THE UNIVERSITY OF CHICAGO

ESSAYS IN HOUSEHOLD AND HOUSING FINANCE

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
JOHN WENDELL HEILBRON

CHICAGO, ILLINOIS
JUNE 2022

Lovingly dedicated to the memory of Joseph Kelly Heilbron,

who was never cowed by the fear of striking out.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to thank my advisors, Amir Sufi, Elisabeth Kempf, Scott Nelson, and Anthony Zhang, for their guidance and feedback during my studies. I would also like to thank Pascal Noel, Constantine Yannelis, Stefan Nagel, Eric Zwick, and John Heaton for their thoughtful discussions of my work. This work has benefitted from the comments and suggestions of my peers in Amir Sufi's PhD Working Group, the Finance Brownbag, and the Third Year Seminar. A special thanks to my Booth Finance cohort-mates: Agustin Hurtado, Shohini Kundu, Kelly Posenau, and Blair Vorsatz.

I am grateful for my family, without whom I cannot imagine having started or finished graduate school - my dad for my love of math, my brothers for unflagging encouragement, and my mom for a certain "stick-to-it-ness". I am also indebted to Josh Morris-Levenson, classmate, companion, and confidant, for many discerning dinnertime conversations, and Rumur Dowling, the (however improbable) increasingly apparent love of my life, for continuing to take my calls.

ABSTRACT

This thesis consists of three essays in household and housing finance. "Credit Constraints at Home Purchase and Bias in Hedonic Amenity Valuations" argues that estimates of amenity values may be biased because households are credit constrained at the time of home purchase. I propose and implement a way to correct the bias in the estimates. "Mortgage Lending Limits and Housing Demand: Evidence from Bunching in FHA Borrowing" adapts the bunching literature to measure the loan-to-value elasticity of housing demand. I implement this estimator in FHA data and argue that this approach better distinguishes the role of credit constraints from the role of beliefs about housing returns. "Housing Wealth Management at Retirement" documents predictable refinancing behavior at the social security claim threshold. I argue that this behavior is driven by a demand for liquid assets in retirement and note that households may store excess wealth in housing relative to a rational benchmark.

CHAPTER 1

INTRODUCTION

In the wake of the Great Financial Crisis of 2007-2008, academic researchers have scrutinized the global financial system to better understand its origins and consequences. In particular, they have devoted considerable attention to the U.S. mortgage and housing markets, where so much of the trouble began. This sprawling agenda, ongoing now for some fifteen years, has appreciably changed our understanding of housing finance in a number of respects.

First and foremost, this literature has underscored the close relationship between lending conditions in the mortgage market and demand for housing. It has investigated and emphasized the importance of borrowing constraints and exuberant housing return expectations, features of household decision-making that deviate from benchmark models. And it has examined the institutional particularities of the U.S. mortgage loan market, the lending incentives inscribed by the GSEs, and their consequences for investor exposure to household credit risk.

This dissertation, therefore, enters on a field of study that has been recently and thoroughly excavated. Contributing to such a literature requires adopting a different lens and agenda. And so, while I draw on insights from the past decade of research, particularly those listed above, I address myself to questions that are more microeconomic in scope than those pursued by scholars of the GFC. In the essays that follow, I focus squarely on household decision-making and I present a detailed portrait of their financial relationships with that most peculiar of assets, housing. Specifically, I consider:

[I] How do households choose a neighborhood to live in?

[II] What are the relative influence of limited credit and expected returns on demand for housing?

[III] How do households use their homes as a store of value for retirement savings?

In a sense, then, these essays represent a thematic return to housing scholarship that predated the GFC, concerned as it was with topics like consumption smoothing and residential sorting. I examine households as they solve precisely these problems, but I do so with an eye to sensibilities that the GFC has made all but impossible to elide: the relevance of financing decisions, the limitations of household rationality and access to credit, and the importance of institutional strictures in mortgage lending.

In the lead essay, "Credit Constraints at Home Purchase and Bias in Hedonic Amenity Valuations," I argue that credit constraints at the time of home purchase introduce bias in hedonic estimates of school quality and that the degree of bias can be captured by households' willingness-to-pay for credit at the time of home purchase. In the absence of frictions, a household would be willing-to-pay, in net present value terms, \$1 for an additional \$1 of credit. I develop a way to measure this willingness-to-pay by exploiting household mortgage menus inherited from the pricing grids offered by the GSEs in the secondary mortgage market. I find that households are willing-to-pay as much as \$1.50 in net present value terms for an additional \$1 of credit and that estimates of school quality may be downward biased by as much as 50pp.

In the second essay, "Mortgage Lending Limits and Housing Demand: Evidence from Bunching in FHA Borrowing," I adapt the econometric bunching framework to measure the loan-to-value elasticity of housing demand among borrowers of FHA loans from 2018-9. In particular, I use information on loan-to-value and debt-to-income limits to construct a kinked choice set of house prices and down-payments and measure bunching at the kink. I find that the loan-to-value elasticity of intensive housing demand is statistically significant but economically small, at 14-25bp. In the absence of frictions, the elasticity with respect to this financing constraint would be zero; thus, the exercise provides a direct test of credit constraints at home purchase. Moreover, I argue that my approach to estimation does a better job than the literature of disentangling the effects of credit constraints from the effects of beliefs.

In the final essay, "Housing Wealth Management at Retirement," I instrument retirement with programmatic Social Security eligibility thresholds and find that retirement makes a household ~12pp more likely to issue any new mortgage debt and ~3pp more likely to extract equity from a home within the following two years. I find evidence that retirement-induced refis adds to liquid savings balances, unlike refis at other stages of the life-cycle, suggesting this is precautionary behavior. Because retirement is predictable and because of the associated transaction costs, I suggest these retirement refis may be evidence that households save excessive funds in housing wealth relative to a rational benchmark.

These research questions are eclectic and so, by necessity, each chapter is self-contained. Still, the chapters hang together because they are all preoccupied by a common set of substantive and methodological concerns. For one, between them, they highlight several distinctive features of housing as an asset class. Moreover, each uses detailed institutional context and granular micro-data to present findings as striking visuals which, where possible, lend themselves to structural interpretation. And finally, each takes care to compare observed household behavior to a frictionless benchmark.

In what follows, I distill and preview the way these themes provide continuity across chapters. For economy, I refer to Chapters (2), (3), and (4) as "Amenities," "Bunching," and "Retirement," respectively.

1.1 Housing as an asset class

Most famously, perhaps, housing is distinctive for being both a consumption and investment good. This poses a variety of interpretive challenges for the researcher. To what extent does the price of a home reflect demand for embedded housing services or beliefs about return expectations? Does a sale or down-sizing reflect portfolio re-balancing, a changing market outlook, or changing preferences of amenities? Each essay accommodates its argument to the challenge of this dual function, though in somewhat different ways.

In "Amenities," I follow the hedonic literature in supposing that, by comparing otherwise identical houses that differ only along a single attribute of interest, such as school quality, I net out the influence of investment prospects on house prices. At the heart of "Bunching" is the insight that bunching designs do not involve an information treatment. This allows me to recover the loan-to-value elasticity of housing demand, while shutting down a channel in which demand changes due to beliefs about housing returns, i.e. its investment value. In "Retirement," I strengthen my argument that households have a demand for liquid assets at the retirement threshold by appealing not just to down-sizing but also to equity extractions. Unlike in the event of down-sizing, in the case of equity extractions, housing services have not changed.

If the tacks they take to overcome this interpretive challenge are substitutes, the relative emphasis of each chapter on housing's different functions is complementary. To begin with, where "Amenities" focuses more on the consumption value of housing, "Bunching" pays more explicit attention to its value as a speculative investment, and "Retirement" considers its nature as a store of value and collateral.

Of course, these delineations of the uses of housing are especially abstract and academic. And what is striking about housing, particularly when viewed from the vantage point of household decision-makers, are the sheer number and variety of its practical uses over the life-cycle. Though this dissertation is certainly non-exhaustive in cataloguing these practical uses, it gives a flavor of the longevity of the relationship between households and homes. "Amenities," for example, acknowledges how young parents may use housing to secure an education for their children whereas "Retirement" touches on how the elderly may live off of or bequeath their housing wealth.

As a final note, housing is distinctive purely for its size and, in particular, its size relative to the balance sheet of owner-occupiers. For many, it is the single largest financial transaction they will enter in their lives. This attribute plays an important role in "Amenities". In that chapter, I argue that while house prices are a reasonable place to look for amenity valuations, the period of home purchase is, ironically, a strange time in the life-cycle to measure household willingness to pay for amenities. Because housing is such a big purchase it requires them to exhaust the credit and savings available to them, distorting our estimates of valuation. The sheer size of housing assets surfaces also in "Retirement". A remarkable feature of retirement dissavings is that, in the aggregate, households tend not to consume their housing wealth. In this chapter, I argue that household may in fact have excess holdings of housing.

1.2 Granular scrutiny of micro data

Analytically, these essays have inherited my preoccupation with granular scrutiny of micro-data. Each features a striking graphic capturing the central finding of the paper. In "Amenities," figure (2.8) plots a bin-scatter of the value of mortgage obligations against mortgage balances in the data. In "Bunching," figure (3.18) depicts a scatterplot of adjusted prices and down-payments in the data. (I then collapse this to a more standard depiction of bunching in figure (3.21).) In "Retirement," I plot the evolution of household balance sheets and refi activity relative to claiming social security in figures (4.3) and (4.4).

These plots are, in a sense, summary statistics of widely available data. The novelty in apprehending them as such is due to a combination of three factors. First, consideration of the household problem at hand so as to lend structural interpretations to features of the data where possible. Second, a careful choice of axes to facilitate this structural interpretation. And, finally, detailed institutional knowledge required to "locate" the household problem, as written down, within the data.

Many variables in finance are isomorphic to each other, and at times half the battle can be choosing the representation that is most revealing, or best facilitates intuition. In "Amenities," for example, moving from mortgage rates and initial balances to present and future consumption bundles serves to capture household willingness-to-pay for credit. In "Bunching," moving from balances and loan-to-value ratios to a space of down-payments and house prices helps to generate bunching that captures the loan-to-value elasticity of housing demand. Even in "Retirement," mapping into event time, defined as years to claiming social security benefits, proved revealing.

I devote considerable time in each chapter detailing (and at times formalizing) the GSE pricing grids, FHA borrowing limits, and social security eligibility thresholds. The advantage of doing so is to enable transformations of the same data that may, in fact, be more revealing. In "Amenities," it is important to combine both mortgage and mortgage insurance obligations, which becomes clear from the institutional setting. In "Bunching," to recover the bunching, it is essential to plot home prices and down-payments relative to reference quantities implicitly defined by household income and borrowing limits.

Finally, I note that when I turn to formally interpret these images, I use the eclectic toolkit of an applied micro-economist. In particular, where "Amenities" uses hedonic regression and discrete choice modeling, "Bunching" relies on a bunching estimator, and "Retirement" on an instrumental variables strategy in the treatment and potential value framework.

1.3 Frictionless benchmarks for households

To argue anywhere that financing matters requires reckoning with the fact that, in a world absent frictions, it is irrelevant. This is true in any context, and household finance is no exception. In each chapter of this dissertation, I clarify and strengthen my reasoning by considering my empirical findings in light of a frictionless benchmark. In "Amenities," an unconstrained borrower would have a revealed willingness-to-pay for credit of $1 implying no distortion in valuation of school quality. In "Bunching," the unconstrained buyer would be unresponsive to the loan-to-value limit on the first mortgage, implying no bunching at the kink point. And in "Retirement," the rational home-owner, anticipating the need for liquid funds in retirement and the costs of refinancing, would save less in housing wealth.

There are a handful of well-known approaches in the literature to thinking about a frictionless benchmark. These include the no-arbitrage approach of Modigliani and Miller, the Permanent Income Hypothesis, and what I call the "choice sets" approach from the behavioral literature. I consider these in turn and their practical application to empirical household finance. I describe and evaluate them here informally. Where I describe limitations, it is not to suggest flawed reasoning or even lack of extensibility. My aim, instead, is to provide some intuition for why, as received, they proved unsuitable for the task at hand.

Most famously, in the corporate context, is the Modigliani-Miller observation that the division of ownership into debt and equity does not change, fundamentally, the value of the assets of a firm, barring financial frictions. One approach to assessing this comes from investment/cash-flow regressions, though these are somewhat impractical in a household context given more limited investment activity. Another approach is an event-study approach around a financing decision where the outcome of interest is the abnormal returns on the firm's equity. This is impractical in a household context because of a lack of data; a household's net worth is not traded as a security.

In the household context, one approach to a frictionless benchmark comes from the permanent income hypothesis (also due, in part, to Modigliani), which observes that access to financing ought to make the timing of consumption independent of the predictable component of income. Evaluation of this hypothesis tends to examine household income and consumption patterns rather than financing choices. In fact, although the availability of financing is instrumental to these results, this approach tends to treat financing abstractly, modeling it and often modeling only a single margin for financing. Of course, what it means for financing to be frictionless is that, even if a given margin for borrowing is restricted, it is possible to finance around this along another margin.

The approach in the behavioral literature to identifying frictionless benchmark is what I term the "choice set" approach. This approach relies on the notion that capital is fungible. For this reason, given a set of well-defined financing arrangements, it may be possible for the researcher to order which is best by which is cheapest. Revealed preferences that deviate from this benchmark may indicate behavioral frictions; revealed preferences indicating indifference provide occasion for nudges. This approach is well-designed to assess household behavior but, because the choices are often taken to be exogenous rather than viewed as market outcomes, has less to say about the function of capital markets.

The approach worked out in this dissertation - most clearly stated in "Amenities" - takes cues form each of these literatures. The main idea, following the "choice set" approach, is to construct a set of choices over financing arrangements for the household. (This is represented by the mortgages available through the GSE grid in "Amenities" or the down-payment and home price pairs available through FHA lending in "Bunching".) This set of choices is supplemented with financing along another margin, a hypothetical frictionless margin priced according to the capital markets. (This is represented by the "non-mortgage" margin in both "Amenities" and "Bunching".) This latter margin captures the intuition, from Modigliani-Miller, that in the absence of frictions, households should be able to access credit in the capital markets by some other means. However many lines of credit we observe available to the household, the frictionless benchmark suggests an 'n+1'th line of credit available directly from the capital markets.

Finally, there are ways in which the analysis of household behavior under credit constraints resembles the more familiar analysis of household preferences more generally. I highlight two of them below.

In hedonics, given the price and characteristics of a choice set, and households' choice of home, economists infer something about households' preferences, which are otherwise inscrutable. Because of capital's fungibility, under the frictionless benchmark, households' willingness-to-pay for credit is pinned down by the market price. Introducing credit constraints once again makes households' preferences for credit inscrutable, but the hedonic approach can be used to reveal these now subjective preferences for credit.

Moreover, because prices are linear in quantities, with recourse to a frictionless line of credit, household willingness-to-pay for credit along any other margin is always linear. When their access to the frictionless line of credit is constrained, however, their willingness-to-pay for credit is diminishing in the amount of credit they get access to because of diminishing returns to consumption. This re-introduction of concavity to the problem plays an important role in "Bunching".

CHAPTER 2

CREDIT CONSTRAINTS AT HOME PURCHASE AND BIAS IN HEDONIC AMENITY VALUATIONS

2.1 Introduction

Where households live is instrumental to their well-being [Kling et al., 2007, Clampet-Lundquist and Massey, 2008, Ludwig et al., 2012] and the future prospects of their children [Chetty et al., 2015, Chyn, 2018]. This is due, in part, to the fact that different neighborhoods offer different amenities, goods like effective schools or safe streets to which households gain access by sheer proximity. Government bodies are often responsible for overseeing provision of these local amenities, many of which are public goods by nature. And as a normative matter, complementing concerns of fairness and externalities, the optimal provision of amenities should be responsive to household preferences and the private benefits that households derive from them.

In this paper, I argue that the standard hedonic techniques used to estimate the private value of amenities are downward biased. Standard techniques compare the prices of otherwise similar homes, one with access to more of the amenity; the difference in prices is interpreted as the household's willingness-to-pay (WTP) for the additional amenity. For credit constrained households, however, payment requires a sacrifice of non-housing consumption that is especially burdensome because it is concentrated at the time of home purchase. What they are willing-to-pay understates what they would be willing-to-pay if the payments could be better financed.

I further argue that household mortgage choice is informative about the size of the bias in amenity valuations. A household's WTP for consumption at the time of home purchase is reflected in the increase in future payment obligations it is willing to accept to increase the size of its mortgage balance. Given a borrower's choice of mortgage, this can be read off the slope of a menu of mortgage contracts quoted in terms of mortgage balances and payment obligations. If a household is unconstrained, it can always smooth consumption along some other margin and will be unwilling to pay more than \$1 in present value of future payment obligations for an increase in a loan balance of \$1. A constrained household, by contrast, may be willing to pay more.

In my empirical analysis, I exploit features of the U.S. agency loan market, the GSE loan-level price adjustment grids and private mortgage insurance requirements, to construct borrower-level mortgage menus. Using these menus and borrowers' chosen contracts, I find that, on average, households will pay \$1.65 in present value of future payment obligations for an increase in a loan balance of \$1. I devise a method for correcting hedonic estimates that accounts for joint heterogeneity in WTP for credit and amenities. Incorporating my borrower-level estimates of WTP for credit, I correct estimates of mean marginal WTP for school quality and find that standard estimates are downward biased by ~50%.

The standard hedonic approach to measuring the private value of amenities is due to Rosen [1974], who regresses rents on housing characteristics, including measures of local amenities. By describing an equilibrium model of housing supply and demand in characteristic space, he gives an interpretation to the coefficients obtained in such a regression, namely, households' mean marginal WTP for amenities. Effectively, the difference in rental rate between two otherwise-identical homes, one of which has access to, say, better schools, describes household WTP for the additional school quality.

Rosen [1974] proposed a regression in housing rents, but hedonic estimates are often obtained in terms of house prices. This is because most households are home-owners and because house price data is more readily available than rental data. By analogy, such estimates purport to recover the capitalized rather than flow value of amenities. But home purchase, unlike rental, requires financing, introducing the possibility of bias. For constrained households, the costs of accessing better amenities cannot be smoothed but instead are concentrated in the period of home purchase. Therefore, what households *will* pay out of *constrained* consumption at home purchase for more amenities understates what they *would* pay as a flow of *unconstrained* consumption over time.

The distinction is important from a policy perspective, the view which hedonic estimates very often inform. A government considering the costs and benefits of a proposed intervention may not be subject to financing constraints of the kind faced by households when choosing a neighborhood. Instead, it may raise the funds for amenity improvements from a flow of taxes. As a consequence, the benefits of improving amenities should be valued relative to the costs of unconstrained consumption. Using "traditional" estimates without correcting for bias would overlook welfare-improving investment opportunities.

To better apprehend the mis-measurement of policy-relevant amenity valuations, I begin by re-interpreting the estimates captured in hedonic regression. What is termed 'willingness-to-pay' in hedonic regressions is, substantively, a marginal rate of substitution (MRS) between housing and non-housing consumption. Somewhat subtly, the housing rents or prices in the regression represent *non-housing* consumption being sacrificed in favor of various amenities. The challenge is that hedonic regressions in prices capture a MRS between housing and *constrained* non-housing consumption, whereas *unconstrained* non-housing consumption is the alternative of policy relevance.

To correct the bias in estimates, I apply the logic of the hedonic regression to a novel domain. Instead of a regression of house prices on home attributes, I consider a regression of mortgage obligations on mortgage balances, which I term a "financial" hedonic regression.[1] The coefficient now captures the MRS between present and future non-housing consumption. To a first degree of approximation, households are credit constrained at the time of home purchase and not in the later course of home-ownership. Therefore, the "financial" hedonic regression delivers the MRS between constrained and un-constrained non-housing consumption. As discussed above, the "traditional" hedonic regression in prices recovers the

1. More generally, this regression could be implemented by regressing the price of a financial product on state- and time- indexed payoffs. To the extent that the menu of contracts features variation along all relevant states and household beliefs are known, it is possible to recover SDFs, which are just a collection of MRS's between consumption in different states and times. This paper sidesteps considerations of household beliefs and state-indexed consumption by focusing only on the MRS between present and future consumption.

MRS between housing and constrained non-housing consumption. Combining the two, we obtain the MRS of policy interest.

I implement my proposed bias correction by exploiting institutional features of the U.S. mortgage market. To obtain the credit guarantees required to securitize a mortgage in the agency market, originators must conform to standards set by the GSEs. In particular, they must pay loan-level fees to the agencies and, for loans originated above an 80 loan-to-value ratio, borrowers must obtain private mortgage insurance (PMI). The size of the fees and the PMI premiums both depend on the leverage of the underlying loan. Jointly, these requirements establish a market-wide menu of mortgage contracts in which borrowers can lever up but must pay higher effective interest rates to do so.

With the GSE requirements establishing an effective menu of contracts, I can conduct two exercises. First, in the spirit of "financial" hedonic regression, I estimate the slope of the mortgage price schedule to find households' WTP for credit, or MRS between present (constrained) and future (unconstrained) consumption, at the time of home purchase. A variety of measurement concerns arise in this setting, including the value of the default and prepayment options embedded in mortgages. I address these by conducting a sensitivity analysis in my estimates.

Second, I turn to the exercise of correcting estimates of WTP for amenities. Because there is cross sectional heterogeneity in WTP for both amenities and credit, the bias correction features a covariance term that cannot be estimated with simply the mean WTP for credit. I use the well-defined nature of the mortgage menu and publicly-available information on GSE pricing grids and PMI rate cards to calculate WTP for credit at the individual level. I devise and implement a strategy that uses this individual-level variation to correct estimates of mean WTP for school quality that accounts for this joint heterogeneity.

Ultimately, I find that the WTP for \$1 of credit at the time of home purchase has mean ~\$1.65. This is the most conservative estimate of the various sensitivity analyses that I run. The mean suggests that bias in hedonic price estimates is on the order of 65%. I use individual level estimates of WTP for credit to correct a "traditional" hedonic regression. When I correct for bias induced by credit constraints, I find that coefficients on district school quality increase by 50%.

The remainder of the paper proceeds as follows. In Section (2.2), I discuss the paper's contribution to each of several strands of literature. In Section (2.3), I formalize the main intuitions of the paper. In Section (2.4), I provide an overview of the GSEs in the mortgage market and the market-wide menu of mortgage contracts they establish through their loan-level fee structure and PMI requirements. In Section (2.5), I describe my data sources and sample construction. Section (2.6) presents results measuring the extent of household credit constraints at the time of home purchase from data on mortgages. Section (2.7) uses information on household credit constraints to correct estimates of household willingness-to-pay for amenities. Section (2.8) concludes.

2.2 Related Literature

This paper revisits a longstanding literature, dating back to the framework of Rosen [1974] and Roback [1982], that uses land rents to value local public amenities like school or air quality. The plausibility of these estimates has been improved by empirical papers using quasi-random variation to reduce selection concerns [Black, 1999, Chay and Greenstone, 2005, Bayer et al., 2007]. Recent papers have raised the possibility of bias in these estimates due to equilibrium effects [Kuminoff and Pope, 2014], limited information and distorted beliefs [Gao et al., 2021], and interest costs in lending [Ouazad and Rancière, 2019]. The hedonic approach continues to be used to inform public policy debates and continues to be evaluated in house prices rather than rents [Currie et al., 2015, Kulka, 2019, Diamond and Mcquade, 2019]. I contribute to this literature in two ways: (i) I argue that borrowing constraints and the resulting shadow-cost of credit may introduce measurement bias and (ii) I propose a technique for measuring this bias by introducing the logic of the hedonic regression to mortgage-choice rather than housing-choice.

In recent years, a large literature has used discrete-choice structural models to investigate household financial product choice. Various papers have studied student debt [Ebrahimian, 2020], mortgage choice [Benetton, 2021, Robles-Garcia, 2018], auto debt [Grunewald et al., 2020], credit card debt [Nelson, 2018], and ETF choice [Egan et al., 2020]. Several of these papers raise the possibility that borrowers may be credit constrained, either a channel of interest or an alternative explanation. None of these papers, however, characterizes product attributes as their time- and state-dependent costs or payoffs. This approach is generally clarifying because the marginal rate of substitution between consumption in various states and times is the household stochastic discount factor (up to information about household beliefs). More specifically, it provides an approach for measuring the effects of credit constraints directly in cases where it is possible to observe a menu of potential financing arrangements.

