,087	-21
Initial income requirement									
None	1.38	0.33	4.60	3.00	-244	53,131	–	11,551	–
1 × FPL	1.41	0.33	4.16	2.20	-216	48,739	-8	10,675	-8
1.25 × FPL	1.42	0.33	3.97	2.08	-201	44,708	-16	9,810	-15
1.5 × FPL	1.41	0.32	3.81	1.93	-178	39,380	-26	8,650	-25
1.75 × FPL	1.42	0.32	3.68	1.81	-152	34,275	-35	7,546	-35
2 × FPL	1.44	0.33	3.57	1.76	-135	29,697	-44	6,551	-43

Initial income requirement is measured in terms of the Federal Poverty Level (FPL). 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.

Table 9. Counterfactual Simulations of Alternative House Price Scenarios

	Termination	Refinance	Default	Negative Equity		Ex-Ante Value	
	%	%	%	% HH	Total $1M	Total	%Δ
House price scenarios							
-8% 1st year change	1.27	0.33	4.72	6.09	-265	54,217	2
-4% 1st year change	1.32	0.32	4.66	4.35	-256	53,650	1
Baseline	1.38	0.33	4.60	3.00	-244	53,131	0
4% 1st year change	1.45	0.39	4.54	2.15	-235	52,643	-1
8% 1st year change	1.54	0.49	4.50	1.46	-232	52,188	-2

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.

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 is 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 key HECM decisions, including refinance, default, and termination. We show how household welfare is influenced by various factors and illustrate the welfare cost of policies that restrict program eligibility, with the aim of reducing defaults and adverse terminations.

A. Proofs

Proof of Lemma 1

Suppose the true utility function is u. First, following Rust (1994) we can find an observationally equivalent utility function $\tilde{u}$ that yields the same observable CCPs σ while still satisfying an identifying restriction such as a "zero normalization". For each state s and choice a, define $\tilde{u}(s,1) = 0$ and $\tilde{u}(s,0) = u(s,0) - u(s,1) + \beta \operatorname{E}[u(s',1) \mid s, a = 0]$. Then, by substituting u and $\tilde{u}$ into (7) above, we can verify that both utility functions yield the same differences in choice-specific value functions and hence the same observable CCPs.

Next, using (5) from the Arcidiacono-Miller Lemma, with termination ($a = 1$) as the reference choice, we can state the ex-ante value function as in (6). 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 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 (12) since $\operatorname{E}[u(s',1) \mid s, a = 2] = c_1$, and then c_2 is also identified. Finally, $u(s,0)$ is identified from (13) 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 (12) 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 (12) 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 (13). Under part i of the assumption, the parameter θ_0^0 is identified.

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

$$s_t^\top(\theta^2 - \theta^1) = \Gamma_2^{-1}(\sigma_t(s_t, 2, \cdot))$$

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

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

Turning to (12), we can premultiply by s_t, substitute to obtain expressions in terms of conditional choice probabilities, and solve to find

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

This separately identifies θ^1 and therefore θ^2. In line with previous arguments, θ^0 is then identified from (13) under the full rank assumption.

d. *Finite state space:* In this case, there are a finite number of payoffs 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 $\sigma_t(s, a_{-1}, a)$, $w\left(\Gamma^{-1}(\sigma_t(s, a_{-1}, \cdot))\right)$, 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 (12) 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 (13).

References

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

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., 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. (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.

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.

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.

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.

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 2015-26.

U.S. Department of Housing and Urban Development. Mortgagee letter 2016-01.

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.