I use the GSE LLPA grid and PMI rate card grids as a basis for the menu of financing arrangements available to borrowers. In the paper, a borrower's choice of cell is revealing about their willingness to substitute non-housing consumption inter-temporally. Conceiving of these grids as a market-wide menu contributes to a literature on their effects and research uses. Fuster et al. [2013] uses these grids to measure lender profitability of loans. Hurst et al. [2016] notes that because the guarantee fees do not account for information on local housing market conditions, the grid effects a large inter-regional insurance program. Bartlett et al. [2021] notes that within LLPA grid cells, lenders are not differentially exposed to borrower credit risk, and uses the grid to study discrimination in mortgage lending.

Finally, the determinants of housing leverage has been an area of active research since the Great Financial Crisis. During the early 2000s, rising leverage was driven by the expansion in availability of sub-prime mortgage credit [Mian and Sufi, 2009] as well as improved housing collateral values against which households could borrow [Mian and Sufi, 2011]. At a microeconomic level, DeFusco and Paciorek [2017] measures the effect of interest rates on first mortgage balances and Bailey et al. [2019] finds a limited role played by household beliefs about housing prices. DeFusco et al. [2020] considers the role played by government policy in limiting mortgage leverage. This paper contributes to the literature by characterizing the mortgage leverage choice when a household (i) has multiple margins for borrowing and (ii) faces rising marginal costs of mortgage credit.

2.3 Framework

2.3.1 The Household Problem

I consider an infinitely-lived HH purchasing rather than renting a unit of housing. The household allocates wealth in the form of initial savings, a_0, and wages, $\{w_t\}_t$, between non-housing consumption each period, $\{c_t\}_t$, and an amenity, s, priced according to an equilibrium hedonic schedule of prices, $P(s)$. The household has a non-mortgage margin for borrowing at the risk-free rate, but faces a credit constraint, which I normalize to zero. The household may also finance the purchase of the home with a mortgage loan. The household chooses initial balance B^o, which it receives in exchange for periodic payment obligations, $B^o r(B^o)$. Here, $r(B^o)$ describes the menu of available mortgage contracts and $r'(B^o) > 0$.

Formally, I write:

$$\begin{aligned}
\max_{\left\{\{c_t\}_{t\geq 0},\{a_t\}_{t\geq 1},s,B^o\right\}} \quad & u(c_0, s_0) + \sum_{t=1}^{\infty} \beta^t u(c_t, s) \\
\text{s.t.} \quad & a_0 + w_0 = c_0 + [P(s) - B^o] + \frac{a_1}{1+r} && (\lambda_0) \\
& a_t + w_t = c_t + B^o r^m(B^o) + \frac{a_{t+1}}{1+r} \qquad \forall t > 0 && (\lambda_t) \\
& a_t \geq 0 \qquad \forall t > 0 && (\mu_t) \\
& r^{m\prime}(B^o) > 0
\end{aligned} \tag{2.1}$$

I describe a solution to the HH problem in Section (A.1.1) using the Kuhn-Tucker conditions. I assume a solution in which the credit constraint is non-binding after the initial savings decision, $\mu_t = 0 \;\forall\; t > 1$. Rearranging the first-order conditions, I obtain two necessary conditions for the behavior of the optimizing HH.

I also define what I term a "credit wedge", $\kappa_t^{t+1} \equiv \frac{MRS_t^{t+1}}{1+r_t}$, as the ratio between the optimizing household's subjective marginal rate of substitution and the market price ratio (or interest rate) between time t and $t+1$.[2] This "credit wedge" may be represented as $\kappa^{t+1} = \frac{\lambda_t/\lambda_{t+1}}{1+r_t} = 1 + \mu_{t+1}/\lambda_{t+1}$, using the Lagrange multipliers and first order conditions of the household's problem. The "credit wedge" may therefore be interpreted as present-valued willingness to pay out of future consumption for additional consumption today. It may also be interpreted as the shadow price of credit. Note that the requirement that $\mu_{t+1} \geq 0$ imposes that $\kappa^{t+1} \geq 1$.

With the "credit wedge" defined, I write the first-order conditions as follows:

$$P'(s)\Big|_{s^*} = \frac{u_s/u_c^1}{r}\frac{1}{\kappa^1} \tag{2.2}$$

$$\Big[\frac{B^o r^m(B^o)}{r}\Big]'\Bigg|_{B^{o*}} = \kappa^1 \tag{2.3}$$

Equation (2.2) interprets the information content in the optimizing household's choice of amenity level from the slope of the hedonic price schedule. This is depicted in Figure (2.1); households optimize by setting their indifference curves tangent to the offer curve in amenity space but their amenity curves are no longer only determined by wealth and preferences.

When the HH credit constraint is non-binding even in the initial savings decision ($\mu_1 = 0$ and $u_c^1 = u_c^0$), then the hedonic price schedule captures the capitalized and consumption-valued service flow from the marginal unit of the amenity, $\frac{u_s/u_c^0}{r}$. This is the motivation for "traditional" hedonic regressions.

However, when the HH credit constraint is initially binding, the slope of the price schedule is biased relative to the HH's WTP for the marginal amenity out of future, *un-constrained* consumption. The magnitude of the bias, somewhat intuitively, relates to the HH shadow price of credit, μ_1/λ_1, which describes how tightly the credit constraint binds. This suggests

2. I write only $\kappa^{t+1} \equiv \kappa_t^{t+1}$ for clarity.

that "traditional" hedonic regressions may be biased. The influence of household credit constraints on their indifference curves is depicted in Figure (2.2).

Equation (2.3) interprets information content in the optimizing household's choice of initial mortgage balance. Note that the information is observable in $\frac{B^o r^m(B^o)}{r}$, which describes the menu of mortgage contracts available in the space of PDV of future payment obligations. Thus, there is a second price schedule, on the financing side, with a slope that captures information about the optimizing household. The mapping between menus in rates and menus in future mortgage obligations is depicted in Figure (2.3).

The slope of the mortgage contract menu captures the borrower's shadow value of credit, κ^1. This equivalence is intuitive. An optimizing household with multiple margins for credit must equate the price of additional credit along each margin. Because mortgage rates are increasing in initial balances, each marginal dollar of mortgage borrowing costs more. The optimizing household will increase mortgage borrowing until its cost exactly equals the shadow-cost of non-mortgage borrowing. Optimization along the margin of mortgage borrowing is depicted in Figure (2.4).

The slope of the mortgage contract menu captures the bias in the hedonic price schedule relative to the value of amenities out of unconstrained consumption. Here, the intuition is that the mortgage menu captures the trade-off between present and future consumption, which are assumed to be constrained and un-constrained, respectively. The distortion in the amenity valuation comes precisely because of the difference in value of constrained and unconstrained consumption.

Figure 2.1: Constrained household amenity optimization

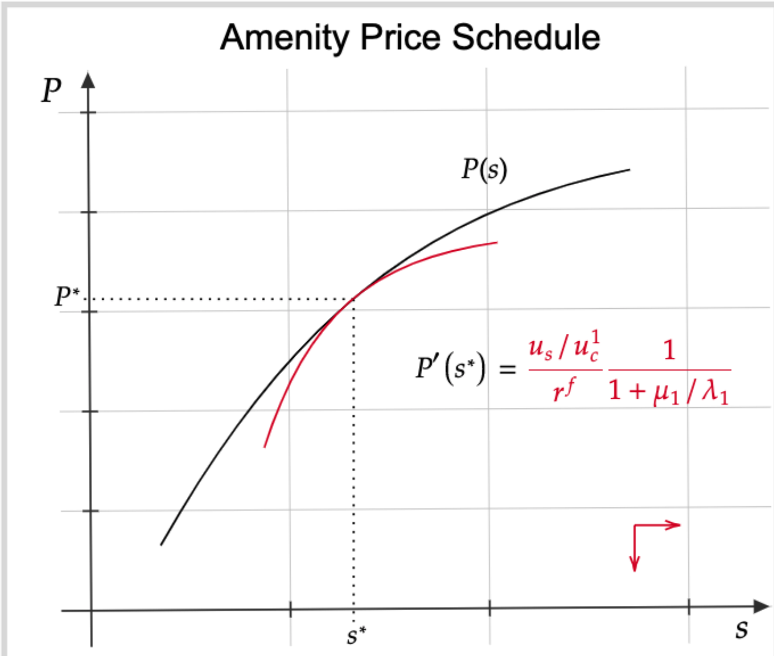

The housing market (black) consists of a menu of units, each characterized by an amenity level, s, and a price, P. The household prefers lower prices and higher levels of the amenity. It optimizes by setting its indifference curve (red) tangent to the amenity price schedule. Because households may be credit constrained at the time of home purchase, the slope of the market menu captures a downward biased version of their willingness-to-pay for the amenity.

Figure 2.2: Constrained household amenity re-optimization

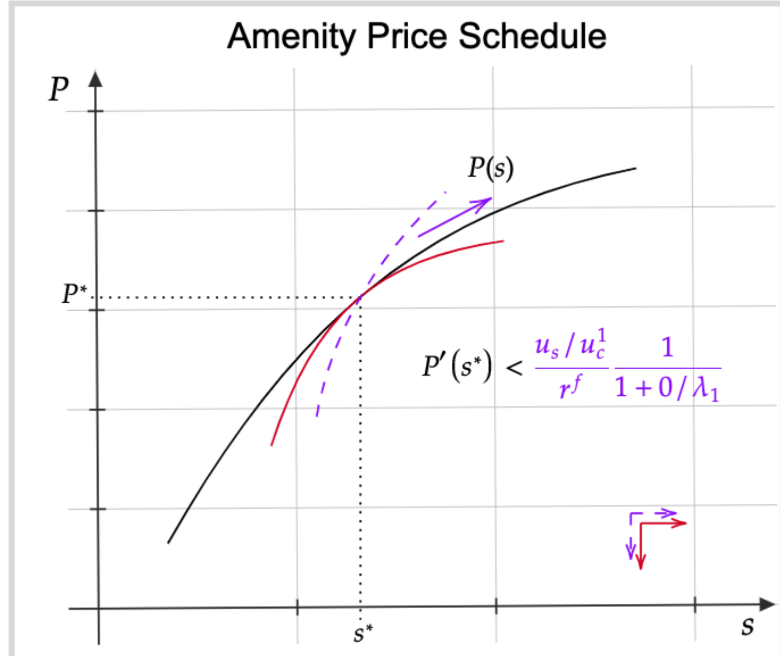

To recover the household's willingness-to-pay for marginal units of the amenity, it helps to consider a thought experiment in which the household solves its constrained optimization problem (red). Then, without changing its allocations, its credit constraint is lifted. The household's new indifference curves (dashed purple) are now steeper than the amenity price schedule. More precisely, the household's indifference curve grows steeper by a factor of $1+\mu_1/\lambda_1$.

Figure 2.3: Convexity of the mortgage price schedule

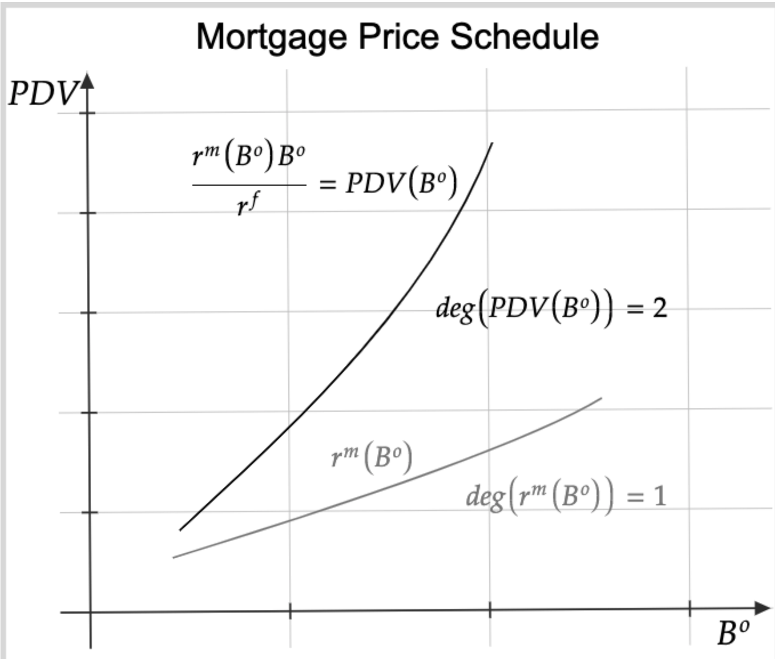

Borrowers are most often quoted the price of a loan in terms of an interest rate, thus facing a menu of loans varying in interest rate by initial balance (gray). This interest rate menu is equivalent to and more easily interpreted as a menu of loan balances and the present discounted value of payment obligations (black). Note that present discounted value scales with mortgage payment, which is the product of loan balance and interest rate. So, if interest rates are first-order in loan balances, payment obligations are second order, and therefore convex.

Figure 2.4: Constrained household mortgage optimization

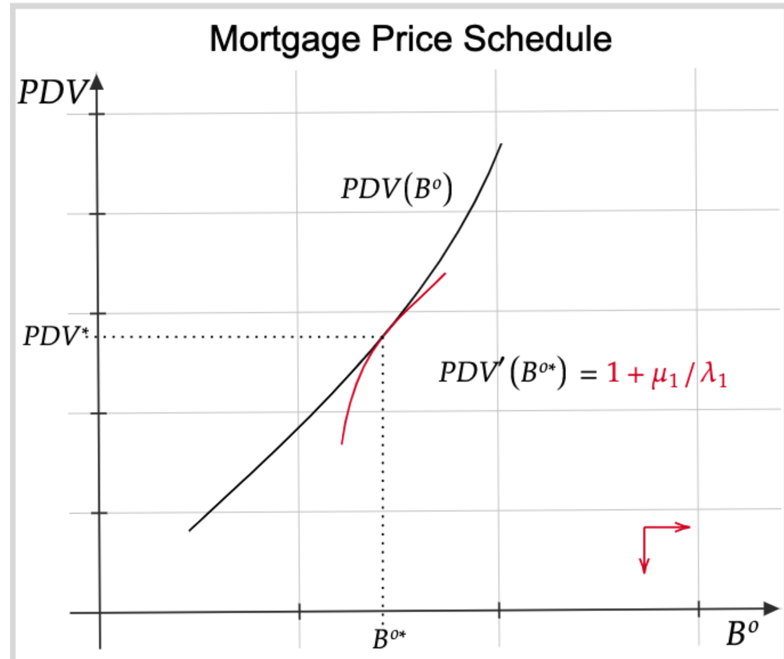

The mortgage market (black) consists of a convex menu of contracts, each defined by an initial balance, B^o, and the present discounted value of future payment obligations, PDV. The household prefers larger initial balances and lower payment obligations. It optimizes by setting their indifference curves (red) tangent to the market menu. The slope of the mortgage menu captures the household's shadow price of credit, $1 + \mu_1/\lambda_1$. This is precisely the downward bias of the amenity price schedule relative to household willingness-to-pay for amenities.

2.3.2 Selection and Moral Hazard in the GSE Problem

In principle, it is possible that rising interest rates in the mortgage offer curve does not reflect any information about the willingness-to-pay of HHs for credit. Suppose, for instance, that as borrowers lever up, they suffer more from moral hazard and repay a lower fraction of their mortgage obligations. The rising interest rates then simply reflect the increased credit risk associated with lending.

There are two potentially mitigating factors here. First, if households do not anticipate the extent to which they will default on loan obligations, their choice of contract may still be informative about their credit constraints. Second, if there are heterogeneous households whom the GSEs cannot or do not distinguish (note that the LLPAs use fairly little information), then whatever the equilibrium pricing schedule, for good types who will not default, the choice of mortgage is still informative about credit constraints.

In this section, I further reconcile the presence of credit risk with the informative-ness of the offer curve slope by sketching a formal (but very rudimentary) model of GSE offer curve choice. Rather than considering moral hazard, I consider a setting of asymmetric information in which the GSE faces a distribution of borrower types who represent varying credit risks. Recent work suggests that selection rather than moral hazard plays a predominant role in driving the correlation between leverage choice and moral hazard [Gupta and Hansman, 2021].

Consider a benevolent lender, the GSE, who aims to maximize some social welfare function subject to a zero-profit condition. The lender faces a distribution of borrowers, f, who vary in their WTP for credit, κ_i, and the fraction of promised obligations they will actually repay, θ_i. The lender states an offer curve, $PDV(B^o)$, of mortgage payment obligations given an initial balance of borrowing and borrowers choose the contract that suits them best. We write:

$$\max_{PDV(B^o)} U(\{B_i^o\})$$

$$\text{s.t.} \quad \int_i \theta_i PDV(B_i^o) f(\theta_i) d\theta_i = \int_i B_i^o f(\theta_i) d\theta_i \qquad (0\Pi) \tag{2.4}$$

$$\theta_i PDV'(B_i^o) = \kappa_i \qquad \text{(IC)}$$

Assuming that the social welfare function places an infinite penalty on redistribution, the lender is now restricted to offer curves that break even, in expectation, loan-by-loan. The zero-profit conditions becomes a more restrictive condition:

$$PDV(B^o)\theta(B^o) = B^o \tag{2.5}$$

I also assume that all borrowers have the same WTP for credit, $\kappa_i = \overline{\kappa}\ \forall i$. Using the (IC) constraint, it is now possible to solve for the type of borrower at a given point on the offer curve. Plugging this into the loan-wise zero-profit condition, we obtain:

$$\overline{\kappa}\frac{PDV(B^o)}{PDV'(B^o)} = B^o \tag{2.6}$$

This differential equation is straightforward to solve and the solution to the GSE's problem is then:

$$PDV(B^o) = (B^o)^{\overline{\kappa}} \tag{2.7}$$

For values of $\overline{\kappa} > 1$, this is convex, which captures a feature of the observed offer curve in the data. Note that as $\overline{\kappa} \to 1$, the curve becomes more and more linear. Although the borrower types drives the sorting behavior, it is the WTP for credit that drives the degree of convexity.

2.3.3 Mis-measurement and Under-investment in Local Amenities

In this section, I formalize the notion that a welfare-maximizing government without borrowing constraints underinvests in amenities if it infers household willingness-to-pay from the house price envelope without correcting for the extent of household credit constraints.

The intuition is that the government can borrow and finance the investment out of future tax income, which falls on the households not when they purchase the home, but in periods when they are unconstrained. The households are therefore willing to pay more for additional amenities out of future taxes than they appear to be at the time they purchase a home.

In the model, households choose housing in municipalities of varying amenity levels. A price envelope forms in equilibrium that makes households indifferent between the municipalities, so that the housing market can clear. The (federal) government has access to an investment project that increases the amenity level of all municipalities. The cost of investment is convex in the quantity of additional amenities. The government may borrow to finance the investment and recover the costs from tax revenue at a later date. The government cannot observe household preferences directly, but they can measure the equilibrium prices in the housing market to infer information about preferences.

The equilibrium price envelope in the housing market depends on whether households are constrained or un-constrained at the time of home purchase. If households are unconstrained, then the average slope of the price envelope is a sufficient statistic for the optimal level of government investment. If households are constrained at the time of home purchase but the government treats them as though they are un-constrained, then it will under-invest in local amenities. In the case that households are constrained, the sufficient statistic for determining the optimal level of government investment is the slope of the price schedule corrected for the extent of credit constraints, the corrected willingness-to-pay put forward in this paper.

Model Set-up

The model contains a unit mass of municipalities, j, a unit mass of households, i. Each municipality has enough room to house a single household and is endowed with a uniformly distributed level of the amenity, $s_j \sim U[0,1]$. The households are identical and have preferences over consumption and the amenity, an endowment of income becoming available in each of two periods, $\{y_0, y_1\}$, and a savings technology with borrowing limit $-\phi$.

The (federal) government in the model has access to an investment project to improve amenity quality by some margin, σ, at cost $I(\sigma)$. It can borrow and save frictionlessly and can impose a uniform flat tax on all households to fund the investment. Finally, the government cannot observe household utility directly but can observe the equilibrium price envelope in the housing market to learn about preferences. This is analogous to the way in which economists estimate hedonic regressions or discrete choice models to measure household subjective willingness-to-pay for amenities.

The model has two periods, $\{0, 1\}$, and the first period features two sub-periods, $\{0a, 0b\}$. At time $0a$, equilibrium is established in the housing market. Households choose municipalities and the equilibrium price schedule forms such that the market clears and households are no better off moving to a different municipality. Households then make their savings decisions and consume their time 0 consumption. At time $0b$, the government observes the price schedule, and chooses its investment and tax policy to maximize the welfare of households. Households do not anticipate this government intervention, nor do they re-optimize in response to it. At time 1, households use their income endowment and savings to pay their tax bill and consume the rest. Amenities also realize improvements due to government investment and households consume these improved amenities.

Housing Market Equilibrium

For the market to clear, municipalities must each host a single household. If two municipalities offer households different utility under any amenity price schedule, no household will prefer to live in the inferior municipality. This will create excess demand elsewhere in the market, implying that the amenity price schedule is not an equilibrium. We can therefore use the household problem to characterize the equilibrium pricing schedule, which must be set so that the optimizing household is indifferent between amenity choices.

Formally, an equilibrium price schedule, $\widetilde{P}(s)$, is implicitly defined by the following problem:

$$
\begin{aligned}
\overline{U} = \max_{\{c_0, c_1, a_1\}} \; & u(c_0) + \beta c_1 + v(s) \\
\text{s.t.} \qquad & y_0 = c_0 + \widetilde{P}(s) + \frac{a_1}{1+r} \qquad (\lambda_0) \\
& y_1 + a_1 = c_1 \qquad (\lambda_1) \\
& a_1 \geq -\phi \qquad (\mu_1)
\end{aligned}
\tag{2.8}
$$

I solve for properties of the equilibrium amenity price schedule in Appendix (A.1.3). The slope of the equilibrium price schedule has the following property:

$$
\widetilde{P}'(s) = \frac{v'(s)}{u'(c_0^*)} = \frac{v'(s)}{1 + \frac{\mu_1}{\lambda_1}} \tag{2.9}
$$

Note that the equilibrium price schedule has the characteristics of the price schedule in the literature on amenities as amended by the household problem introduced in Section (2.3.1). Namely, it reflects the marginal willingness-to-pay for the amenity out of present non-housing consumption, which is constrained. Alternatively it is a downward biased measure of the willingness-to-pay for the amenity out of future, unconstrained non-housing consumption. (N.B. The marginal utility of future consumption before subjective time discount is 1.)

Optimal Government Investment

The government has access to a project that will improve amenity quality by some continuous margin, σ, at the expense of some convex investment costs, $I(\sigma)$ with $I', I'' > 0$. The government funds the investment through a flat tax that falls equally on all households and has frictionless access to financing. The government then chooses how much to invest in improving the amenity in order to maximize welfare subject to the constraint that it must raise the revenue required for the investment from taxes:

Formally, the government's problem is:

$$\max_{\{\sigma,t\}} \int_i U_i(\sigma,t)di$$

$$\text{s.t.} \qquad \int_i \frac{t}{1+r} di = I(\sigma) \qquad \text{(BC)}$$

$$\text{where:} \qquad U_i(\sigma,t) \equiv u(c_0^{i*}) + \beta(c_1^{i*} - t) + v(s^{i*} + \sigma)$$

$$\{c_0^{i*}, c_1^{i*}, a^{i*}, s^{i*}\} \equiv \arg \max_{\{c_0^i, c_1^i, a_1^i, s^i\}} u(c_0^i) + \beta c_1^i + v(s^i) \tag{2.10}$$

$$\text{s.t.} \qquad y_0 = c_0^i + \widetilde{P}(s^i) + \frac{a_1^i}{1+r} \quad (\lambda_0)$$

$$y_1 + a_1^i = c_1^i \quad (\lambda_1)$$

$$a_1^i \geq -\phi \quad (\mu_1)$$

$$s^{i*} = s^i = s_j \quad \text{for} \quad i = j$$

In Appendix (A.1.4), I obtain the following first-order condition for the government that pins down the optimal level of amenity improvement:

$$\sigma^* = I'^{-1}\Big(\mathbb{E}^i\big[\widetilde{P}'(s^{i*})(1 + \mu_1^i/\lambda_1^i)\big]\Big) \tag{2.11}$$

I also define the level of amenity improvements, σ^g, undertaken by a government with a potentially misspecified model. The government measures the slope of the price envelope but assumes that households are unconstrained at the time of home purchase. This level of amenities is given by:

$$\sigma^g \equiv I'^{-1}\Big(\mathbb{E}^i\big[\widetilde{P}'(s^{i*})\big]\Big) \tag{2.12}$$

In Appendix (A.1.4), I show that if households are truly unconstrained, and the government's assumption is correct, that the level of government investment is equal to the optimal level of investment, $\sigma^g = \sigma^*$. If the households are constrained, however, the government invests less than is optimal, $\sigma^g < \sigma^*$. The intuition is simply that the slope of the price envelope is made less steep when households are constrained at home purchase. Households pay for the government funded amenities out of future consumption rather than present consumption and are therefore willing to pay more for the amenities.

2.4 Institutional Setting

In this section, I describe the institutional setting of the agency loan market. In Section (2.4.1), I describe the loan-level pricing adjustments and private mortgage insurance requirements required by the GSEs. In Section (2.4.2), I assume that GSE fees are passed through to borrowers and describe how to construct the menu of contracts available to borrowers at the time of origination using borrower FICO scores; mortgage balance, loan-to-value and interest rate; LLPA grids; and insurance rate cards. In Section (2.4.3), I provide evidence that these fees are passed through from lenders to borrowers at the time of origination.

In describing the institutional setting, I document how requirements imposed by the GSEs at the level of the market generate mortgage menus at the level of the borrower. These menus correspond to the interest rate menu of the household problem, $r^m(B^o)$. In the empirical setting, though, the effective rate on the loan balance will be a combination of the rate due to the mortgage, r^m, and the rate due to the mortgage insurance, r^{mi}. Additionally, the rate menu in the empirical setting is a step-wise rather than continuous function. In line with the assumption of the household problem that $r^{m\prime} > 0$, the rate menu is an increasing function.

The description of the institutional setting here also facilitates the empirical analysis conducted in Section (2.7). Ultimately, I use information on the borrower's chosen mortgage relative to non-chosen alternatives to extract information about the extent of borrower credit constraints. I am able to construct borrower-level mortgage menus precisely because of the standardized and transparent way in which fees and PMI requirements are assigned and passed through.

2.4.1 Agency Loan LLPAs and Required PMI Premiums

In the secondary mortgage market, originators sell loans into collateral pools that are divided into tranches and sold to investors as mortgage-backed securities (MBS). The vast majority of these pools are "agency" pools, which require that loans enjoy a credit guarantee from the GSEs, Fannie Mae and Freddie Mac. In the event of borrower default, the guarantee ensures that the GSEs will step in to cover required principle and interest payments. This arrangement, backed implicitly by the fiscal power of the federal government, insulates the ultimate investors from credit risk and thereby supports the functioning of the mortgage market.

To be eligible for the GSE credit guarantee, a mortgage loan must be "conforming", that is, it must meet standards set forth by the GSEs. These standards include the so-called "jumbo" limit on the size of the mortgage, restrictions on the borrower credit score at origination, and limits on borrower debt-to-income ratios. The GSEs also require that loans with loan-to-value ratios above 80 must be covered by mortgage insurance and they specify the level of required coverage.

Originators must also pay a variety of fees mandated by the GSEs in order to sell mortgages into agency pools. They pay two fees to the GSEs: a one-time up-front insurance premium known as a loan-level price adjustment (LLPA) and an ongoing, monthly "g-fee" that is a small fraction of the loan balance. The GSEs permit originators to trade off between the two by "buying down" or "buying up" the "g-fee" at specified multiples. The GSEs also require originators to pay a 25bp minimum servicing fee to the party servicing the loan. [Fuster et al., 2013]

The LLPAs are determined by a wide variety of loan characteristics, the most prominent of which are the loan-to-value ratio on the loan and the FICO score of the borrower at origination. For vanilla, 30-year, fixed rate mortgages, these completely characterize the

LLPA requirement and the mapping is published and periodically updated in what are known as the GSE “pricing grids”. I hand-collect the contents of the Fannie Mae “pricing grids” from 2009 to present and plot an example in the left panel of Figure (2.5). Loans to less credit-worthy borrowers tend to have higher LLPAs. LLPAs are also increasing in the loan-to-value ratio of the loan, though only to the 80 LTV threshold at which PMI requirements kick in.

For loans above the 80 LTV threshold, the GSEs specify the amount of PMI coverage required for the loan to meet “conforming” standards. Insurers publish “rate cards” describing the premiums (as a percentage of loan balance) of monthly, borrower-paid mortgage insurance at various levels of coverage. These “rate cards” resemble the GSE “pricing grids” in that mortgage loan-to-value ratio and borrower credit score completely characterize the premium, given the required level of coverage. The binning of loan-to-value ratios and credit scores is, in fact, identical to that used to determine LLPAs. For one mortgage insurer, Essent, I hand-collect premiums from rate cards corresponding to the minimum insurance coverage required by the GSEs and plot an example in the right panel of Figure (2.5). Premiums are higher for less credit-worthy borrowers; there are no premiums required below the 80 LTV threshold but premiums are increasing in the loan-to-value ratio above this.

Figure 2.5: Example Pricing Regime (Jan.–Mar. 2011)

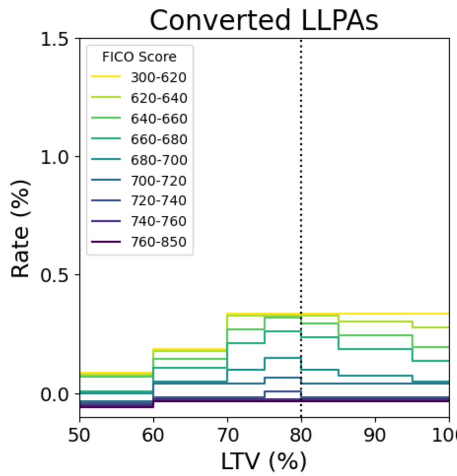

(a) Converted Fannie Mae LLPAs

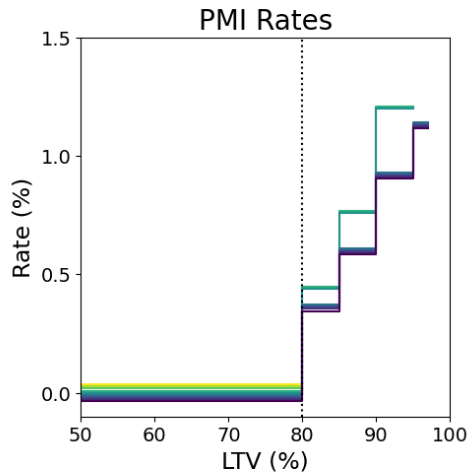

(b) Essent PMI Rates

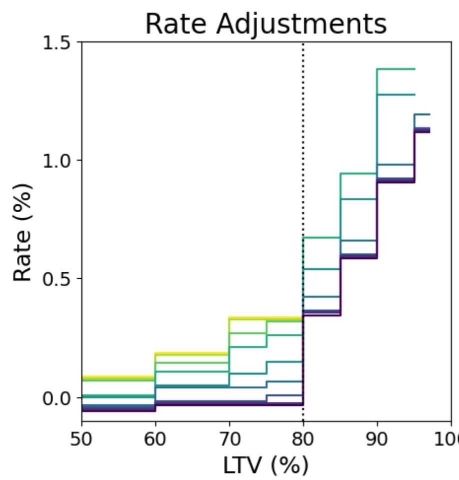

(c) Combined LLPAs and PMI Rates

This figure depicts a sample pricing regime in force from (as early as) January through March 2011.

- Figure (2.5a) contains LLPAs from the Fannie Mae pricing grid. These vary by mortgage LTV and borrower FICO at origination. I convert from LLPA to interest rate adjustment using a conservative factor of 10.

- Figure (2.5b) contains the annual PMI rates quoted on Essent rate cards for the minimum Fannie Mae PMI requirement. These also vary by LTV and borrower FICO.

- Figure (2.5c) contains the sum of the converted LLPAs and PMI rates. This corresponds to the menu of rate adjustment options faced by the borrower at origination of a mortgage.

The LLPAs are increasing to the 80 LTV threshold and thereafter constant or decreasing; the PMI rates are sharply increasing above 80 LTV. Because the rise in PMI more than offsets the fall in LLPAs, the combination is monotonically increasing.

2.4.2 The Borrower-Level Mortgage Menu

Because of the central role of the agency market in the US context, the LLPA and PMI requirements define near market-wide menus for mortgage credit at the borrower level. In particular, information in the rate cards can be used to approximate the cost to a household of levering up on a mortgage in terms of the additional monthly expenses to cover both the LLPA and, possibly, additional PMI payments. Below, I detail how information in a borrower's realized mortgage loan, the GSE "pricing grid", and insurer "rate cards" can be used to approximate the menu of mortgage contracts available to the borrower at the time of home purchase.

Consider a borrower, i, with credit score, S_i, who is purchasing a house at price, P_i, and choosing the amount of initial loan balance, B_i^o, to take out to finance the transaction. Given the choice of loan balance, the loan will have a loan-to-value ratio, $L_i = \frac{B_i^o}{P_i}$. Given the loan balance and loan-to-value ratio, an interest rate will be assigned to the loan, r_i^m, and the borrower may also be required to pay monthly PMI premiums, r_i^{mi}. Given the initial balance, interest rate, and PMI obligations, the borrower will have monthly payment obligations, $(r_i^m + r_i^{mi})B_i^o$.

The GSE pricing grids and PMI rate cards return LLPAs and monthly premium rates, respectively, as a function of mortgage loan-to-value ratios and borrower credit scores. Both are comprised of a sequence of loan-to-value ratio limits, $\{\overline{L}_j\}_j$, and maximum credit scores, $\{\overline{S}_k\}_k$, that divide borrowers into bins. The borrower-level limit is the smallest one greater than the corresponding loan characteristic, $\overline{L}_i = \min_j\{\overline{L}_j \mid \overline{L}_j \geq B_i^o/P_i\}$ and $\overline{S}_i = \min_k\{\overline{S}_k \mid \overline{S}_k \geq S_i\}$.

Estimating the borrower's minimum PMI rates required to enter the agency market is fairly straightforward. GSE minimum PMI coverage requirements, $\underline{f}_t(\overline{L}_j)$, and insurer rate-card matrices, $r_t^{mi}(\overline{L}_j, \overline{S}_k, \underline{f}_t)$, define a matrix of minimum required insurance rates. The borrower's required PMI rate is then:

$$r_{it}^{mi}(B_i^o \mid P_i, S_i) = r_t^{mi}\big(\overline{L}_i, \overline{S}_i, \underline{f}_t(\overline{L}_i)\big) \tag{2.13}$$

Together, the base g-fee requirement, g_t^b, base LLPA, $LLPA_t^b$, excess LLPA matrix, $LLPA_t(\overline{L}_j, \overline{S}_k)$, and buy-up multiples, ϕ_t, imply a 0-UIP g-fee matrix. This matrix, g_t, specifies the g-fee required for a loan to enter the agency market without up-front payments according to the loan LTV and borrower credit score. This can be written as $g_t(B_i^o/P_i, S_i) = g_t^b + \phi_t * \big[LLPA_t^o + LLPA_t(\overline{L}_i, \overline{S}_i)\big]$. Assuming the lender passes along the 0-UIP g-fee to the borrower 1-for-1, and assuming the lender does not otherwise apply any risk-pricing to the loan (the g-fee, after all, secures insurance for the loan), we obtain the menu of mortgage interest rates available to borrowers for different balances at origination:

$$r_{it}^{m}(B_i^o \mid P_i, S_i) = r_i^{om} + g_t^b + \phi_t\Big[LLPA_t^b + LLPA_t\big(\overline{L}_i, \overline{S}_i\big)\Big] \tag{2.14}$$

In the data, we can observe the borrower's chosen mortgage contract, $\{B_{it}^{o*}, r_{it}^{m*}\}$. If the borrower were to borrow some alternate amount, B_i^o, against the same house, we can estimate the consequent mortgage rate. (Note that many GSE pricing parameters fall out in differences.)

$$r_{it}^{m}(B_i^o \mid P_i, S_i) = r_{it}^{m*} + \Delta_{B^{o*}}^{B^o} r_{it}^{m} = r_{it}^{m*} + \phi_t\Big[LLPA_t(\overline{L}_i, \overline{S}_i) - LLPA_t(\overline{L}_i^*, \overline{S}_i)\Big] \tag{2.15}$$

2.4.3 Empirical Evidence on LLPA Pass-through

To estimate household's MRS between present and future consumption, it would be best to observe the actual menu of contracts available to households when entering a bank. The approach of this paper, by comparison, is to construct a household-level choice set of financial contracts on the basis of parameters governing the market-rate securitization of mortgage contracts. Though feasible, this approach is rudimentary; after all, banks are not required pay the 0-UIP g-fee or to pass along g-fees to borrowers one-to-one.

Despite these limitations, I document that banks pass along the g-fees assessed by the GSEs according to the step-wise function in the LLPA matrix. Using a sample of loans sold to Fannie Mae (the "merged sample", described below), I examine interest rates in the cross section of LTV ratios. I capture this variation by regressing loan interest rates on LTV bins, including a large number of controls, including fixed effects for MSA, month of origination, DTI, FICO, race, sex, age, income percentile, and purchase price percentile:

$$r_i \sim \alpha + \sum_{\ell=50}^{97} \beta_\ell 1\{L_i = \ell\} + \gamma' X_i + \varepsilon_i \tag{2.16}$$

To emphasize the breaks where the LLPAs change, I also estimate a piecewise-polynomial defined function on the same data. The domain interval cut-points are defined by the GSE pricing grid, $\overline{L} \in \mathcal{L}_{GSE} = \{60, 70, 75, 80, 85, 90, 95, 97\}$. The estimated function is quartic and the level, cubic, and quartic terms are allowed to vary between interval. The regression specification is:

$$r_i \sim \alpha + \sum_{j=1}^{4} \beta_{j,50}(L_i - 50)^j + \sum_{\overline{L} \in \mathcal{L}_{GSE}} \sum_{j \in \{0,3,4\}} \beta_{j,\overline{L}} 1\{L_i \leq \overline{L}\}(L_i - \overline{L})^j + \gamma' X_i + \varepsilon_i \tag{2.17}$$

The results of this analysis are plotted in Figure (2.6). The step-wise change in rates at the cut-points is salient, particularly when including the high-dimensional fixed-effects as controls. It is also noteworthy that above the 80 LTV threshold, the interest rates charged on loans declines, consistent with pass-through of the declining LLPAs. Borrower monthly obligations in these contracts are higher on net, however, because of PMI obligations, which are changing at the 80 LTV threshold.

I also document considerable bunching in mortgage loans at the LTV thresholds at which the GSEs impose higher g-fees on originators or require additional PMI of borrowers in Figure (2.7). The bunching is consistent with the loan becomes discretely more expensive as borrowers lever up above these thresholds. The bunching suggests a pricing interpretation rather than an alternative interpretation that borrowers use heuristics and reference numbers, e.g. multiples of 5, when choosing their leverage in a home. In particular, bunching does take place at LTV 97 and not at LTV 65. Although 97 is not one of these candidate reference numbers, it is an LTV above which private mortgage insurers charge more for coverage. Conversely, 65 is a reference number but sees no increase in the GSE fees and exhibits no bunching. (Note that although the density is fairly slim, there is observable bunching below this, at 60).

The bunching appears to be particularly pronounced for loans that are sold to the GSEs ex-post. Comparing Figure (2.7a), which contains the CRISM Sample of all conventional, conforming loans, and Figure (2.7b), which contains only loans from HMDA that can be matched to the FNMA loan performance data, it is clear that there is more bunching for loans sold ex-post to the GSEs. To the extent that the originator knows ex-ante whether or not the conforming loan will in fact be sold to the GSEs, this makes sense.

Figure 2.6: Pass-through of Loan-Level Pricing Adjustments

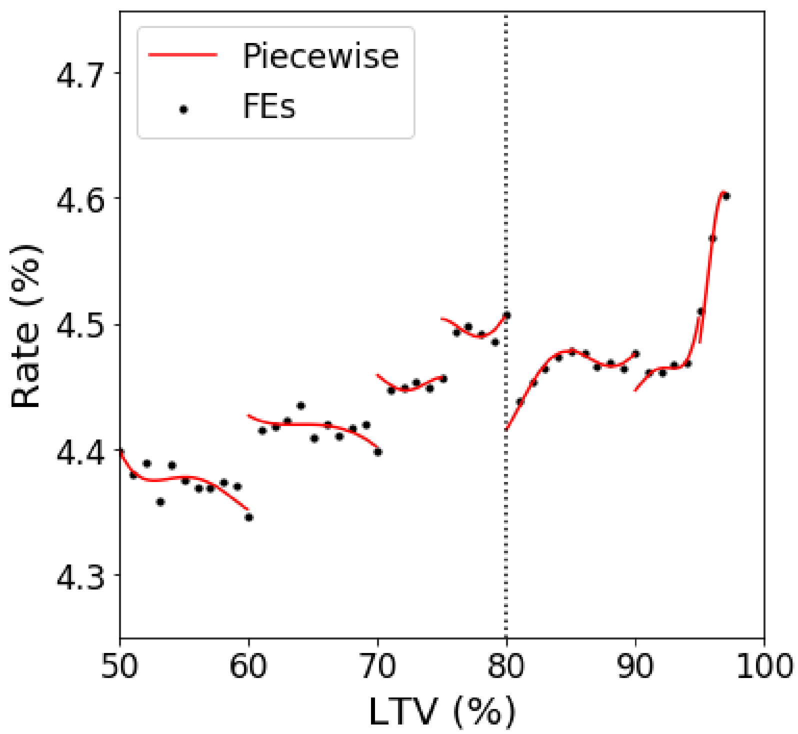

The figure above uses the FNMA-HMDA-SEDA Merged Sample to plot interest rates in the cross section of LTV at origination. The black dots are estimates from the regression with LTV bin fixed-effects; the red lines are estimates from a regression with a piecewise-polynomial defined function. I control flexibly for an array of confounders, including zip code, origination month, and borrower FICO. The rates offered by lenders reflects the step-wise adjustments to LLPAs established by the GSEs.

Figure 2.7: Bunching at LLPA and PMI thresholds

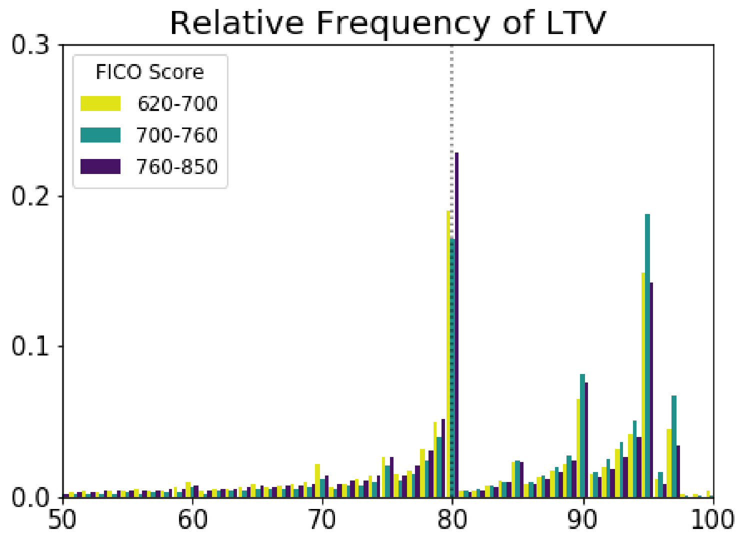

(a) CRISM Sample

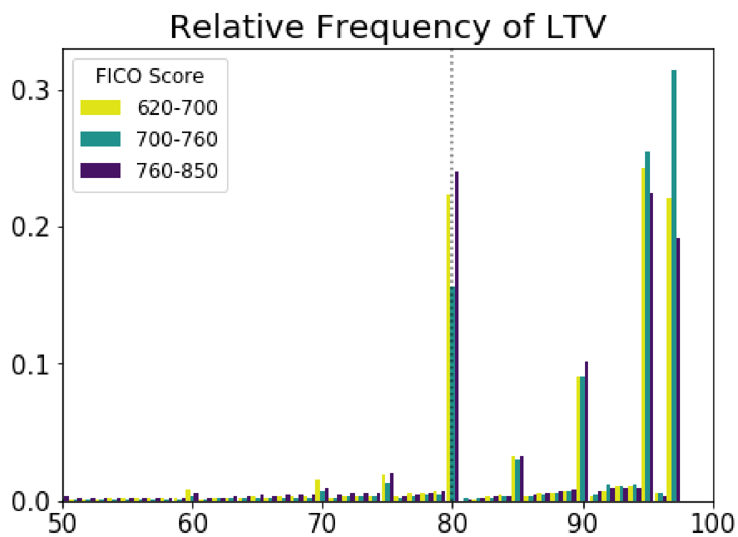

(b) Merged Sample

Histogram of loan-to-value ratios for loans in the CRISM Sample (2.7a) and FNMA-HMDA-SEDA Merged Sample (2.7b). The left plots show bunching over the entire distribution, the right plots show that it is present at all FICO levels. Bunching occurs at loan-to-value amounts where g-fees increase, insurance becomes required, and where insurance premiums increase. Notably, there is no bunching at 65 loan-to-value, a possible reference value where there is no increase in g-fees. Moreover, there is bunching at 97, unlikely to be a reference value that does entail an increase in mortgage insurance premiums or decrease in availability.

2.5 Data

2.5.1 Sources

I use the following data sources in various analyses:

- *Credit Risk Insights Servicing McDash (CRISM), 2005-2020* consists of credit bureau data fields from Equifax merged to servicing records from McDash covering loan origination and performance. Since 2005, CRISM covers roughly 60% of mortgage originations [Adelino et al., 2013].

- *Fannie Mae (FNMA) Single-Family Loan Performance Data, 2005-2020* describes origination and performance characteristics of loans securitized through Fannie Mae.

- *Home Mortgage Disclosure Act (HMDA) Loan Application Register (LAR), 2018-2019* describes origination characteristics of mortgages subject to HMDA reporting requirements. The data has near universal coverage of the US mortgage market.

- *Fannie Mae (FNMA) Loan Level Pricing Adjustment (LLPA) Grids, 2009-2020* publicize the fees that originators must pay in order to obtain FNMA insurance (and securitize the loan through the GSEs). I hand collected grids dated to late 2012 and beyond from the Wayback Machine (https://archive.org/web/)[3].

- *Essent Private Mortgage Insurance (PMI) Rate Cards, 2011-2020* disclose the fees that borrowers must pay in order to obtain PMI through the insurer Essent. I hand collected these rate cards back to 2011 using links to historical rate cards available on Essent's web-page.

- *The Stanford Education Data Archive (SEDA), 2009-2018* is a data product constructed by the Educational Opportunity Project at Stanford University. It includes measures of average standardized test performance at the school level for the school

3. Thank you to Andreas Fuster for providing me with grids from 2009-2012.

years 2008-2009 to 2017-2018. State-level standardized test performance is converted so that measures may be compared nationally.

- *Census Shapefiles, 2010* are a census product describing geographic boundaries of administrative units including school districts and census tracts.

2.5.2 Sample Construction

I construct the following samples for analysis:

- The *FNMA Sample* is the primary sample in the analysis. It consists of all 30-year, fixed rate, conventional, conforming, first-lien mortgages originated between 2005 and 2020 for the purchase of single family, owner-occupied, 1-4 unit dwellings that were subsequently sold to Fannie Mae for securitization. I describe summary statistics of the FNMA Sample in Table (2.1). The characteristics of these mortgages are standard.

- I construct the *HMDA-FNMA-SEDA Merge* by (i) merging HMDA records to information on school district quality in SEDA (ii) merging FNMA records to information on FNMA LLPA grids and Essent rate cards and (iii) merging HMDA and FNMA records. My approach is similar to the approach in Bartlett et al. [2021] and uses the same data sources as the more recent Buchak and Jørring [2021]. Below, I describe the sample contents and construction more closely.

   i. I subset HMDA records to retain only vanilla (non-negatively amortizing, non-balloon, non-interest only), conventional, conforming, fixed-rate, 30 year, first mortgages originated in 2018 or 2019 for purchase of single-family, 1-4 unit, owner-occupied housing. I use the 2010 Census Shapefiles to assign the Census Tract of a transaction in HMDA to a school district. I use SEDA to obtain district characteristics, including school performance.

   ii. I subset FNMA records in the same manner. From the FNMA LLPA grids, I record the LLPA rate for FNMA insured loans according to the date, loan LTV, and borrower FICO at origination. From the Essent rate cards, I record monthly premiums for borrower PMI according to the date, borrower LTV and FICO at origination. Using the loan origination month, LTV, and borrower FICO, I merge these to the subset of FNMA records.

iii. To merge records, I require that loans share an identical state, MSA, 3-digit zip code, origination year, originator if available, and debt-to-income bin. I then conduct a fuzzy merge using the original balance, loan-to-value ratio, and rate at origination. The quality of the fuzzy merge along merge variables is depicted in Figure (A.2). The merge rate and representativeness of the merged sample is depicted in Table (2.2).

- The *CRISM Sample* consists of a 7% random sample of loans in CRISM that are vanilla (non-IO and non-Balloon), conventional, conforming, 30-year, fixed-rate, first mortgages originated between 2005 and 2020 against 1-4 unit, single-family, owner-occupied housing. I discard a handful of loans that are not onboarded promptly to CRISM or that have unreasonably high loan-to-value ratios. I describe summary statistics for the CRISM Sample in Tables (A.1) and (A.2).

Table 2.1: Summary Statistics for the Fannie Mae Sample†

	N	$\overline{x}$	s_x	min	p25	p50	p75	max
Loan Characteristics:								
Loan Amount ($k)	52.6k	235	126	10	140	209	307	794
Property Value ($k)	52.6k	0.29	0.18	0.02	0.17	0.25	0.37	2.61
Interest Rate (%)	52.6k	4.41	0.87	2	3.88	4.25	4.88	8.62
LTV (%)	52.6k	83.6	13.1	5	80	85	95	97
DTI (%)	52.4k	35.1	9.55	2	28	36	42	64
FICO Score	52.5k	751	46.1	503	721	761	789	838
PMI Characteristics:								
Has PMI	52.6k	0.5	0.5	0	0	1	1	1
Has Borrower PMI	52.6k	0.47	0.5	0	0	0	1	1

The Fannie Mae Sample consists of loans from the Fannie Mae Single Family Loan Performance dataset with the following properties: originated 2005-2020, purchase loans, first liens, single-family, owner-occupied, 1-4 unit dwelling, conventional, conforming, fixed-rate, and 30 year terms. For computational ease, the summary statistics above are computed on a 1pp random sub-sample of the Fannie Mae Sample.

Table 2.2: Summary Statistics for HMDA-FNMA-SEDA Merge

	FNMA Sample† N=10066	HMDA Sample† N=27478	Merge† N=3523
Loan Characteristics:			
Loan Amount ($k)	245 (126)	261 (130)	225 (107)
Property Value ($k)	. (.)	322 (181)	328 (3538)
Interest Rate (%)	4.54 (0.53)	4.50 (0.59)	4.50 (0.51)
Total Costs ($k)	. (.)	4.21 (4.01)	3.48 (2.04)
Rate Spread (%)	. (.)	0.49 (0.58)	0.50 (0.49)
LTV (%)	86.2 (12.7)	84.5 (14.1)	86.8 (12.3)
DTI (%)	36.6 (9.05)	35.7 (9.44)	34.7 (9.09)
Has PMI: Yes	62.1%	.%	63.5%
PMI Coverage	26.2 (5.79)	. (.)	26.4 (5.58)
Purchaser Type:			
Balance Sheet	0.00%	49.3%	19.4%
FNMA	100%	50.7%	69.4%
Borrower Characteristics:			
Annual Income ($k)	. (.)	160 (2661)	108 (1059)
FICO Score	748 (45.1)	. (.)	750 (43.8)
Age (y)	. (.)	41.7 (13.7)	39.9 (13.5)
Ethnicity: Hispanic or Latino	.%	13.4%	11.3%
Race:			
Asian	.%	7.37%	3.96%
Black or African American	.%	5.10%	4.43%
White	.%	84.7%	89.6%
Sex:			
Female	.%	25.2%	25.1%
Joint	.%	41.1%	40.9%
Male	.%	33.7%	34.0%
First Home: Yes	55.6%	.%	56.7%
Tract Characteristics:			
Population (k)	. (.)	5.41 (3.63)	5.65 (3.31)
Minority Population (%)	. (.)	27.3 (23.9)	25.8 (22.0)
Median Income ($k)	. (.)	82.9 (41.4)	82.5 (28.3)
1-4 Unit Homeownership Rate (%)	. (.)	74.3 (14.8)	73.3 (14.4)
Median Age of Housing Units (y)	. (.)	31.6 (19.7)	37.2 (17.3)
District Characteristics:			
Enrollment (k)	. (.)	23.5 (52.0)	13.9 (20.6)
Minority Population (%)	. (.)	40.9 (27.0)	37.8 (25.9)
Median Income ($k)	. (.)	62.3 (19.2)	58.8 (15.4)
Test Performance (z)	. (.)	0.07 (0.31)	0.04 (0.31)
Segregation (W/B)	. (.)	0.13 (0.15)	0.12 (0.12)

continued on next page

Table 2.2 – *continued from previous page*

	FNMA Sample N=10066	HMDA Sample N=27478	MSAMP0 N=3523

I merge data from the Home Mortgage Disclosure Act (HMDA) Loan Application Record, the Fannie Mae (FNMA) Single Family Loan Performance data, and the Stanford Education Data Archive (SEDA). The SEDA data are merged to HMDA prior to the merge summarized above. The samples from HMDA and FNMA include loans with the following characteristics: originated 2018-2019, purchase loans, first liens, single-family, owner-occupied, 1-4 unit dwelling, conventional, conforming, vanilla (no balloon payments, no interest only payments, no negatively amortizing loans, no other non-amortizing features, no reverse mortgages), fixed-rate, and 30 year terms. For computational ease, the summary statistics above are computed on a 1pp random sub-sample of the FNMA, HMDA, and merged data.

2.6 Results I: The Average Shadow Price of Credit

2.6.1 The Mortgage Price Schedule

I construct a representation of the mortgage price schedule from data on mortgage originations in the *FNMA Sample*. Having linked these data to *Essent Rate Cards* using origination date, borrower FICO, and loan LTV, I can observe the following in the data: annual mortgage interest rate, r_i^m, minimum required annual PMI rate, r_i^{pmi}, home price at origination, P_i^o, and initial balance, B_i^o. I also choose a time-varying discount factor, r_t. I use the 30-year treasury yield at loan origination as a measure of a risk-free rate at the time horizon of the mortgage contract tenor. This choice is merely illustrative and I consider alternatives in the sensitivity analysis below.

For mortgage i originated at time t, I construct the present discounted value of mortgage obligations, PDV_i:

$$PDV_i = \sum_{\tau=1}^{360} \frac{m_i + pmi_{i,t+\tau}}{(1 + r_t/12)^{\tau}} \tag{2.18}$$

These obligations are a function of monthly mortgage payments, m_i, and monthly mortgage insurance payments, $pmi_{i,t+\tau}$. Monthly mortgage payments are computed according to the fixed-rate mortgage amortization formula. PMI is no longer required when borrowers' current LTV reaches 80, so PMI payments are assumed to terminate when the outstanding principle reaches 80% of the original home price:

$$m_i = r_i^m/12 * B_i^o * \frac{1}{1 - (1 + r_i^m/12)^{-360}} \tag{2.19}$$

$$pmi_{i,t+\tau} = r_{it}^{pmi}/12 * B_i^o * \mathbb{1}\left\{B_{i,t+\tau}/P_i^o > 0.80\right\} \tag{2.20}$$

Adjustments to the mortgage and PMI rates due to the LLPA and rate card grids are quoted relative to the loan-to-value ratio. This means that the menu of payment obligations is increasing in the mortgage balance only after fixing the price of the home. To ensure that I am depicting household locational choices within a given menu and not comparing borrowers along different menus, I bin the data according to home price percentiles. I plot a bin-scatter of mortgage payment obligations, PDV_i, against initial balances, B_i^o, within home price bin.

I present the results of this exercise for the 10th, 30th, 50th, 70th, and 90th quantile of the home price distribution in Figure (2.8). The bin-scatter has two salient features, a "backbone" and a series of "tails". The "backbone" consists of the lower portions of each mortgage price schedule, which are approximately collinear across home price bins. The presence of the backbone is consistent with a constant price of mortgage credit, so that marginal dollars are no more expensive. The "tails" consist of the upper portions of each mortgage price schedule, which are convex and slope more steeply upward and away from the "backbone". The tails are consistent with the increasing mortgage and PMI rates as the loan-to-value on the mortgage increases.

Section (2.3.1) suggests the interpretation that the shadow price of credit can be read off of the slope of these mortgage price schedules - the steeper the slope, and the more mass on steeper portions of the schedule, the higher is households' shadow price of credit. Loosely, two factors determine the steepness of the curves: the slope of the "backbone" and the convexity of the "tails". In a frictionless world, the present discounted value of mortgage obligations must always be exactly equal to the mortgage balance for any contract to be freely entered. The slope of the "backbone" would be unity and the convexity of the "tails" would be nil and the figure would trace a 45 degree line.

Mathematically, the slope of the "backbone" is influenced to a great extent by the choice of the discount factor. Economically, this discount factor should capture the relevant time horizon of the contract and account for the option value of the pre-payment option embedded in the mortgage contract. The convexity of the "tails" is influenced by the payments the borrower anticipates making relative to their mortgage obligations. Economically, to the extent that borrowers anticipate defaulting on mortgage obligations in certain bad states of the world, this may reduce the convexity of the curves.

Figure 2.8: Empirical Mortgage Price Schedule

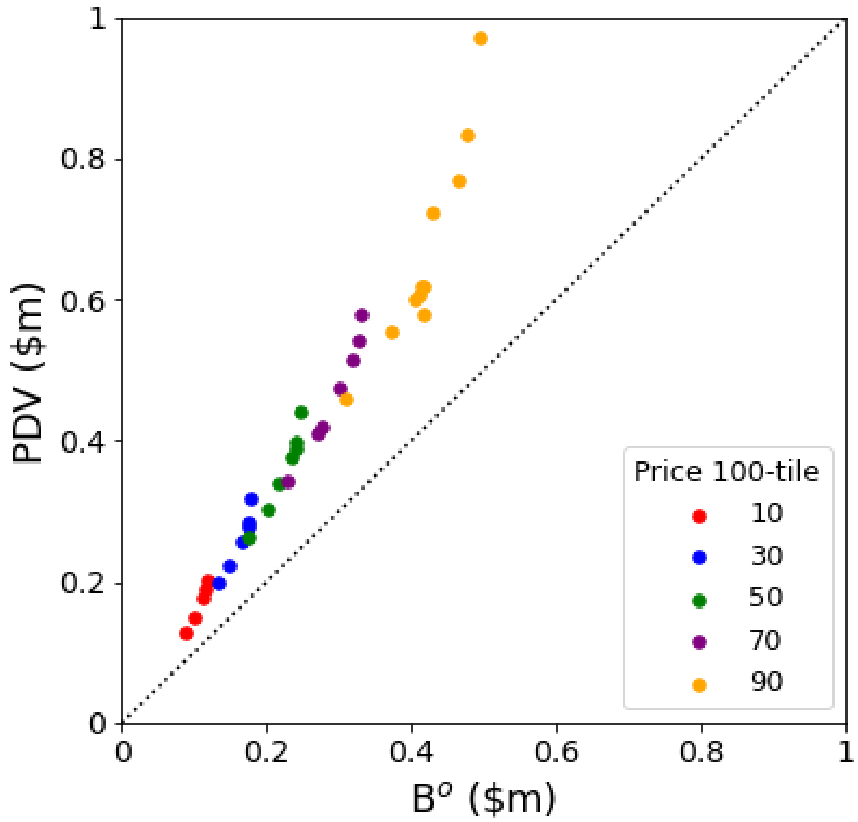

This binscatter shows the value of mortgage obligations relative to the initial mortgage balance at different percentiles of the home purchase price distribution. Within each price percentile, the binscatter traces a convex menu, consistent with the interest rates that increase in loan-to-value. The lower portions of these convex menus are approximately collinear across price quantiles. The slope of these menus, where the borrower chooses to locate, corresponds to borrowers' willingness-to-pay for credit.

2.6.2 Sensitivity Analysis of Average WTP for Credit

My approach to measuring the average willingness-to-pay for credit at the time of home purchase is to measure the average slope of the mortgage price schedule, weighted by the frequency of mortgages in the data appearing along it. My tack is to take adjacent points within home price quantile on the binscatter and find the slope of the line between the two.

I then assign this slope to all mortgages in that price quantile with balances between those two points. Having assigned an approximate slope to each mortgage, I take a simple average and interpret it as willingness to pay for credit.

This procedure would accurately capture willingness-to-pay in the simple model I've written down to motivate the analysis. Capturing plausible estimates of willingness-to-pay for credit outside the model requires addressing additional factors in mortgage pricing. Two of the most salient are the default and prepayment options. While I do not build a sophisticated pricing model to account for these factors, I aim to address them by conducting a sensitivity analysis.

My sensitivity analysis proceeds as follows:

1. I choose a recovery rate on mortgage payment obligations, θ, and a discount rate, r_t

    i. I construct the present discounted value of mortgage obligations as:

    $$PDV_i = \theta \sum_{\tau=1}^{360} \frac{m_i + pmi_{i,t+\tau}}{(1 + r_t/12)^{\tau}} \tag{2.21}$$

    Where m_i and $pmi_{i,t+\tau}$ are defined as before.

    ii. For some p and q, I bin the data by p-tiles in home price, P_i^o, and q-tiles in initial balance, B_i^o

        I. For each p-tile, I compute local slopes between adjacent points in a bin-scatter.

        II. I compute the average slope as the observation-weighted average local slope.

    iii. I search for values of p and q where estimates of the average slope stabilize. (In particular, I use 100-tiles in P_i^o and 30-tiles in B_i^o.)

    iv. My estimate for a given (θ, r_t) pair is the estimate from this stabilized cell.

2. I report estimates for each (θ, r_t) pair.

I choose the following values for the recovery rate, θ, and discount rate, r_t. I use the 10-year treasury yield, $y_t^{(10)}$, to capture the common realized term of the contract given events of repayment and default. I use the 30-year treasury yield, $y_t^{(30)}$, to capture the term written on the contract. And I use an ad-hoc discount rate, r^*, which I compute at the 30-year treasury yield plus a spread. I choose the spread so that the "backbone" of the mortgage price schedule has a slope of ~1. Thus, I conservatively assume that prepayment risk in mortgages is perfectly priced.

I consider a recovery rate of 1 and 0.99. The recovery rate of 0.99 reflects the fact that 99% of outstanding balances are recovered on average, i.e. a foreclosure rate of 2% and a loss given default of 40% [An and Cordell, 2021]. The recovery rate of 1 is counter-factual but underscores the fact that borrower impressions of their loan is what matters. If borrowers anticipate that they will repay 100% of their balances at the time of mortgage origination, then it may be more accurate to be somewhat less conservative.

I report results of this analysis in Table (2.3). For each combination of parameter choices, I report means and, in parentheses, standard deviations of the local slopes. The choice of discount factor tends to move the estimates by a greater extent than the choice of the recovery rate. The most conservative estimate, with discount rates at $r_t^* = y_t^{(30)} + 1.5\%$ and a recovery rate of $\theta = 0.99$ results in an average slope of 1.65. This suggests that, to a first degree of approximation, hedonic estimates may be downward biased by as much as 65%.

2.7 Results II: Correcting Hedonic Estimates

2.7.1 Distribution of Shadow Prices

I begin by constructing hypothetical mortgage contract menus facing borrowers at the time of origination. To do so, I am guided by the presentation of institutional details in Section (2.4.2) and evidence in Section (2.4.3). I assume that, given the price of the home, each

Table 2.3: Sensitivity Analysis of Credit Wedge, $\kappa = 1 + \frac{\mu}{\lambda}$

Discount rate, r_t	Recovery rate, θ	
	1	0.99
$y_t^{(10)}$	2.017 (4.73)	1.996 (4.68)
$y_t^{(30)}$	1.996 (3.05)	1.976 (3.02)
r_t^*	1.665 (2.36)	1.648 (2.34)

Credit wedges are computed as the local slopes in binscatters of mortgage payment obligations against origination balances. The computational procedure is described in the text. For these results, the data are binned by 100-tiles in initial home price, P_i^o, and within this, 30-tiles in initial mortgage balance, B_i^o. Per equation (2.21), the value of mortgage obligations is computed as $\theta \sum_{\tau=1}^{360} \frac{m_i + pmi_{i,t+\tau}}{(1+r_t/12)^\tau}$, where θ is the fraction of payments borrowers expect to make and r_t is the choice of discount factor on obligations.

borrower faces a menu of contracts comprised of initial balances and interest rates. The initial balances are determined by the loan-to-value thresholds used by the GSEs and the changes in the interest rate are pinned down by the LLPA adjustments and PMI rates. The level of the interest rates is pinned down by the contract chosen by the borrower that I can observe in the data.

This notion of a borrower-level menu is depicted in Figure (2.9). The borrower's choice of contract, observable in the data, is depicted in red. Two alternative contracts are constructed and depicted in blue. These represent what a borrower might plausibly have faced if they levered up or down.

I use these menus to estimate borrowers' shadow price of credit, i.e. their willingness-to-pay for an additional dollar of funds today out of funds tomorrow. The intuition of this exercise is to convert the menu of rates into a menu of payment obligations, as depicted in Figure (2.3), and find the slope at the point where the borrower locates. In practice, I consider a menu of finite options. I formalize the borrower's choice among these finite options to derive bounds on the borrower WTP for credit in Appendix (A.1.2). From Equation (A.8), we have the following, where j^* is the chosen mortgage contract and j^*-1 and j^*+1 are adjacent non-chosen alternatives.

$$\frac{r^m_{j^*-1}}{r}\left[1+\frac{\%\Delta^{j^*}_{j^*-1}r^m}{\%\Delta^{j^*}_{j^*-1}B^o}\right] \leq \kappa_i \leq \frac{r^m_{j^*}}{r}\left[1+\frac{\%\Delta^{j^*+1}_{j^*}r^m}{\%\Delta^{j^*+1}_{j^*}B^o}\right] \tag{2.22}$$

Figure 2.9: Imputation of borrower-level mortgage menus

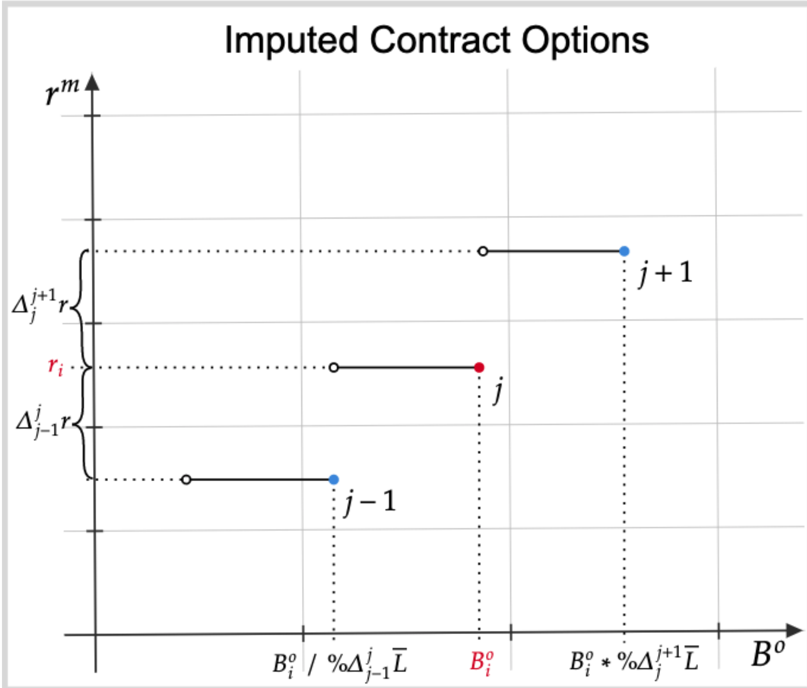

I consider available alternative mortgage contracts for borrower i (red) using the borrower's chosen mortgage rate, r_i, and initial balance, B_i^o, the Fannie Mae LLPA grids, and Essent rate card grids. Considering the borrower's mortgage as contract j, I construct more and less levered alternatives, $j+1$ and $j-1$ respectively:

- The alternative available balance is the observed balance, B_i^o, scaled by the change in the maximum allowable leverage for the loan to qualify under a different grid cell, $\%\Delta_j^{j\pm1}\overline{L}$.
- The alternative contract rate is the observed rate, r_i, plus a component that captures the costs of covering LLPA fees or additional PMI required to remain conforming, $\Delta_j^{j\pm1}r$.

Some assumptions are required to implement this analysis, which I consider here.

- I can only obtain historical mortgage rate cards from Essent, so I must assume that these rates are fairly representative. Table (A.1) reports that Essent provides PMI for about 5% of loans in the mortgage market suggesting that they are a large but certainly not dominant player. In inspect rate cards from other insurers in more recent years, i.e. 2019-, and find that both insurance rates and the matrix structure of the cards are comparable.

- Similarly, I can only observe interest rates for borrower-paid mortgage insurance, so it must be the case that this is a common way of structuring the contract. Table (2.1) confirms this. Note that nearly half of borrowers have PMI and nearly all borrowers who have PMI have "borrower-paid" PMI, the standard form of insurance.

- Converting the LLPA fee reported in the FNMA grids to interest rate increases also requires assumptions. Technically, the grids contain up-front premiums assessed to the originator. Originators commonly convert these up-front obligations to a stream of payment obligations and then assess borrowers a higher rate to cover the additional expenses. The conversion multiple by which lenders convert the up-front premium into a g-fee are difficult to collect. Inspecting one reveals that the buy-up multiple fluctuates around 5-10. I use a factor of 10 to convert the LLPA fees to estimated interest rate rises faced by borrowers. I aim to err on the conservative side, suggesting that there is a fairly small rise in the interest rate faced by borrowers as they lever up. A smaller interest rate rise will correspond with lower estimated willingness to pay for credit.

- Again, I use the 30-year Treasury yield for the risk-free rate. This is motivated by (i) considering the borrower's alternative margin to be priced by a savings rate and (ii) the notion that the service quality is invariant to the state of the market and therefore that the relevant discount rate for pricing the value of these services is a risk-free rate. However, the value of mortgage obligations may be affected by taxes or default or prepayment propensity, and discounting at the 30-year Treasury rate does not account for these.

I implement the analysis suggested above to obtain lower and upper bounds on the shadow values of additional borrowing in this simplified framework. This distribution of lower bounds is plotted in Figure (2.10). The x-axis should be interpreted as willingness to pay in dollars to relax the borrowing constraint by a dollar. Households appear to be considerably credit constrained and there appears to be a high degree of heterogeneity, with a variance in the estimates of about 0.5.

The borrowers populating the CDF loosely increase in LTV along the x axis from left to right. This corresponds with the convexity of the mortgage-offer curve. The high upper bounds on the willingness-to-pay for credit, whose values fall between 6 and 8, come from the borrowers with LTVs of 97. Note that the increase in the rate for PMI at 97 is considerable and that these borrowers tend to have high interest rates on their mortgage already, both of which contribute to the high estimated upper bound.

2.7.2 Corrected Hedonic Estimates

I consider a cross-section of households, i, each purchasing a home at some date, $t(i)$, and located in a census tract, $k(i)$, school district, $d(i)$, and county, $c(i)$. I run a standard hedonic regression of home prices, $\ln P_{id}$, on measures of school quality at the district level, s_d, and other amenity measures at the census tract level including median income, rate of owner-occupancy, and minority percentage, $X_{k(i)}$. Finally, I include fixed effects for the county, $\alpha_{c(i)}$, and month of home purchase, $\alpha_{t(i)}$. The regression specification is below:

$$\ln P_{id} \sim \tilde{\alpha}_{c(i)} + \tilde{\alpha}_{t(i)} + \tilde{\beta} s_d + \tilde{\gamma}' X_{k(i)} + \varepsilon_{id} \tag{2.23}$$

The interpretation of the hedonic coefficient, as amended by the analysis in this paper, is the cross-sectional mean marginal willingness-to-pay for the amenity, downward biased by the credit constraint wedge:

$$\tilde{\beta} = \mathbb{E}^i \left[\frac{u_s^i / u_c^i}{r} \frac{1}{\kappa_i} \right] \tag{2.24}$$

The object of policy relevance, provided that government borrowing and taxation can be used to smooth the incidence of costs, is simply the marginal willingness-to-pay for the amenity,

Figure 2.10: Distribution of Credit Wedges

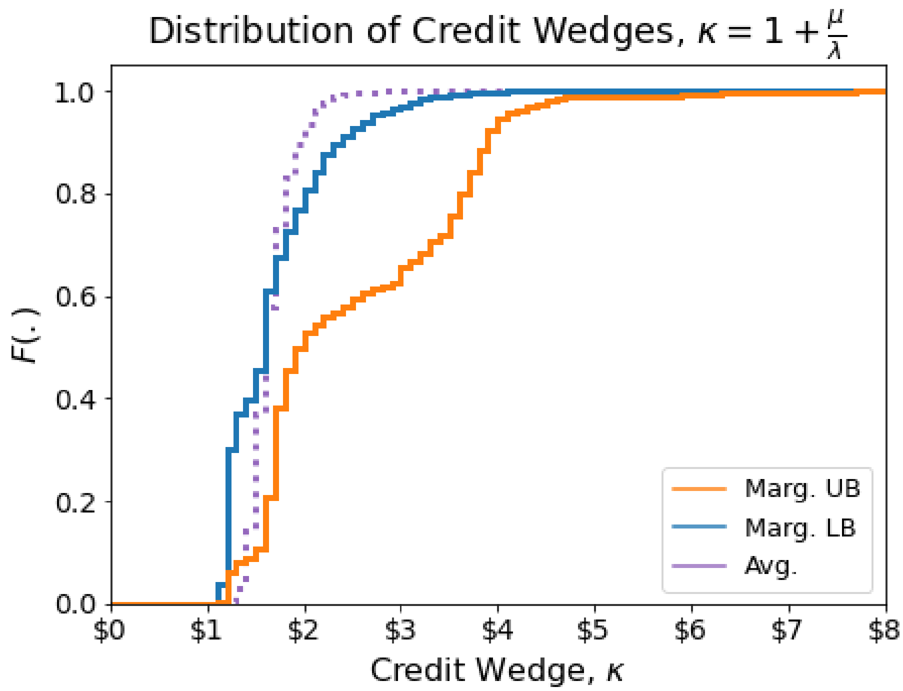

This plot depicts the estimated distribution of credit wedges, κ, for borrowers in the HMDA-FNMA-SEDA merged sample. The blue and orange lines depict the CDFs of upper and lower bounds on the credit wedge, per Equation (2.22). Note that this formulation captures the borrower's WTP for the *marginal* dollar of credit. As a benchmark, the dotted purple line depicts the distribution of credit wedges assuming that the borrower's WTP for the *marginal* and *average* dollar of credit are the same. In particular, I compute a ratio of the borrower's effective interest rate to the 30-year treasury yield, $(r_i^m + r_i^{pmi})/y_t^{(30)}$. The WTP for the *average* dollar of credit understates the extent of credit constraints in roughly half the population.

$\mathbb{E}^i\left[\frac{u_s^i/u_c^i}{r^f}\right]$.

Because of the possibility of cross-sectional heterogeneity, correcting the bias in the estimate is not as simple as finding the average credit wedge. Using a covariance decomposition, the policy-relevant estimate can be shown to be:

$$E^i\left[\frac{u_s^i/u_c^i}{r}\right] = \left(E^i\left[\frac{1}{\kappa_i}\right]\right)^{-1}\left(E^i\left[\frac{u_s^i/u_c^i}{r}\frac{1}{\kappa_i}\right] + Cov^i\left(\frac{u_s^i/u_c^i}{r}, \frac{1}{\kappa_i}\right)\right) \tag{2.25}$$

Beyond estimating the credit wedge, it is necessary to know something about the population covariance between the marginal willingness-to-pay for amenities and credit in equilibrium.

My tack for correcting hedonic estimates is to exploit the fact that, having constructed menus facing borrowers, it is possible to estimate willingness-to-pay for credit at the borrower level. Effectively, I bin borrowers by their willingness-to-pay for credit, estimate hedonic coefficients and take a weighted average. In Section (A.2.2), I sketch the sense in which a single fixed-effects regression captures the object of interest.

Returning to the same cross section of households, I introduce a borrower's credit wedge, κ_i, as well as their credit wedge bin, $j(i)$. I normalize the school quality measures by the borrower credit wedge, $\frac{s_d}{\kappa_i}$, and I include fixed effects for the borrower's credit wedge bin, $\alpha_{j(i)}$.

$$\ln P_{id} \sim \alpha_{j(i)} + \alpha_{c(i)} + \alpha_{t(i)} + \beta\frac{s_d}{\kappa_i} + \gamma' X_{k(i)} + \varepsilon_{id} \tag{2.26}$$

The interpretation of this corrected regression coefficient, is the cross-sectional mean marginal willingness-to-pay for the amenity. This is the value of policy interest.

$$\beta = \mathbb{E}^i\left[\frac{u_s^i/u_c^i}{r}\right] \tag{2.27}$$

Table (2.4) presents the results of this regression exercise. In the first two columns, I conduct the exercise with limited controls and in the latter two I introduce more tract-level controls. In columns two and four I run the constraint correction specification. I cluster

Table 2.4: 'Corrected' Hedonic Regressions

Dependent Variable:	ln Home Price$_{id}$			
	(1)	(2)	(3)	(4)
Variables				
Test Score$_d$	0.442***		0.156***	
	(0.017)		(0.013)	
Adj. Test Score$_{id}$		0.666***		0.230***
		(0.028)		(0.020)
Fixed-effects				
County FE	✓	✓	✓	✓
Month FE	✓	✓	✓	✓
Shadow Price FE		✓		✓
% Minority FE			✓	✓
q50 Income FE			✓	✓
% Ownership FE			✓	✓
Fit statistics				
Observations	382,407	382,407	382,407	382,407
R^2	0.461	0.474	0.554	0.566
Within R^2	0.065	0.049	0.009	0.006

Clustered (County) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

at the county level. In both cases, the estimates of mean marginal willingness-to-pay for amenities increases by ~ 50%.

As a caveat, traditional hedonic specifications are known not to be well identified, for one, because of unobserved covariates. Limitations on my ability to merge data prevent me from instrumenting in the manner most common in school quality applications, border discontinuities at within-district catchment area borders (due to Black [1999]). Instrumenting in this way tends to lower the estimate of willingness-to-pay for school quality. The exercise presented here is intended to suggest the effect of correcting for credit constraints on estimates of amenity value.

2.8 Conclusion

In this paper, I argue that frictions in household credit markets may systematically bias estimates of local public amenities derived from traditional hedonic regressions. Credit constrained households bid less for marginal amenities than they would in the absence of such constraints, flattening the hedonic price schedule of amenities. I also argue that the bias in hedonic regressions, which is related to the household's shadow price of credit, may be estimated from households' choice of mortgage product, because of the increasing costs of a marginal dollar of borrowing. I estimate the "credit wedge", or household willingness-to-pay for a marginal dollar of borrowing, in the US mortgage market. And I propose and (on a preliminary basis) implement a strategy for correcting bias in hedonic estimates that accounts for the cross-section heterogeneity.

CHAPTER 3

MORTGAGE LENDING LIMITS AND HOUSING DEMAND: EVIDENCE FROM BUNCHING IN FHA BORROWING

3.1 Introduction

Mortgage credit supply and expectations of future house price growth are prominent among explanations of the housing boom and bust in the 2000s [Mian and Sufi, 2009, Kaplan et al., 2020]. Under the credit supply story, households with newly relaxed credit-constraints demanded more housing services and, consequently, more housing. Under the expectations story, household demand for housing services did not change, per se, but demand for housing increased because of a pervasive and counter-factual belief in high returns on housing as an asset class.

In principle, both channels may have contributed to the housing cycle. Empirically determining the relative contribution of each is fraught, however, because the channels are potentially interrelated: changing aggregate credit conditions could affect rational or irrational beliefs and vice versa. In particular, the credit supply story has been most rigorously established by instrumental variable approaches that identify the effect of increased credit supply on local housing price indices in the cross-section [Favara and Imbs, 2015, Johnson, 2019]. But this approach captures both a direct effect on demand for housing services as well as a knock-on effect of demand for housing due to changing beliefs.

In this paper, I develop and employ a different tack to better establish the relationship between credit supply and demand for housing by homeowners, the intuition of which I explain here. Households face two primary constraints when borrowing through major institutional channels to finance a home purchase: the loan-to-value limit, $\overline{L}$, and initial loan balance limit, $\overline{B^o}$, implied, under fairly weak assumptions, by an institutional debt-to-income limit. For the first \$1 of housing, and for every \$1 a household demands before reaching the initial loan balance limit, $\overline{B^o}$, the household must pay $\$(1-\overline{L})$ more in down-payment. For the first \$1 of housing after reaching $\overline{B^o}$, and for every \$1 thereafter, the household must pay \$1 more in down-payment.

If, outside these institutional lending programs, households are not constrained, and can substitute without cost to another line of credit, e.g. a second mortgage, credit card, or personal loan, then this discontinuity in marginal down-payment requirements should not affect their demand for housing. In fact, by borrowing from other lines of credit, households can eliminate this programmatic discontinuity. For example, having reached the initial loan balance limit, $\overline{B^o}$, to finance an additional \$1 of housing, households could simply make the \$1 down-payment required by the lender, and borrow $\$\overline{L}$ on another line of credit, for a net marginal down-payment at the time of purchase of $\$(1-\overline{L})$.

However, if households are elsewhere constrained, e.g. when home purchase may be financed only with savings and a mortgage loan obtained from a major institutional channel, the lending limits ensure discontinuously higher marginal down-payment and therefore a discontinuously higher utility cost associated with an additional dollar of housing. As a result, households for whom the marginal benefit of additional housing falls between the utility costs associated with a marginal downpayment of $\$(1-\overline{L})$ and \$1 would demand additional housing until they reached this discontinuity in marginal down-payment requirements. The effect of changing down-payment requirements can therefore be observed in the data as the extent of bunching at the down-payment amount, D^*, associated with this marginal down-payment discontinuity. In a later section, I formalize and clarify the interpretation of this bunching behavior.

This approach to analyzing the link between demand for housing and credit availability has several advantages over those in the literature. Chief among these, it is implausible that identified behavioral responses in this setting are the result of changing expectations. For one, this strategy does not rely on aggregate (e.g. zip code or county-level) treatment, which might affect home-buyer beliefs through its effect on other households or investors. Moreover, even at the individual level, the information treatment and credit constraint quasitreatment are disentangled. Borrowers may update their beliefs about housing markets as they come to learn the rules that govern their own and others' borrowing limits, but this change in beliefs occurs separately from the selection of a loan from a kinked opportunity set. These features of the empirical design allow for isolation of the effects of credit constraints on housing demand, holding constant households' beliefs about the market.

There are ancillary benefits to this approach as well. For one, the fact that bunching is measured from individual-level outcomes, it is possible to look at different segments of the borrower population to understand the extent to which changing mortgage borrowing constraints influences their demand for housing. This is productive in light of research suggesting that capital gains on housing during the early 2000s varied in the cross-section of housing quality [Landvoigt et al., 2015]. Second, this strategy isolates effects of borrowing requirements on housing demand, a real outcome, rather than first mortgage balances, a financial outcome which may be offset by borrowing along additional margins [DeFusco and Paciorek, 2017].

This bunching design is not without limitations. Primarily, there is a challenge in comparing the results of this paper, which quote the effect of credit supply on individual quantities of housing demanded, to the existing literature, which quotes the effect of credit supply on equilibrium prices. To make this comparison requires two steps, each of which requires additional assumptions. First, it is necessary to aggregate changes in housing demand resulting from changing credit conditions in the early 2000s. Second it is necessary to map from aggregate change in quantities demanded to equilibrium price changes. Additionally, the analysis requires merging, at the individual level, household incomes, back-end debt-to-income ratios, loan balances, initial interest rates, transaction prices, as well as information on their county of residence and up-front mortgage insurance premia. My ability to do this in the data, at present, extends only to the years 2018-2019 and only for FHA loans, for which subsample both HMDA and McDash provide enough fields to merge the required variables.

In Section (3.2), I consider in closer detail how my analysis builds on two robust literatures, housing demand in the 2000s and bunching estimators. In Section (3.3), I provide context for my analysis by describing the history and features of the FHA lending program; I describe how I use the features of this program to construct the loan-to-value limit and the initial loan balance limit for each borrower. In Section (3.4), I describe my empirical approach by sketching the household problem being examined (3.4.1), how observable quantities can be used to estimate the relevant elasticity (3.4.2), and how to interpret this elasticity (3.4.3). In Section (3.5), I describe the way I construct a sample from available data that suits the needs of my estimator (3.5.1) and technical details of implementing the bunching estimator (3.5.2). Section (3.6) presents the results of the analysis and Section (3.7) concludes.

3.2 Related Literature

3.2.1 Beliefs and Credit Constraints

Since the housing crisis of the 2000s, a central line of inquiry has sought to understand the determinants of housing demand and the housing cycle. Early research substantiated the role of credit supply in generating both the boom and bust in housing prices. By disaggregating to the level of the zip code, Mian and Sufi [2009] note that areas with high mortgage credit growth in the boom years tended to have relatively low income growth, suggesting that credit supply rather than demand drove the rise in household leverage and house prices. The credit supply story has theoretical as well as empirical appeal, as it plays an important role in building a cohesive narrative of the major facts surrounding the housing cycle in the early 2000s. Justiniano et al. [2019] observes that for increased leverage to be consistent with falling interest rates, as was true in the early 2000s, credit supply must have been expanding rather than credit demand.

Other scholarship has pursued the role of household expectations; under this hypothesis, home prices increased because optimistic beliefs about the path of future home prices in the future generated demand for housing access in the present. Kaplan et al. [2020] argues that, though credit supply helps to explain changes in homeownership (see also, Acolin et al. [2016]) and default rates during the crisis, it provides little explanatory power for house prices. Constrained households may obtain equivalent housing services from the rental market, they argue, such that relaxed credit constraints do not generate increased demand for housing units. Rather, they attribute the boom to a counter-factual belief that housing would appreciate in value. Mian and Sufi [2018] studies home purchases by investors and notes that investors' optimistic beliefs drive their purchasing behavior and therefore home prices during the boom.

A central challenge in empirical scholarship on this topic is disentangling beliefs from credit constraints, which are easily confounded: changing aggregate credit conditions could affect rational or irrational beliefs and vice versa. The credit supply story has been most rigorously established by instrumental variable approaches that identify the effect of increased credit supply on local housing price indices in the cross-section [Favara and Imbs, 2015, Johnson, 2019]. But this approach captures both a direct effect on demand for housing services as well as a knock-on effect of demand for housing due to changing beliefs.

Some empirical work has attempted to disentangle credit access from beliefs in generating housing demand. Empirical tests to rule out the beliefs story have appealed to regions with high housing supply elasticity in which a belief in long-run price fluctuations might be unreasonable [Mian and Sufi, 2009]. But because these exercises have not been able to directly measure beliefs, it is more difficult to control for unreasonable beliefs. Bailey et al. [2019] use social networks to identify shocks to beliefs about housing values. They argue that changes in beliefs do not generate enough change in housing demand to account for the housing cycle of the early 2000s. They suggest that credit access may play a role instead but their research design does not enable them to measure this in the data.

Finally, some work has examined bunching in data at loan-limits, an approach which, though not the studies' explicit intent, is more persuasive in controlling for beliefs. Anenberg et al. [2017] looks for a mass point at the "loan frontier", an empirically imputed maximum loan balance available to a borrower of a given credit-score, down-payment, and income. This mass point is a poor test of credit constraints, though, as down-payments are chosen contemporaneously with home purchase. DeFusco and Paciorek [2017] identifies bunching at the jumbo loan threshold where interest rates increase, suggesting that first mortgage demand is responsive to interest rates. This measures a different elasticity altogether, though, and the link from mortgage demand to housing demand is unclear because of the possibility of putting other lines of credit to use.

3.2.2 Bunching Estimators

There is widespread use of bunching estimators in the economics literature to identify elasticities in the data. An older literature considers information content available in kinked opportunity sets, but the availability of administrative data set off the present use of bunching estimators, first formalized by Saez [2010]. Kleven and Waseem [2013] provides a theoretical analysis of notches and relates the empirically estimated elasticities to structural elasticities of interest. These estimators have been used in a wide variety of "real" contexts [Yelowitz, 1995, Sallee and Slemrod, 2012, Ramnath, 2013, Manoli and Weber, 2016, Blundell and Shephard, 2012, Blundell and Hoynes, 2004, Kleven, 2016].

The introduction of bunching estimators to the finance literature is somewhat more recent. DeFusco and Paciorek [2017] estimates first mortgage borrowing elasticity relative to the interest rate by exploiting the interest rate variation at the GSE conforming loan limit. Best et al. [2018] estimates the elasticity of inter-temporal substitution using interest rate variation at particular LTV ratios in the British mortgage market. Dagostino [2019] estimates the elasticity of municipal bond issuance size using changes in available yields due to bank qualification status. Bachas et al. [2020] estimate the elasticity of loan size using a discrete change in the loan guarantee rate. Ebrahimian [2020] uses bunching as a moment to identify a structural IO model of student debt.

Relative to this literature, I exploit policy parameters that generate a kink from borrowing rather than budget constraints. A households' additional costs enter through increasing down-payment requirements at the kink point rather than increasing interest expense. My analysis explains how it is possible to use a bunching estimator to estimate and interpret a financial elasticity of interest, the loan-to-value elasticity of housing demand.

3.2.3 FHA Lending

There is a small literature on FHA lending, particularly in the aftermath of the financial crisis, during which it served as an important source of mortgage lending as the private market contracted. Hwang et al. [2016] study the effects of HERA, which raised FHA loan limits and induced more borrowing, but primarily through cash-out refinancing and higher LTV ratios at purchase rather than the purchase of larger home. Park [2016] studies the same limit increases and finds an increase in FHA-qualified loan originations in 2008 but no corresponding decrease in 2014 when the limit increases expire; he concludes that borrowers substitute toward the private market as it recovers in the interim years. Passmore and Sherlund [2016] find real effects of FHA lending during this period; in particular, they document that counties with ex-ante higher FHA lending experience smaller declines in mortgage purchase originations, house prices, and new automobile purchases as well as smaller increases in unemployment rates during the crisis. DeFusco et al. [2020] document a crisis-era FHA policy excluding unemployed borrowers from refinancing severely reduced the refinancing activity of those with greatest demand to do so, curtailing the effects of monetary policy. I document several new facts relative to this literature, namely, the bunching behavior of borrowers at the county loan limit, and the responsiveness of borrowers in the upper quantiles of the borrowing distribution to the cross-sectional increase in the limits.

3.3 The Federal Housing Administration

3.3.1 *The History and Function of the FHA*

The Federal Housing Administration, or FHA, was founded in 1934 to help stabilize the housing market during the Great Depression. By providing insurance for mortgage principle in the event of borrower default, the government provided assurances to help keep lenders operating in the market [DeFusco et al., 2020]. The program was designed to be self-funding, with insurance premiums sufficient to cover program costs due to mortgage default.

Since its creation, the prominence of FHA lending in the mortgage market has waxed and waned. In the aftermath of World War II, the Veteran's Administration entered mortgage lending and the FHA lost market share. Amid civil rights legislation during the 1960s, the FHA was brought under the regulation of the Department of Housing and Urban Development, or HUD, and authorized to expand in scope by congress. In the late 20th century, the FHA lost market share as the growth of securitization and private mortgage insurance made mortgage credit more easily available to sub-prime borrowers in the private market. During the housing crisis of the late 2000s, FHA market share grew substantially from 4.5% to about 25% of mortgage originations. Since then, it has dwindled again, but remains around 15-20%. [Immergluck, 2011]

The FHA program has remained self-sufficient, with a single exception in the aftermath of the housing crisis [Puzzanghera, 2013]. Insurance premiums through the program are two-fold, and consist of an up-front mortgage insurance premium (UFMIP) and an annual premium. The size of these premiums may be updated by regulators and is disclosed in the FHA's Mortgagee Letters. This historical values of these premium rates are depicted in Figure (3.4); at present, the UFMIP rate is 1.75% and the annual MIP rate is 0.85%.

The Government National Mortgage Association, or GNMA, was founded in 1968. Since its founding, it has served to guarantee timely payment of principal and interest on MBS backed by pools of government-insured mortgage loans, most frequently loans with FHA insurance. Although it does not buy mortgages or issue MBS itself, this guarantee provides FHA loans with access to the secondary market. [FDIC, 2018]

At present, loans that qualify for FHA insurance tend to serve low-income borrowers and first-time home buyers who cannot afford down-payments for conventional mortgages and do not qualify for private mortgage insurance. Table (3.1) summarizes a 1pp sample from HMDA including all first mortgages on owner-occupied homes in 1-4 family units originated during 2018 and 2019. FHA-insured loans tend to be smaller and more standardized in repayment structure, and to have higher LTV and DTI ratios. Borrowers obtaining these loans tend to be younger, have lower income, and be hispanic or Black. The purchased homes are in poorer census tracts with higher minority populations and lower home-ownership rates.

Table 3.1: HMDA Sample† Summary Statistics

	Conventional N=93936	FHA-insured N=21299
Loan Characteristics:		
Loan Amount ($k)	283 (264)	215 (107)
Property Value ($k)	472 (8881)	243 (550)
Interest Rate (%)	4.43 (0.92)	4.50 (0.65)
Total Costs ($k)	4.12 (3.93)	7.58 (6.74)
Rate Spread (%)	0.48 (0.86)	1.37 (8.15)
LTV (%)	73.3 (20.3)	92.4 (10.3)
DTI (%)	34.7 (9.90)	42.9 (10.1)
Purpose: Purchase	53.8%	70.8%
Term: 30y	77.0%	96.1%
Rate Type: Fixed Rate	87.6%	98.1%
Repayment Type: Vanilla	91.8%	96.6%
Borrower Characteristics:		
Annual Income ($k)	178 (6486)	113 (2272)
Age (y)	47.0 (14.7)	42.4 (13.8)
Ethnicity: Hispanic or Latino	10.9%	21.7%
Race:		
Asian	7.31%	2.65%
Black or African American	4.37%	14.9%
White	85.4%	79.1%
Sex:		
Female	23.2%	29.1%
Joint	44.7%	34.3%
Male	32.1%	36.6%
Census Tract Characteristics:		
Population (k)	5.70 (3.18)	5.94 (3.34)
Minority Population (%)	28.7 (23.8)	36.0 (27.5)
Median Income ($k)	93.8 (39.7)	75.4 (26.8)
1-4 Unit Homeownership Rate (%)	74.3 (14.8)	70.4 (14.6)
Median Age of Housing Units (y)	35.5 (17.8)	34.9 (17.6)

†1pp; 2018-19; 1st Lien; Owner-Occupied; 1-4 Unit Dwelling

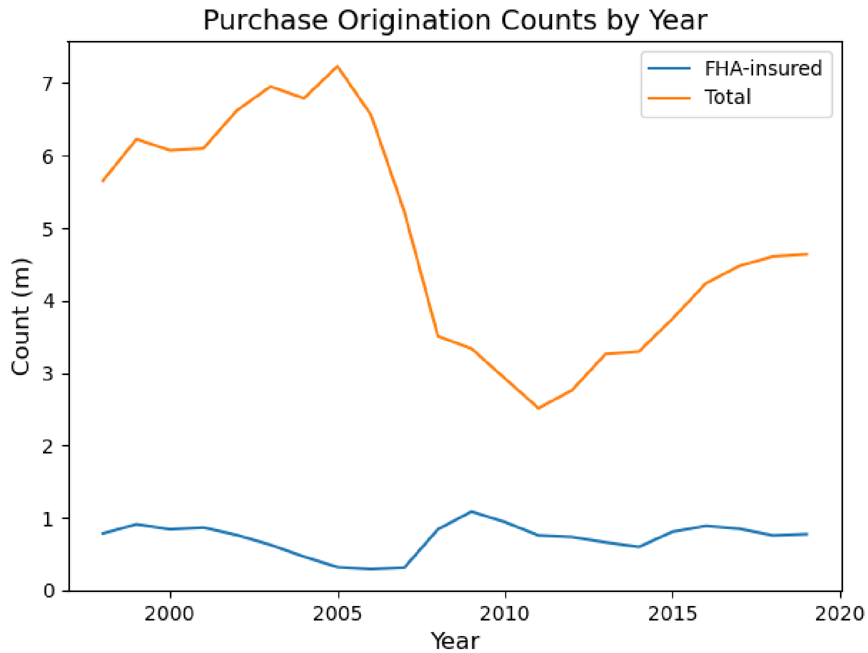

Figure 3.1: FHA loans have historically constituted 10%-30% of purchase originations. During the Great Recession, they grew to about 1 in 3 purchase originations. They remain around 20% of the market at present.

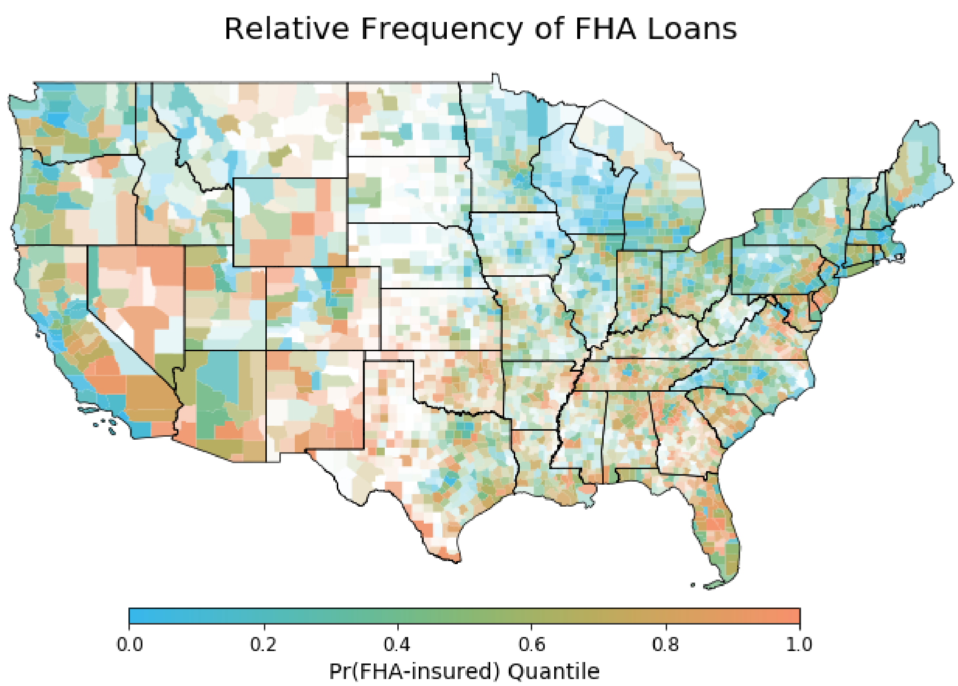

Figure 3.2: Darker counties have a higher count of FHA purchase originations. Redder counties have a higher relative rate of FHA originations. Urban centers have the most FHA lending though it tends to be relatively high at urban center peripheries.

3.3.2 FHA Loan Limits and the Bunching Framework

The FHA restricts access to mortgage credit in a variety of different ways including characteristics of the borrower, such as their FICO score, characteristics of the purchased housing, which must meet basic sanitary and safety requirements, and the characteristics of the loan itself. The main identification strategy pursued in this paper relies on the presence of two limits, a loan-to-value limit, $\overline{L}$, and an initial balance limit, $\overline{B^o}$. In this section, I describe the borrowing limits imposed on FHA loans and how these limits correspond to the limits central to the identification strategy.

To qualify for insurance through the FHA, loans must conform to three limits: a county-level cap on the initial loan balance, $\overline{B^o}_{jt}$, a cap on the loan-to-value ratio of the loan, $\overline{L}_t$, and a cap on the resulting debt-to-income ratio of the borrower following the loan, $\overline{D}_{ijt}$. Consider borrower i in county j and year t with annual (monthly) income y_{ijt} ($\hat{y}_{ijt}$) and non-mortgage debt service expenses κ_{ijt} ($\hat{\kappa}_{ijt}$) who purchases a home at price P_{ijt} with a 30-year fixed-rate mortgage characterized by initial balance B^o_{ijt} and annual interest rate r_{ijt} (and implied monthly amortizing factor $\hat{r}_{ijt} \equiv \frac{r_{ijt}/12}{1-(1+r_{ijt}/12)^{-30*12}}$). The FHA restrictions may be represented as follows:

$$B^o_{ijt} \leq \overline{B^o}_{jt} \tag{3.1}$$

$$\frac{B^o_{ijt}}{P_{ijt}} \leq \overline{L}_t \tag{3.2}$$

$$\frac{\hat{r}_{ijt}B^o_{ijt} + \hat{\kappa}_{ijt}}{\hat{y}_{ijt}} \leq \overline{D}_{ijt} \tag{3.3}$$

Assume that the rate of the loan is pinned down by borrower rather than transaction characteristics (i.e. that it does not depend on the loan size). We can then solve Equation (3.3) for B^o_{ijt} and combine this with Equation (3.1) to obtain an individual-level loan balance limit:

$$B^o_{ijt} \leq \overline{B^o}_{ijt} \equiv \min\left\{\overline{B^o}_{jt}, \frac{1}{\hat{r}_{ijt}}\left[\overline{D}_{ijt}y_{ijt} - \kappa_{ijt}\right]\right\} \tag{3.4}$$

We can rewrite Equations (3.2) and (3.4) in terms of prices and down-payments so that they correspond more closely with the bunching framework outlined in Section (3.4.1). For this change of variables, we use the fact that the down-payment and initial balance together make the transaction price, $P_{ijt} = D_{ijt} + B^o_{ijt}$.

$$P_{ijt} \overset{(3.2)}{\leq} P_{ijt} \frac{1 - B^o_{ijt}/P_{ijt}}{1 - \overline{L}_{ijt}} = \frac{D_{ijt}}{1 - \overline{L}_{ijt}} \tag{3.5}$$

$$P_{ijt} = D_{ijt} + B^o_{ijt} \overset{(3.4)}{\leq} D_{ijt} + \overline{B^o}_{ijt} \tag{3.6}$$

Finally, we can combine Equations (3.5) and (3.6), differentiate with respect to the price of the home purchased, and write the piece-wise function for the marginal down-payment relative to a dollar increase in housing demand. Note that we define a reference house price, $P^*_{ijt} = \frac{\overline{B^o}_{ijt}}{\overline{L}_{ijt}}$, and reference down-payment, $D^*_{ijt} = (1 - \overline{L}_{ijt})P^*_{ijt}$, as the price and down-payment at which (3.5) and (3.6) are both binding.

$$\frac{dD_{ijt}}{dP_{ijt}} \geq \begin{cases} 1 - \overline{L}_{ijt} & P_{ijt} < P^*_{ijt} \\ 1 & P_{ijt} \geq P^*_{ijt} \end{cases} \tag{3.7}$$

3.3.3 FHA Loan Limit Assignment Rules

Identifying the choice set in the space of prices and down-payments available to an individual borrower in the data requires imputing the loan-to-value and borrower-level initial balance limits assigned to them in the FHA program. The loan-to-value limit was 0.965 for all borrowers over the entire sample period. Borrower-level initial balance limits are more sophisticated. Per Equation (3.4), I assign these as $\overline{B^o}_{ijt} \equiv \min\left\{\overline{B^o}_{jt}, \frac{1}{\hat{r}_{ijt}}[\overline{D}_{ijt}y_{ijt} - \kappa_{ijt}]\right\}$. Below, I review the method for assigning county-level loan limits, $\overline{B^o}_{jt}$, and debt-to-income limits, $\overline{D}_{ijt}$, in the FHA and in my data. I review how I identify household debt-service, κ_{ijt}, which is not explicitly available as a field in the merged data.

The county-level loan limit, $\overline{B^o}_{jt}$, is computed as the product of HUD's county-median home price index the previous year, $\tilde{P}_{j,t-1}$, and a median home price multiplier, θ_t, unless this value is below the FHA national loan limit floor, $\underline{\overline{B^o}}_t$, or above the FHA national loan limit ceiling, $\overline{\overline{B^o}}_t$, in which case it takes on those values, respectively. The FHA national loan limit floor and ceiling are set according to the conventional conforming national loan limit floor, $\underline{\overline{B^o}}_{G,t}$, as well as multipliers for the floor, $\underline{\varphi}_t$, and ceiling, $\overline{\varphi}_t$. (In particular, $\underline{\overline{B^o}}_t = \underline{\varphi}_t \underline{\overline{B^o}}_{G,t}$ and $\overline{\overline{B^o}}_t = \overline{\varphi}_t \underline{\overline{B^o}}_{G,t}$.) Formally, this may be represented as follows:

$$\overline{B^o}_{jt} = \begin{cases} \underline{\overline{B^o}}_t & \theta_t \tilde{P}_{j,t-1} \le \underline{\overline{B^o}}_t \\ \theta_t \tilde{P}_{j,t-1} & \theta_t \tilde{P}_{j,t-1} \in \left[\underline{\overline{B^o}}_t, \overline{\overline{B^o}}_t\right] \\ \overline{\overline{B^o}}_t & \theta_t \tilde{P}_{j,t-1} \ge \overline{\overline{B^o}}_t \end{cases} \tag{3.8}$$

Figure (3.3) depicts the historical values of the median price multiplier, θ_t, floor and ceiling multipliers, $\underline{\varphi}_t$ and $\overline{\varphi}_t$, as well as the conventional conforming national loan limit floor, $\underline{\overline{B^o}}_t^{GSE}$, and FHA national loan limit floor and ceiling, $\underline{\overline{B^o}}_t$ and $\overline{\overline{B^o}}_t$. Figure (3.5) depicts loan limits assigned to each county by the FHA against the median home price used to assign the limit. A separate panel is used for every period in which HUD enforced different limits. It is clear that the assigned limits largely follow the truncated linear assignment rule. In the immediate aftermath of the financial crisis, 2008-2013, additional provisions preventing a fall in limits are responsible for the observable noise in the limits from the stated rules. I can merge these limits into the data using the county and origination month or year of the transaction.

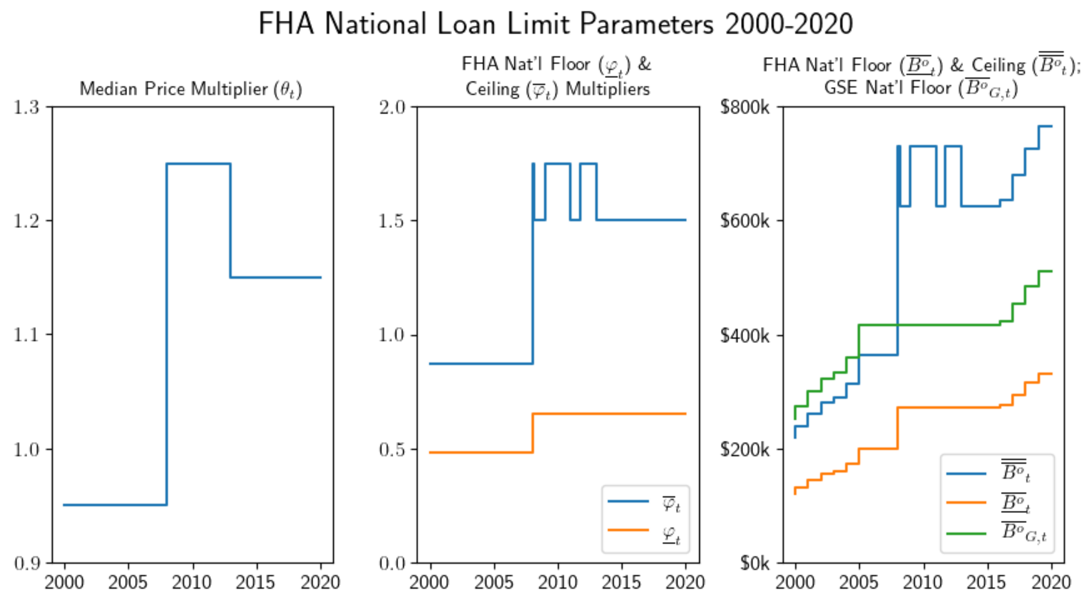

Figure 3.3: Parameters for FHA initial balance limit. Source: HUD website.

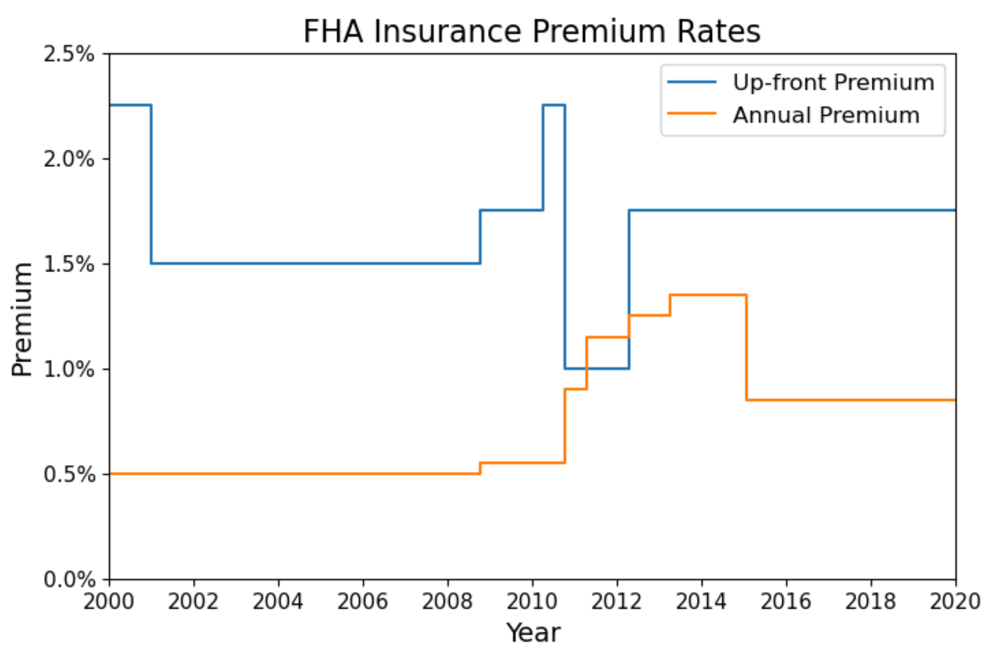

Figure 3.4: Up-front and annual mortgage insurance premium rates for FHA loans. Source: HUD Mortgagee Letters.

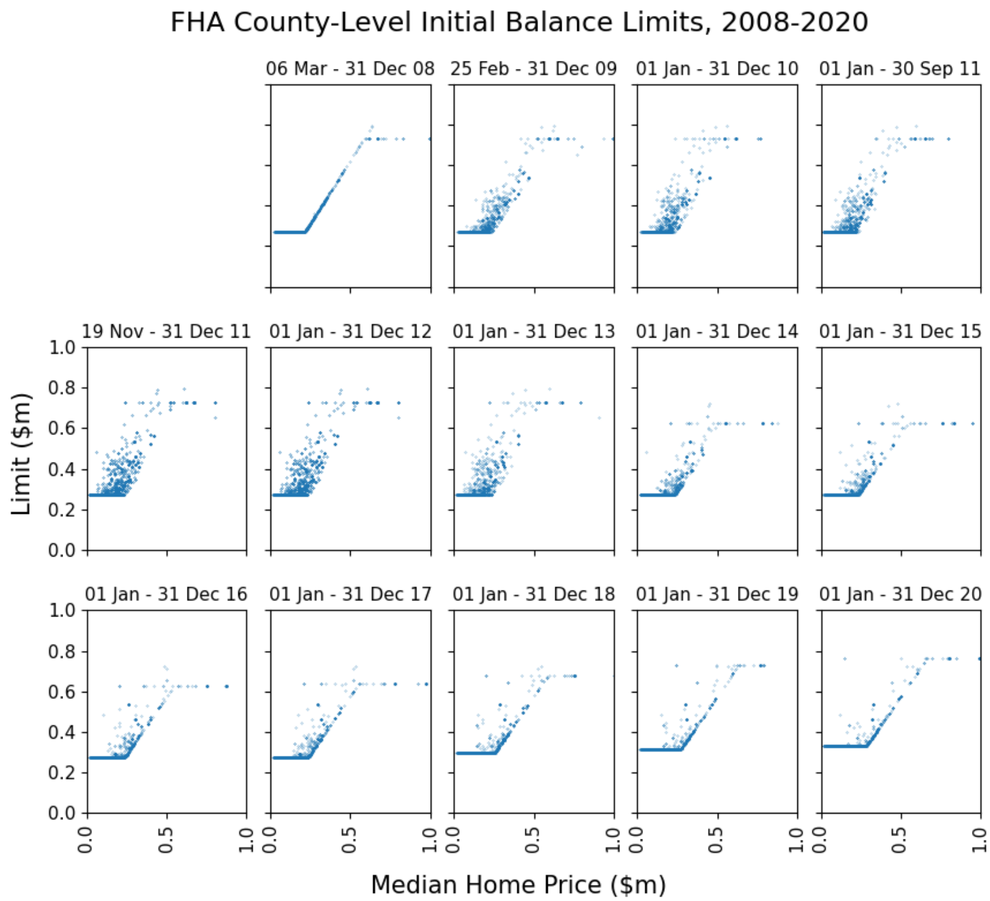

Figure 3.5: County-level FHA loan limit by median house price and time in force, as designated by the Mortgagee Letter. County limits largely follow assignment rules; noise follows exceptions introduced in the aftermath of the financial crisis. Source: HUD website and HUD Mortgagee Letters.

The debt-to-income limit imposed by the FHA varies according to whether the borrower is approved through the Fannie Mae Automated Underwriting System (AUS) or requires approval through the Manual Underwriting System (MUS). Typically, lenders will attempt to obtain approval through the AUS, which has a maximum debt-to-income limit of 0.569. If borrowers do not qualify under the AUS, then the MUS determines a borrower's debt-to-income limit according to borrower FICO score and a number of "compensating factors", which are indications that the borrower may have sufficient assets or income to afford the loan. In particular, for borrowers with a FICO score between 500 and 580, the MUS sets a debt-to-income limit of 0.43. For borrowers with a FICO score between 580 and 850, the MUS sets a debt-to-income limit of 0.43 for no compensating factors, 0.47 for one, and 0.50 for two compensating factors.

In practice, it is difficult to identify whether a loan was underwritten by the MUS or AUS. It is also difficult to identify the number of compensating factors enjoyed by a borrower in the absence of the administrative data. Although I have the FICO score of borrowers in the data, there are no borrowers in the merged sample with a FICO between 500 and 580, so this is of no value in discriminating debt-to-income limits. Instead, I assign borrowers the lowest debt-to-income limit that is still above that which they received in the data.

Although non-mortgage debt service, κ_{ijt}, is not a field in the merged data, it can be obtained by subtracting the monthly mortgage expense on the observed loan from the total debt service implied by the the debt-to-income ratio and income of the borrower, $\kappa_{ijt} = y_{ijt}D_{ijt} - \hat{r}_{ijt}B^{o}_{ijt}$.

3.4 Empirical Strategy

3.4.1 The Household Problem

A household has preferences over housing, H_t, and non-housing consumption, X_t, in each of two periods, t_0 and t_1. To simplify the derivations shown in Appendix (B.1), I assume standard concavity, that housing and non-housing consumption are separable within period, and that the borrower discounts with factor β. For exogenous reasons, the borrower opts to purchase a home at t_0, and sell the home in favor of renting in t_1. The borrower is endowed with an income made available in each period, y_0 and y_1, and chooses how to allocate this wealth as consumption. Of particular interest is the intensive margin of household demand for housing, H_0.

I normalize the initial price of non-housing consumption to 1; it rises according to the risk-free rate to $(1+r_f)$ at t_1. The price of housing is initially p_H and evolves as $p_H(1+\tilde{r}_H)$. It is the borrower's beliefs that matter when choosing how much housing to purchase and, for simplicity in the derivations, I assume these are dogmatic and degenerate, $\mathbb{E}_i^0[\tilde{r}_H] = \nu_i$.

Households have two financial vehicles to facilitate their consumption decision. They may save or borrow at the risk-free rate along what I call a "liquid" margin, analogous to a checking account or credit card. Along this margin, they face a standard borrowing constraint, which I normalize to 0 for simplicity. They may also take out a mortgage loan to finance the purchase of their housing, also at the risk-free rate, r_f. As is standard in mortgage lending, this loan is subject to two constraints, a loan-to-value limit, $\overline{L}$, and a debt-to-income limit, the latter of which I rewrite as an initial mortgage balance limit, $\overline{B^0}$.

To choose the optimal quantity of housing at t_0, the household considers the marginal costs and benefits, in utility terms, of an additional unit of housing. I depict this choice in Figure (3.6). The marginal benefits to an additional unit of housing are standard. They are the sum total of utility from t_0 housing services, U_H^0, and the discounted utility from additional t_1 non-housing consumption the borrower anticipates purchasing out of excess returns on housing relative to the risk-free asset, $\beta U_X^1 p_H \mathbb{E}_i^0[\tilde{r}_H - r_f]$. In general these are falling due to concavity in both X and H.

The household's marginal costs feature a potential non-standard discontinuity due to the two constraints on the mortgage loan. I depict this discontinuity in Figures (3.8) and (3.9). A household exhausting their loan-to-value limit can rely on additional mortgage financing until they also exhaust their initial balance limit. This occurs at the reference quantity of housing, $H^* = \frac{1}{p_H}\frac{\overline{B^0}}{\overline{L}}$. Below this quantity, the utility cost of an additional unit of housing derive from lost t_0 consumption due to additional down-payment and lost t_1 consumption due to additional costs of eventual mortgage repayment, $p_H[(1-\overline{L})U_X^0 + \overline{L}U_X^1]$. Beyond this quantity, the price of an additional unit of housing is financed entirely from a down-payment out of t_0 consumption, and the utility costs are $p_H U_X^0$.

The household's additional margin for borrowing and saving outside of the mortgage contract plays a key role in generating the discontinuity in marginal costs. If households are not at their "liquid" borrowing constraint, then they may smooth consumption by equating marginal utility of non-housing consumption in both periods. In this case, there will be no discontinuity in marginal costs, $p_H[(1-\overline{L})U_X^0 + \overline{L}U_X^1] \overset{U_X^0=U_X^1}{=} p_H U_X^0$. This is akin to the Modigliani-Miller benchmark. However, if borrowers are constrained along other margins, then the household would benefit from more consumption smoothing and so $U_X^0 > U_X^1$, generating the upward discontinuity in marginal costs.

In the absence of constrained households, therefore, there will be a smooth distribution of households across the reference quantity of housing, H^*. In Figure (3.7), in the counterfactual case, borrowers b and m choose different quantities of housing. Denote this counterfactual distribution as $f(H)$. By contrast, when households are constrained, they tend to bunch at the reference quantity. In Figure (3.10), buncher b and marginal household m both choose the reference quantity of housing, H^*.

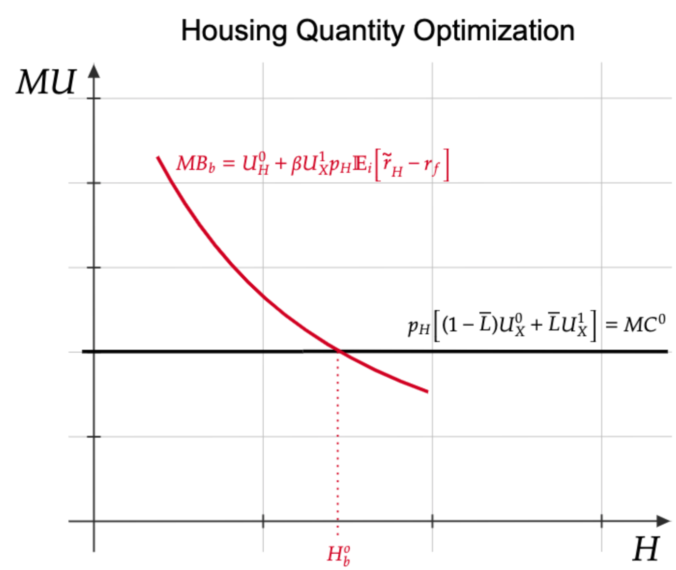

Figure 3.6: The marginal benefits from housing are the sum of utility due to housing services, U_H^0, and utility due to future consumption afforded by expected excess returns on housing investment, $p_H \mathbb{E}_i\left[\tilde{r}_H - r_f\right] * \beta U_X^1$. For the bunching counterfactual, in which a borrower is not subject to a debt-to-income or initial balance limit, the marginal costs are due to a downpayment requirement out of present consumption, $p_H(1-\overline{L}) * U_X^0$, and mortgage repayment out of future consumption, $p_H \overline{L}(1+r_f) * \beta U_X^1$. Borrowers obtain the optimal quantity of housing by equating the marginal costs and benefits of an additional unit, H_b^o.

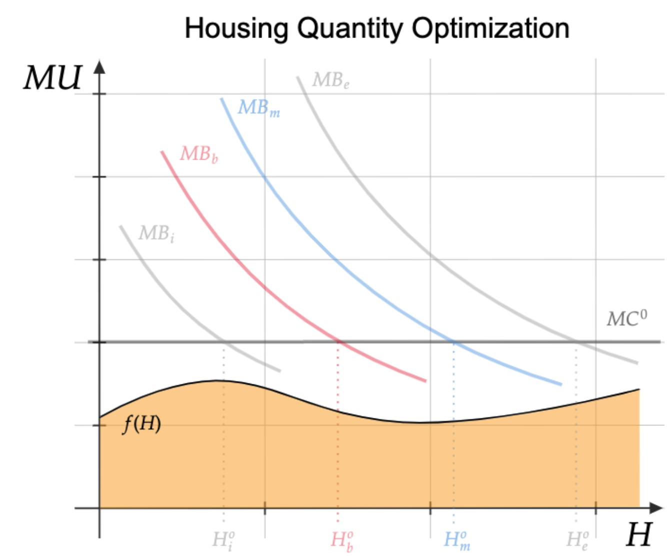

Figure 3.7: For the bunching counterfactual, in which a borrower is not subject to a debt-to-income or initial balance limit, households exhibit a distribution of demand for quantities of housing. This heterogeneity may be driven by preferences for housing relative to non-housing consumption, or even wealth orthogonal to the extent of household credit constraints. The optimal choices of four households, $\{i, b, m, e\}$, are depicted. The distribution of optimal housing quantities in the counterfactual is summarized as $f(H)$.

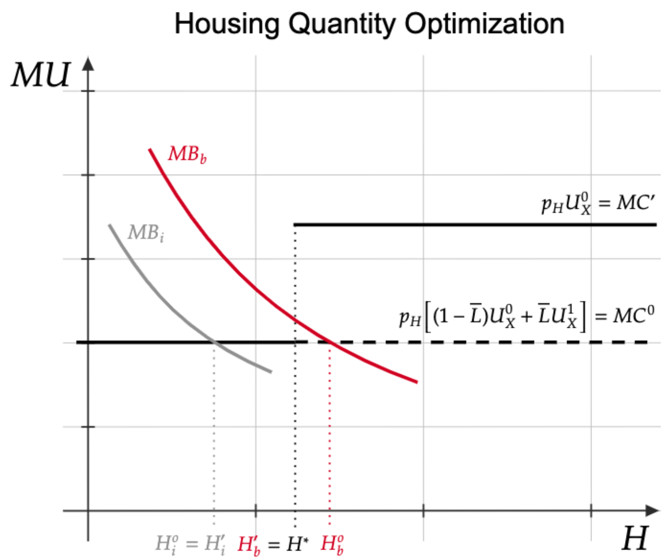

Figure 3.8: With the introduction of the initial balance constraint, the marginal costs of housing are unchanged under the reference quantity, H^*. Above the reference quantity, the marginal costs are due entirely to down-payment requirements out of present consumption, $p_H * U_X^0$. For unconstrained households, consumption smoothing delivers $U_X^0 = U_X^1$ and $MC' = MC^o$. For constrained households, the constraint implies $U_X^0 > U_X^1$ and $MC' > MC^o$. Borrowers adjust their housing consumption if their counterfactual optimum consumption lies above the reference quantity. Borrower i, in grey, is unaffected. Borrower b, in red, reduces demand from H_b^0 to $H_b' = H^*$.

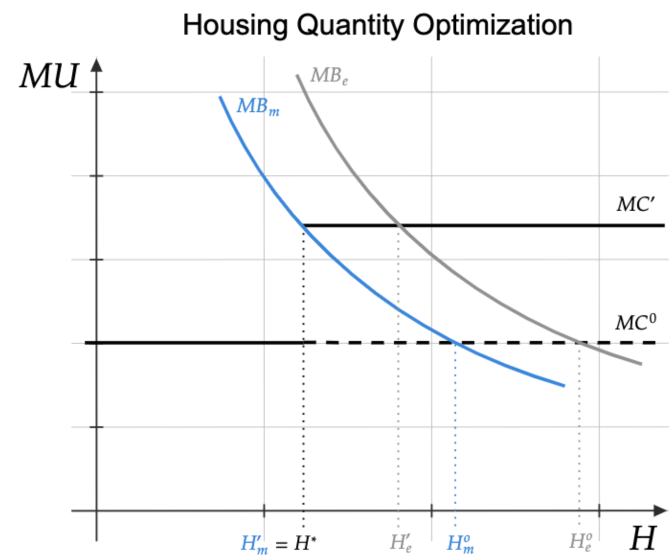

Figure 3.9: The marginal borrower, m, in blue, is the borrower who adjusts to the reference quantity of housing, H^*, in the presence of the initial balance constraint, and whose marginal costs and benefits of housing are equal after adjustment. Extra-marginal borrowers, e.g. e, may adjust in the presence of the DTI constraint, but does not adjust to the reference quantity.

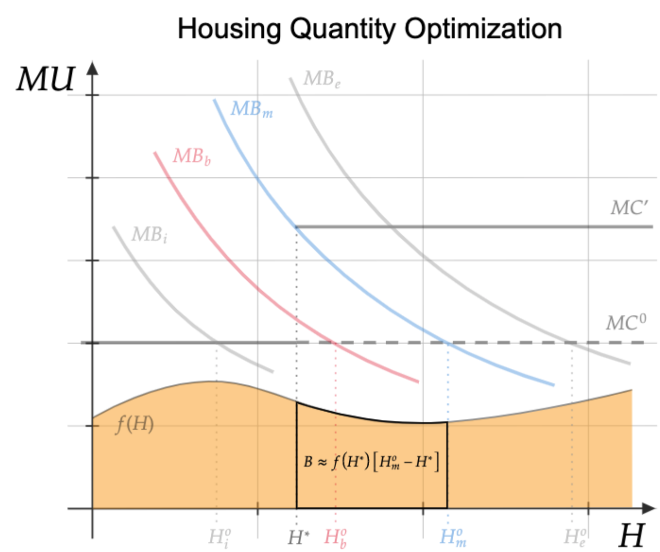

Figure 3.10: Bunching is due to all borrowers with housing demand in the counterfactual regime between the reference quantity, H^*, and the marginal buncher's demand, H^o_m. The mass of bunching may be measured as the mass of such borrowers and approximated as $B = f(H^*)[H^o_m - H^*]$. This is the density of borrowers at the reference quantity (under the counter-factual distribution) times the marginal borrower's behavioral response.

3.4.2 Identification

To identify the loan-to-value elasticity of housing demand, I rely on as-good-as-random assignment of counterfactual household demand for housing across the reference housing quantity. At and above the reference quantity, the decreased effective loan-to-value limit induce an increase in utility costs of additional housing for borrowing constrained households. For this reason, households with counterfactual demand above the reference quantity may be considered the "treated" households and they respond by adjusting their demand downward. Households with counterfactual demand below the reference quantity, the "control" group, do not respond. The size of the resulting bunching relative to the unperturbed distribution measures the behavioral response of the "treated" households.

I measure the loan-to-value limit elasticity of housing demand, $\widehat{\varepsilon}^{H}_{\overline{L}} = \frac{\%\Delta H}{\%\Delta \overline{L}}$. I consider the behavioral response of the marginal buncher to the reduction in the loan-to-value limit on housing units above the reference quantity from $\overline{L}$ to 0. I also observe that house prices are simply house quantities scaled by a constant price level. The estimator can therefore be rewritten as:

$$\frac{\%\Delta H}{\%\Delta \overline{L}} = \frac{\frac{H^* - H^o_m}{H^o_m}}{\frac{0 - \overline{L}}{\overline{L}}} = \frac{[H^o_m - H^*]}{H^* + [H^o_m - H^*]} = \frac{[P^o_m - P^*]}{P^* + [P^o_m - P^*]}$$

The reference house price, P^*, can be obtained from the household's borrowing limits. The adjustment of the marginal borrower, $[P^o_m - P^*]$, can be obtained from the bunching equation:

$$B = \int_{P^*}^{P^o_m} f(P)dP \approx f(P^*)[P^o_m - P^*]$$

The bunching mass, B, can be measured in the realized distribution. And the counterfactual density at the reference house price $f(P^*)$, can be approximated from the realized density as well.

Below, I discuss the appropriate interpretation of the estimate. I distinguish between loan-to-value limits on average and marginal units of housing, establish the nature of the measured elasticity relative to these, and describe the relationship of the measured elasticity relative to loan-to-value and debt-to-income policy parameters. I then distinguish between partial and general equilibrium loan-to-value limit elasticities and characterize one advantage of estimating partial equilibrium elasticities.

3.4.3 Interpretation

Marginal vs. Average LTV Limit Elasticity

Because loan-to-value limits determine the trade-off between down-payment supply and housing demand, they function as prices in the household problem.[1] In this capacity, their impact resembles compensated and uncompensated demand elasticities from the standard consumer theory. Households may reduce housing demand in response to a lower loan-to-value limit on the entire loan (i.e. the average unit of housing purchased), akin to an uncompensated price elasticity. Households may also reduce housing demand in response to a lower loan-to-value limit on merely the last unit of housing, akin to a compensated price elasticity. These behavioral responses are depicted in Figure (3.12).

1. Between the price level of housing, p_H, and the loan-to-value limit, $\overline{L}$, this suggests that households face two prices associated with housing. This comports with the intuition that households make credit constrained down-payments at the time of purchase and unconstrained mortgage payments at later dates. In other words, there are two separate margins along which households sacrifice consumption.

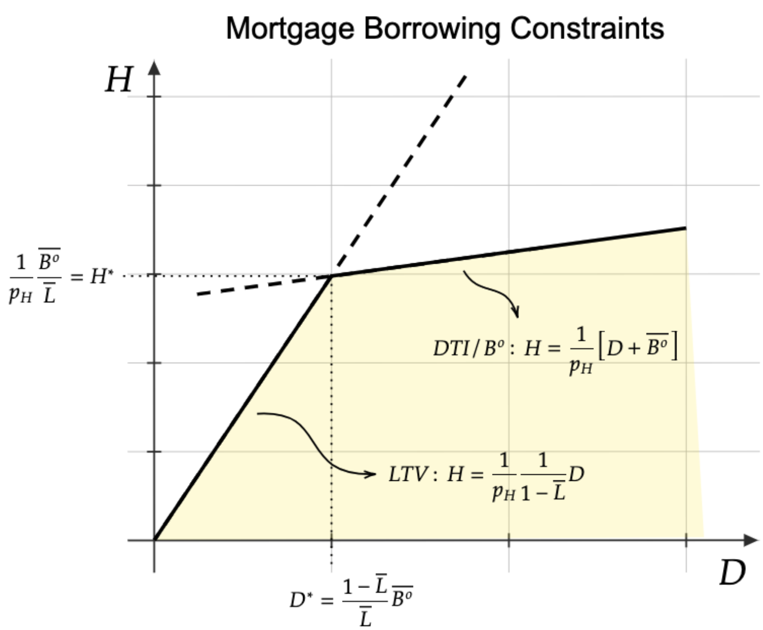

Figure 3.11: The HH choice set in downpayment-housing space is depicted in yellow. It is bounded on the left by the the LTV constraint and above by the DTI constraint, which is really an initial balance constraint. The reference down-payment, D^*, and housing quantity, H^*, are depicted at the point of intersection of the two constraints. Solving the constraints for D and differentiating in H, we obtain the marginal down-payment requirement for an additional unit of housing:

$$\frac{dD}{dH} = \begin{cases} p_H * 1 & H < H^* \\ p_H * (1 - \overline{L}) & H \geq H^* \end{cases}$$

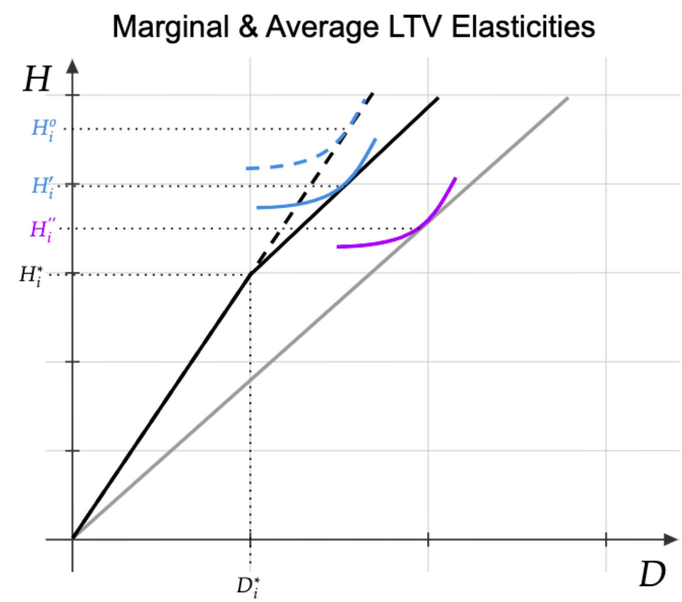

Figure 3.12: Changing LTV requirements on a loan may apply either to marginal units of housing or to the average unit of housing, i.e. the entire loan. The dashed black line depicts a baseline in which borrowers must comply with a LTV limit, $\overline{L}^o$, over the entire loan. The solid black line depicts a regime in which borrowers must comply with the lower LTV limit, $\overline{L}'$, for marginal units above a reference quantity, H^*. The solid grey line depicts a regime in which borrowers must comply with the lower LTV limit, $\overline{L}'$, on the entire loan. The behavioral response of households is depicted and, in principle, may differ.

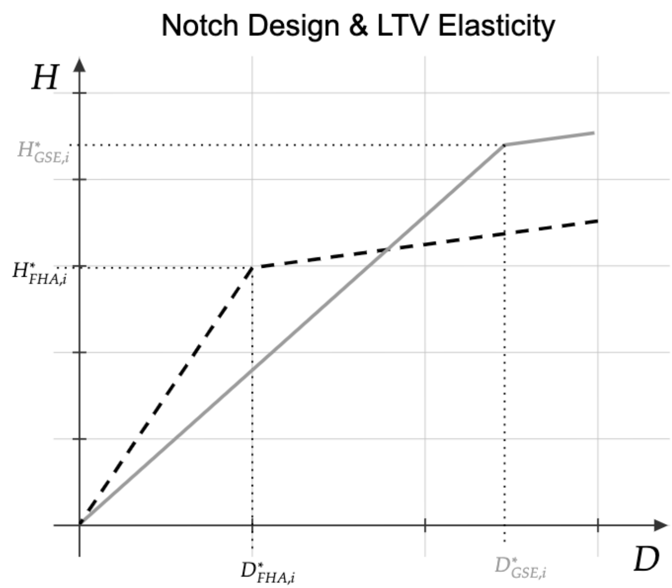

Figure 3.13: The dashed black line outlines the choice set due to lending limits of the FHA program. The solid gray line outlines that due to lending limits of the GSEs (in the conforming segment). The reference down payments, D^*, and housing quantities, H^*, are labeled for each program. Note that this nearly represents a notch design from the bunching literature, where the average loan-to-value required increases above downpayment, D^*_{FHA}.

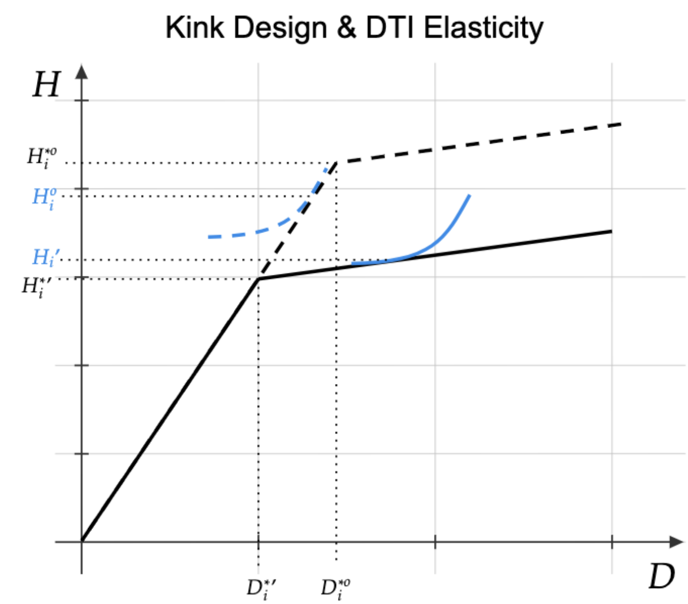

Figure 3.14: The dashed black line shows an initial scenario with a relaxed debt-to-income limit. The solid black line shows a counterfactual with a more restricted debt-to-income limit. The reference down-payments, D^*, and housing quantities, H^*, are labeled. The behavior of borrower i in the initial and counterfactual scenario are depicted in the dashed and solid blue lines, respectively. A kink design therefore approximates the change in housing demand for the marginal agent when the debt-to-income ratio is relaxed.

Compensated and uncompensated price elasticities are related in the famous Slutsky equation by income effects. In this context, the marginal, $\varepsilon_{\overline{L}}^{H,mg.}$, and average, $\varepsilon_{\overline{L}}^{H,avg.}$, loan-to-value limit elasticities are related analogously by a quasi-income effect, where the quasi-income is the loan balance obtainable before making any down-payment. Following the sketch in Kleven [2016], I use this Slutsky-type relation to show in Appendix (**??**) that the estimated elasticity is a weighted average of these elasticities:

$$\widehat{\varepsilon}_{\overline{L}}^{H} = \left[1 - \frac{\Delta a}{\Delta \mu}\right]\varepsilon_{\overline{L}}^{H,mg.} + \left[\frac{\Delta a}{\Delta \mu}\right]\varepsilon_{\overline{L}}^{H,avg.} \tag{3.9}$$

The estimator contains some information about average loan-to-value limit elasticity, but there is clearly a margin for bias. Concerns about bias may be mitigated for two reasons. First, the value of $\frac{\Delta a}{\Delta \mu}$, which ranges from 0 to 1, increases in the sharpness of the kink and in the present setting the kink is quite sharp. Second, to the extent that quasi-income effects are null, we have that the marginal and average loan-to-value limit elasticities are identical and the estimated elasticity captures both. In the bunching literature on tax elasticities of earnings, structural models tend to impose this.

Policy LTV and DTI Limit Elasticities

Loan-to-value and debt-to-income limits are policy parameters that can be used to regulate the mortgage and housing markets. Below, I review the nature of these policy elasticities and the ways in which they relate to and differ from the conceptual elasticities introduced above. Though the estimator in this setting captures some information about behavioral responses to changes in policy loan-to-value limits, it does a better job of estimating responses to changes in policy debt-to-income limits. A related research design exploiting features of the mortgage market may better capture responses to policy loan-to-value limits.

The effect of adjusting policy loan-to-value limits bears closest resemblance to the average loan-to-value limit elasticity because it affects down-payment requirements on all units of housing up to the reference quantity. To the extent that debt-to-income limits restrict household behavioral response to a change in the policy limit, the effect of policy may understate the average loan-to-value limit elasticity. Setting aside this caveat, the decomposition of the estimated elasticity above indicates that it captures some information about effects of changing loan-to-value policy limits, with associated caveats.

In spite of its shortcomings, this research design suggests another that might better capture the policy effects of changing loan-to-value limits. Conventional mortgages and FHA mortgages have different loan to value limits and, as described, FHA mortgages have limits on initial balances. This variation is depicted in Figure (3.13) and serves as a close analogy to the bunching notch designs in the literature, which do a better job measuring average rather than marginal effects.

The effect of adjusting policy debt-to-income limits is more similar to the marginal loan-to-value limit because they affect effective loan-to-value limits only on units of housing beyond the reference quantity and not on initial units. Again, if the adjusted debt-to-income limit remains binding, the policy effect may understate the elasticity. Additionally, because non-infinitesimal changes in debt-to-income limits imply quasi-income effects, the policy effect may overstate the marginal loan-to-value limit elasticity. These quasi-income effects are precisely the reason for the non-zero weight placed on the average loan-to-value limit elasticity in Equation (3.9). For this reason, the policy effects of changing debt-to-income limits are those best captured by this framework. This notion is depicted in Figure (3.14).

Partial vs. General Equilibrium LTV Limit Elasticity

In partial equilibrium, an increase in the availability of mortgage credit may increase housing demand by relaxing household credit constraints, all else equal. In general equilibrium, aggregate increases in the availability of mortgage credit may affect housing demand through effects in other markets. For example, as household wealth increases due to wages or non-housing asset values, wealth effects may further drive demand. The strategy in this paper identifies the partial equilibrium effect of loan-to-value limits on housing demand.

An advantage of measuring the partial equilibrium elasticity is that it can better isolate the effects of easy credit on household demand for housing without capturing potentially confounding effects, such as beliefs about housing returns. Indeed, the relative roles of changing beliefs and easy credit remains an enduring question about the housing cycle of the early 2000s. A predominant empirical approach in the literature has been to exploit heterogeneity in the cross-section of geographies and to find instruments at the geographic level for the availability of credit. A shortcoming of this approach is that it can only identify the pure effects of relaxed credit constraints under fairly strong assumptions about household beliefs.

To see this, consider a simple model of housing demand in which demand is a function of county-level loan-to-value limits and individual beliefs about housing:

$$H_{ic} \sim \beta_0 + \beta_L \overline{L}_c + \beta_r \mathbb{E}_i[\tilde{r}_{H,c}] + \varepsilon_{ic} \tag{3.10}$$

Further specify a fairly general model of household beliefs as a linear combination of everything in the information set of the household:

$$\mathbb{E}_i[\tilde{r}_{H,c}] \sim \rho_0 + \rho_L \overline{L}_c + \rho_F' \mathcal{F}_i \backslash \overline{L}_c + \eta_{ic} \tag{3.11}$$

The county-level loan-to-value limit is part of the household's information set and is represented because of its possible dual role in mitigating credit constraints and shaping beliefs.

The parameter of interest is β_L, which, as described above, is the parameter for which this paper provides an unbiased estimate. Ordinarily, using a geographic instrument to resolve the problem of an omitted variable, beliefs, would also provide an unbiased estimate. However, because any county-level treatment that eases a household's credit constraints doubles as an information treatment that may shape their beliefs about housing returns in the area, this strategy may produce a biased estimate of the pure credit availability effects.

To see this, construct a candidate instrument, z_c, that satisfies relevance:

$$\overline{L}_c \sim \lambda_0 + \lambda_z z_c + u_c \qquad \lambda_z \neq 0$$

Further stipulate that the candidate is nearly exogenous in the sense that it is orthogonal to all other shocks and information:

$$z_c \perp \mathcal{F}_i \backslash \overline{L}_c, \eta_{ic}, \varepsilon_{ic}$$

The candidate instrument cannot be orthogonal to beliefs because it is non-orthogonal to loan-to-value limits by construction and, as information, these enter beliefs. For this reason, the instrumental variable estimator may include bias:

$$\hat{\beta}_L^{IV} = \frac{\frac{Cov(H_{ic}, z_c)}{Var(z_c)}}{\frac{Cov(\overline{L}_c, z_c)}{Var(z_c)}} = \ldots = \beta_L + \beta_r \rho_L$$

In practice, there may be no bias if either return expectations have no effect on housing demand, $\beta_r = 0$, or local credit availability has no effect on return expectations, $\rho_L = 0$. Because a household with more optimistic return expectations can increase their perceived wealth by weighting their portfolio toward housing, it is unlikely that the former holds. If the local housing market has perfectly elastic housing supply, it is long-run expectations that enter the demand equation, and households have rational expectations, then the latter may hold.

Rational expectations is a useful benchmark but a fairly strong assumption. In the present setting, it is possible to weaken this assumption because the effective treatment of loan-to-value limits takes place at the individual rather than aggregate level. The availability of credit experienced by the borrower in no way differentially informs them about the availability of credit experienced by other borrowers and therefore it is less plausible that they might draw conclusions about the future price of housing in the area.

3.5 Estimation

3.5.1 Sample Construction

A key obstacle for implementing the estimator outlined in the previous section is the ability to reconstruct the borrower's choice set from their characteristics so that their transaction may be located in that choice set relative to the kink. Loan-to-value limits are fairly standard across borrowers, but to impute effective initial balance limits requires knowledge of both a borrower's debt-to-income limit and their income at the time of origination. I overcome data limitations by merging HMDA and CRISM data. In doing so, I follow the observation from Bartlett et al. [2021] that these data are mergeable; I have a somewhat different purpose in doing so, and am somewhat more limited in my ability to do so extensively.

HMDA data cover 90% of mortgage originations in the United States [Bartlett et al., 2021], provide borrower-level data, and include minimally-redacted information on borrower income. However, HMDA does not cover debt-to-income limits or loan realizations until 2018 and these are heavily redacted. CRISM is a merge of Equifax credit bureau fields to BlackKnight mortgage loan origination and performance data; since 2005, it covers roughly 60% of mortgage originations [Adelino et al., 2013].

By merging the datasets, I obtain both income and realized debt-to-income values on a sample of FHA loans. My sample is limited to FHA loans in the years 2018 and 2019 for reasons of data availability. CRISM tends to report 5-digit, as opposed to 3-digit, zip codes only for FHA loans; HMDA contains enough origination variables, like home price and interest rate, only beginning in 2018. In principle, with better data, it would be possible to compute the analysis on any loan with both a loan-to-value and debt-to-income limit. I describe the merge procedure in more detail below.

I begin by constructing a large sample of FHA loans, which I refer to as the CRISM

2009-2020 Sample. This consists of loans in CRISM that are vanilla (non-IO and non-balloon), fixed-rate, 15 or 30 year, first mortgages for purchase of single-family, owner-occupied housing and are originated through the FHA between 2009 and 2020. I discard a handful of observations which are not onboarded to CRISM promptly, or have unrealistically high LTV ratios at origination. I use HUD crosswalks to identify the counties associated with the 5-digit zip codes in these samples, and use county and month of origination to merge in up-front mortgage insurance premium rates, the HUD county median home price index, county-level FHA loan limits and county-level GSE conforming loan limits, all hand-collected from mortgagee letters on HUD's website.[2]

Next, I construct another sample of FHA loans from HMDA, which I refer to as the HMDA 2018-2019 Sample. I begin by identifying loans in HMDA that are vanilla (non-IO, non-balloon, non-HELOC, non-reverse, and non-negatively amortizing), fixed-rate, 15 or 30 year, first mortgages for purchase of single-family, owner-occupied housing and are originated through the FHA between 2018 and 2019.

Finally, I assemble a sample for the main analysis in the paper, which I refer to as MSAMP3. For each loan in the CRISM 2009-2020 Sample, I search for a loan in the HDMA 2018-2019 Sample according to the following procedure. First, I consider any loan in this sample that with the same state, county, zip code, origination year, and loan term (15 or 30 years). Next, score the proximity of the loans along four characteristics, the initial balance, the transaction price, the loan-to-value ratio, and the interest rate. Because of censoring in both HMDA and CRISM, I do not penalize differences in the initial balance or transaction price up to $5k. Outside of this bandwidth, and in the case of the loan-to-value and interest rate, I apply a linear penalty for deviations, until at a large-enough bandwidth, I say there is no match between the characteristics. I require the candidate HMDA loan to match the

2. Note that it is possible, in principle, to extend this sample to 2005; the ease of collecting FHA and GSE conforming loan limits beginning in 2009 was the motivation for the sample start date. Furthermore, obtaining 5-digit zip codes in CRISM, which is necessary both for identifying county-level limits and for a successful merge to HMDA data, requires focusing on FHA loans rather than conventional mortgages in the version of CRISM to which I have access.

CRISM loan along all four characteristics and then keep a single match from any that remain that performs the best on a combination of all four dimensions.

I depict the quality of the merge in Figure (3.15). The vast majority of these merges are exact matches in interest rates and loan-to-value ratios. Further evidence of the match quality can be obtained by observing the sharp decrease in matches for which the CRISM loan amount or house price is more than $5k away from its HMDA counterpart. This is consistent with an exact match given the redaction methods in the HMDA data.

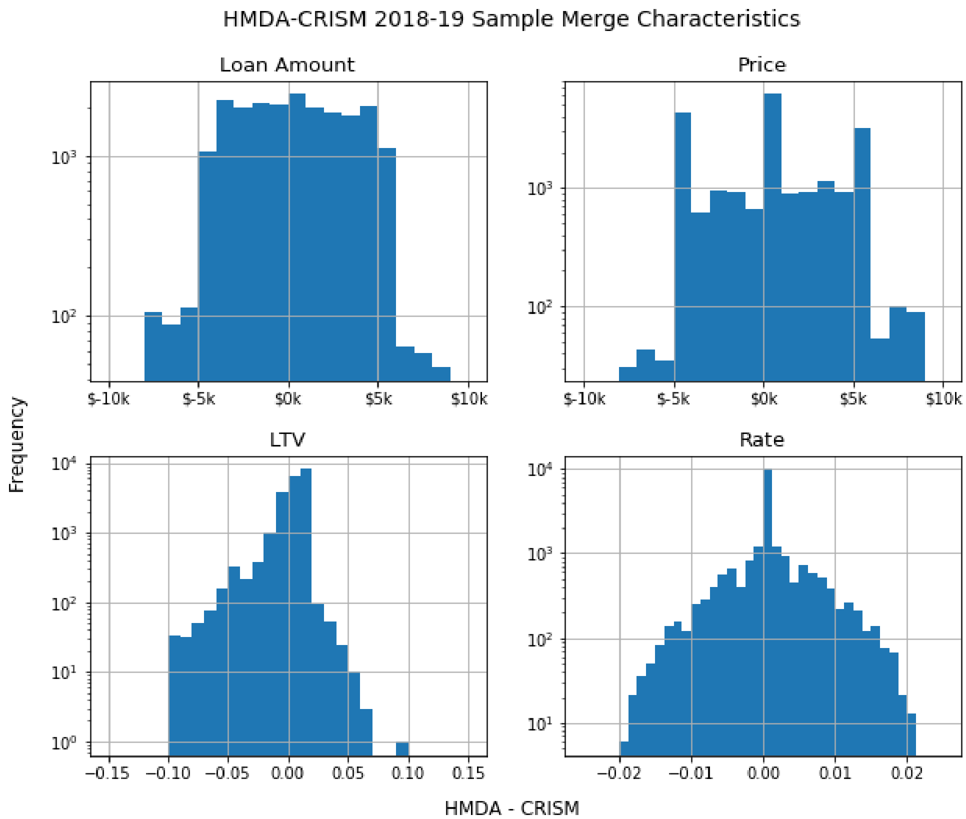

Figure 3.15: This figure depicts the merged sample in the state of California. The four panels depict differences in the HMDA and CRISM recorded values of the Loan Amount, Price, LTV, and Rate for each loan in the final sample. [**I**] The y-axis is a log-scale so that the modest central spikes, for which merged characteristics are almost identical, vastly outnumber those instances of mis-matched errors. [**II**] The steep decrease in match frequency of loans with either loan amounts or loan prices more than $5k apart. This drop in match frequency is not the result of our match scoring procedure, which moves continuously across the threshold. Instead, it is likely due to the fact that HMDA censors loan amounts and transaction prices at $5k (CRISM censors at $1k). Therefore, the censored values of correctly matched loans would have a difference of anywhere from -$5k to $5k, depending on their pre-censored values. The merged characteristics are consistent with this.

I describe the characteristics of the merged sample in Table (3.2). I am able to identify a match for 75% of the loans in CRISM. Because the CRISM sample is so small, however, this is only about 15% of the near-universal HMDA sample. Still, the contents of the CRISM sample look broadly similar to the more representative HMDA sample on the attributes available in CRISM (columns 1 and 2). And the merged sample also appears broadly similar, though it is slightly weighted towards loans that remain on balance sheet than the HDMA sample.

Table 3.2: Merge Process Summary Statistics

	CRISM Sample† N=2934	HMDA Sample† N=14282	MSAMP3 N=2184
Loan Characteristics:			
Loan Amount ($k)	206 (100)	213 (101)	206 (102)
Property Value ($k)	218 (109)	221 (107)	215 (107)
Interest Rate (%)	4.66 (0.64)	4.59 (0.64)	4.65 (0.63)
Total Costs ($k)	. (.)	7.65 (9.54)	7.24 (3.45)
Rate Spread (%)	. (.)	1.40 (0.56)	1.44 (0.55)
LTV (%)	94.9 (6.48)	96.0 (5.31)	95.8 (4.62)
DTI (%)	41.2 (9.02)	43.3 (9.61)	40.4 (8.99)
Borrower Characteristics:			
Annual Income ($k)	. (.)	71.9 (37.4)	70.0 (36.0)
FICO Score	673 (45.9)	. (.)	675 (46.3)
Age (y)	. (.)	38.9 (11.9)	38.3 (11.9)
Ethnicity: Hispanic or Latino	.%	23.8%	24.1%
Race:			
Asian	.%	2.50%	2.00%
Black or African American	.%	15.2%	15.1%
White	.%	78.9%	80.1%
Sex:			
Female	.%	28.2%	28.3%
Joint	.%	33.7%	31.9%
Male	.%	38.1%	39.7%
Funding Characteristics:			
Purchaser Type:			
Balance Sheet	.%	41.3%	47.7%
FNMA & FDMC	.%	0.21%	0.10%
GNMA	.%	44.6%	38.3%
Source Type:			
Correspondent	33.8%	.%	37.1%
Retail	50.5%	.%	46.6%
Wholesale	15.6%	.%	16.3%

†1pp; FHA-insured; 2018-19; 1st Lien; Owner-Occupied; 1-4 Unit Dwelling; Purchase; 30-year; Fixed Rate; Vanilla; Conforming

Table 3.3: Detailed Summary Statistics for MSAMP3

	N	$\overline{x}$	s_x	p0	p25	p50	p75	p100
Loan Characteristics:								
Loan Amount ($k)	218k	206	98.6	20	137	187	257	967
Property Value ($k)	218k	215	103	20	142	195	267	1150
Interest Rate (%)	218k	4.65	0.64	2	4.25	4.62	5.12	7.12
Total Costs ($k)	208k	7.35	6.69	0	5.04	6.82	9.05	1710
Rate Spread (%)	215k	1.44	0.56	-4.96	1.05	1.38	1.78	21
LTV (%)	218k	95.9	4.41	13.9	95.5	97.5	98.2	117
DTI (%)	68.9k	41.2	8.97	1	36	42	47	95
Borrower Characteristics:								
Annual Income ($k)	218k	69.9	36.3	-40	46	63	86	990
FICO Score	191k	675	46.6	342	643	668	702	850
Age (y)	218k	38.5	11.8	20	30	40	50	80
Census Tract Characteristics:								
Population (k)	218k	5.6	2.89	0.18	3.85	5.1	6.66	53.8
Minority Population (%)	218k	36.8	27.9	0	13.5	28.9	56.1	100
Median Income ($k)	218k	71.5	24.6	0	54.1	68.4	85.3	282
1-4 Unit Homeownership Rate (%)	217k	69.3	14.8	0.31	60	71	80.3	100
Median Age of Housing Units (y)	216k	39.4	17.3	4	26	38	52	76

3.5.2 Estimator Implementation

Taking the identification strategy to the data requires that three assumptions hold. First, that the counterfactual distribution would be smooth if not for the presence of the notch. Second, that bunchers SSecond, that bunchers come from a continuous set such that there exists a well-defined marginal buncher. Third,

To calculate identify $P^o_m - P^*$, I estimate the counterfactual distribution that would have occurred in the absence of a kink. I identify each transaction's "relative price", the distance between the transaction price and the household's reference price, where the reference price is imputed for each household according to the procedures explicated above. I estimate the counterfactual distribution of relative prices by fitting a polynomial to the counts, but excluding data near the kink, with the following specification:

$$N_j = \sum_{i=0}^{q} \beta_k^0 (p_i)^k + \sum_{i=R_\ell}^{R_r} \delta_j^0 \mathbb{1}\{i = j\} + \varepsilon_j^0 \tag{3.12}$$

Above, transactions are divided into 250 bins indexed by $j \in \{-200, ..., 0, ..., 49\}$. Each bin has a width of $400 so that the relative prices used in the estimation range from -$80k to $20k. For each bin, j, the number of loans is denoted N_j and the relative price at the midpoint of the bin is denoted p_j. R_ℓ and R_r denote the index of the left and right-most bins deemed to contain bunching mass, which I set at -12 and 0, respectively. And q is the order of the polynomial, which I set to 15.

I use the predicted values from Equation (3.12) to estimate the counterfactual distribution:

$$\hat{N}_j = \sum_{k=0}^{15} \hat{\beta}_k (d_j)^k \tag{3.13}$$

I use the difference between observed and counterfactual bin counts in the bunching region, $i \in \{-12, 0\}$, to estimate the excess mass:

$$\hat{B} = \sum_{j=-12}^{0} (N_j - \hat{N}_j) = \sum_{j=-12}^{0} \hat{\delta}_j \tag{3.14}$$

I do not observe salient missing mass above the kink point, a common feature of the bunching estimator literature. Because the bunching mass in the realized distribution must derive from mass above the kink point in the counterfactual distribution, a more accurate estimate of the counterfactual distribution would increase its mass by that of the estimated excess bunching mass. This is a common feature of bunching strategies and I employ an iterative procedure to identify teh counterfactual distribution used in the literature [Bachas et al., 2020, Chetty et al., 2011, Kleven and Waseem, 2013]. Having estimated a counterfactual distribution, I inflate the mass of transactions above the bunching region by the amount of the bunching mass. In particular, having estimated $\hat{B}$ from equations (3.12) and (3.14), we estimate the following:

$$N_j * \left[1 + \mathbb{1}\{j > R_r\} \frac{\hat{B}}{\sum_{j=R_r+1}^{\infty} N_j}\right] = \sum_{i=0}^{q} \beta_k (p_i)^k + \sum_{i=R_\ell}^{R_r} \delta_j \mathbb{1}\{i = j\} + \varepsilon_j \tag{3.15}$$

This yields a new value of $\hat{B}$ which can be used again in Equation (3.15) iteratively until the value of $\hat{B}$ reaches a fixed point.

I then define the adjustment of the marginal buncher, $P_m^o - P^*$, as the ratio of the bunching mass to the average density of the counterfactual distribution in the bunching region:

$$P_m^o - P^* = \frac{\hat{B}}{\sum_{j=R_\ell}^{R_r} \hat{N}_j / (R_r - R_\ell + 1)} \tag{3.16}$$

I compute standard errors by following a non-parametric bootstrap procedure. In particular, following Bachas et al. [2020], Chetty et al. [2011], I create new bins of loans by sampling from the estimated vector ε_j and adding these to the estimated distribution. I then use the outlined procedure to compute a new estimate of the marginal buncher's adjustment. To compute the elasticity, I use the median reference price as the reference price. I define the standard error of the elasticity estimate as the standard deviation of estimates due to repeated sampling.

Because down-payment and house prices are related by a constant, I can repeat the estimation procedure with down-payments. In this specification, I consider 200 bins indexed by $j \in \{-100, ..., 0, ..., 99\}$, each with a width of \$60 so that the relative down-payments are in the range of -\$6k to \$6k. I set R_ℓ to -1 and R_r to 0. I retain a polynomial of order 15.

3.6 Results

3.6.1 Suggestive Evidence

Suggestive evidence of bunching in the cross-section of relative housing quantity comes from simple histograms of realized loan-to-value and debt-to-income ratios on loans. I plot these in Figure (3.16a) for the full HMDA FHA sample and in Figure (3.16b) for the merged sample. There is considerable bunching in both samples at the loan-to-value limit of 96.5% and some modest bunching at the lowest of debt-to-income limits of 43%.[3] Although neither form of bunching is exactly the bunching of interest, the fact that both forms show up in the full sample mitigates concerns about using a selected sample for the primary analysis.

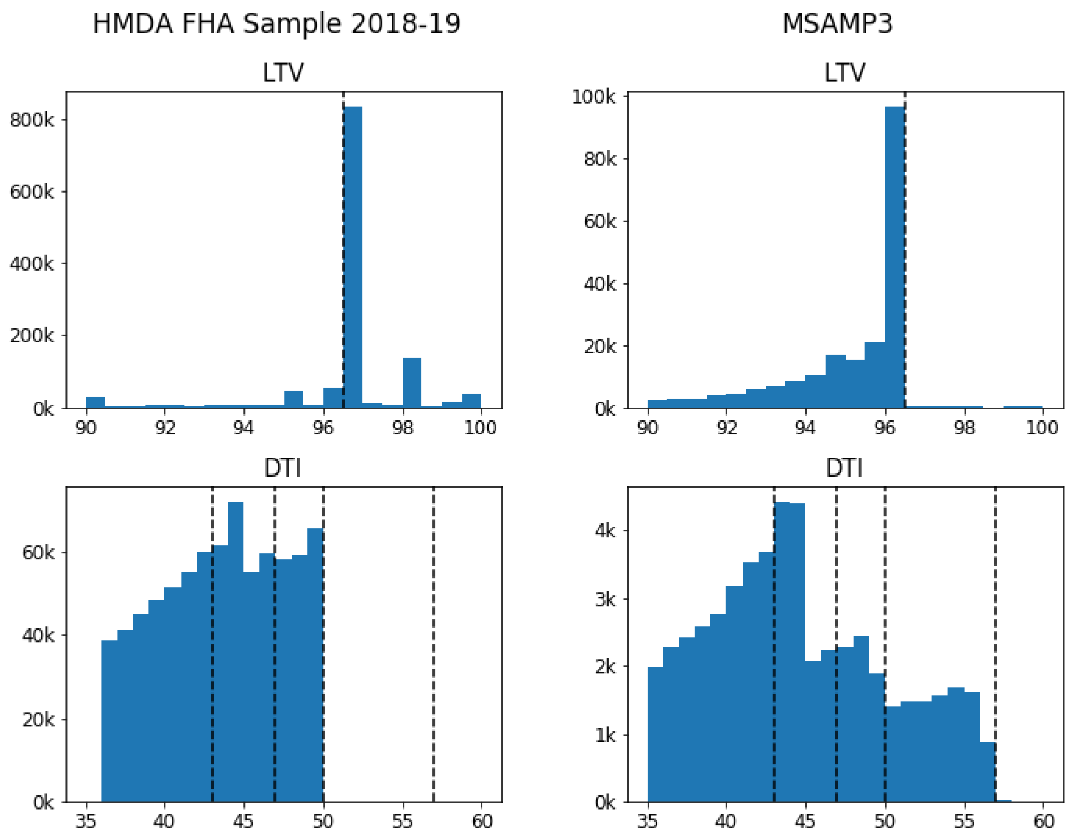

(a) (b)

Figure 3.16: Above are histograms of LTV and DTI for both the HMDA FHA sample and the merged sample. In both samples, there is bunching in both the LTV and the DTI at limit amounts, denoted in dashed black lines. The disagreement between the limits and the location of the bunching is due to adjustments made for mortgage-insurance premiums. Neither form of bunching is the bunching of interest for the purpose of the analysis but both are suggestive of the existence of such bunching. Finding comparable suggestive evidence in the HMDA sample mitigates concerns about the representativeness of the merged sample.

3. Where the bunching in debt-to-income realizations appears somewhat off the standard debt-to-income limits, this is because the FHA allows a borrower to roll the mortgage insurance premium into the balance of the loan without counting it toward the debt-to-income limit.

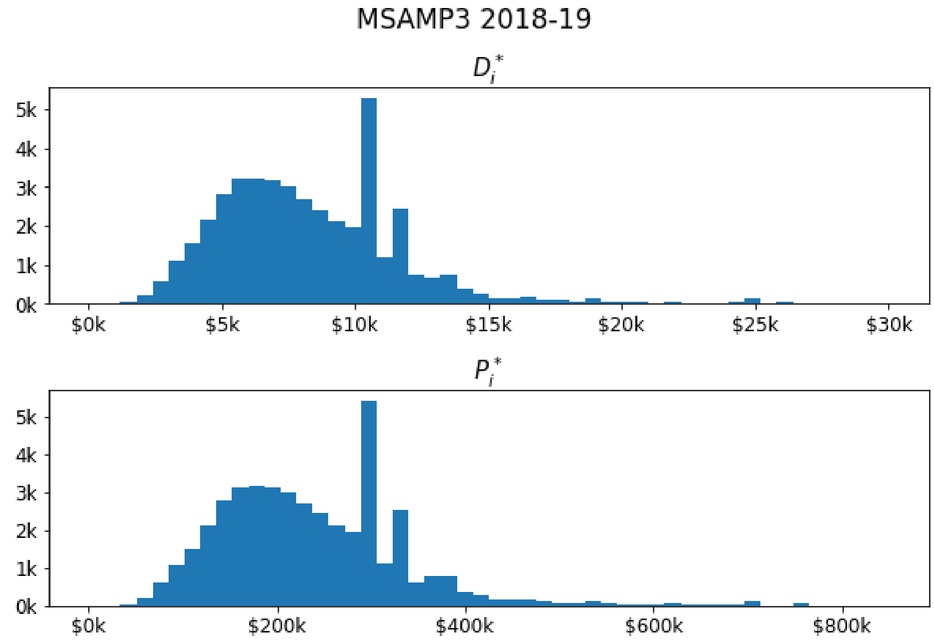

Figure 3.17: Above are histograms of the reference downpayment and house price for every transaction in the merged sample. There is considerable variation in these values, suggesting that any bunching obtained is not the result of a coincident policy. The apparent bunching in these plots is the result of county-level loan limits which are binding for borrowers of sufficiently high income in the county. This is not the bunching of interest but does serve to improve the variability of the kink point relative to borrower characteristics.

The intuition conveyed in these histograms is reiterated in Figure (3.18). In this figure, I plot the relative down-payments and house prices for transactions in MSAMP3 that take place in Cook County, Illinois. The density of transactions near an apparent left and upper boundary comports with the nature of the household problem described in the section on the empirical approach. The transactions near the left-most boundary are a depiction of bunching in the LTV constraint; transactions near the upper boundary are a depiction of bunching in the DTI constraint. The bunching of interest in our estimator is excess density where the two constraints intersect.

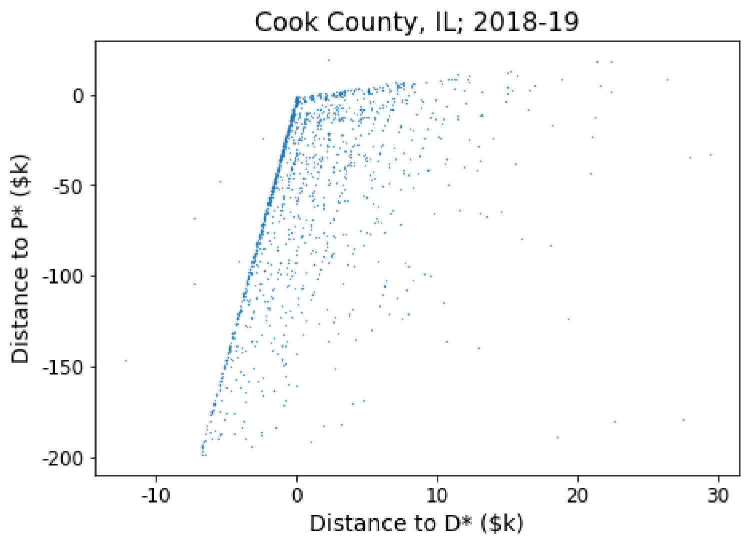

Figure 3.18: Above are transactions from the merged sample in Cook County, IL, plotted in downpayment-house price space and normalized to the reference downpayment and house price. The household choice set becomes apparent empirically from this plot and resembles the choice set established in the household problem. The bunching in LTV and DTI are apparent in the relative frequency of borrowers along the left and upper edges of the choice set, respectively. The bunching of interest in this paper concerns the relative frequency of borrowers at the origin, specifically.

3.6.2 Loan-to-value Elasticities of Housing Demand

More direct primary evidence comes from Figure (3.19), which plots down-payments and house prices relative to the reference amount.[4] Note that the upper and lower panel are really the result of collapsing a full-sample version of Figure (3.18) onto the x- and y-axes, respectively. The bunching at the reference amount suggests that households are, in fact, credit constrained and therefore rely on financing terms available through the FHA for marginal units of housing purchase.

Figure (3.17) plots the distribution of reference down-payments and housing prices for the individuals whose transactions are plotted in Figure (3.19). It is no coincidence that the upper and lower panels look nearly identical; recall that the reference down-payment is the reference price scaled by a function of the loan-to-value limit, and this loan-to-value limit is

4. These plots can only be constructed for the merged sample and so sample selection concerns remain.

identical for nearly all FHA transactions. The presence of bunching in this plot is due to the presence of absolute county initial balance limits, which bind instead of DTI limits for all individuals of sufficiently high income in a county. This plot provides evidence that the reference points aligned to generate Figure (3.19) are heterogeneous due to both differences in income and geography. It allays concerns that the documented bunching is due to some coincident threshold.

Implementation of the bunching estimator is depicted in Figure (3.20). The dotted black lines designate the assigned area of bunching, the orange line depicts the counterfactual distribution, and the red line depicts the counterfactual location of the marginal buncher. I use the median reference price and down-payment in the formula for the bunching estimator. Using down-payment amounts, I obtain an elasticity estimate of 0.025 ± 0.007; using house price amounts, I obtain 0.014 ± 0.004.

These suggest fairly small adjustments to housing demand in response to changes in credit availability. For a hypothetical household formerly bound by a debt-to-income requirement who experiences a relaxation of that requirement, this estimate suggests they would only increase their housing demand by 2.5%. This effect, however, is restricted to households bound by the debt-to-income requirements. Hypothetical households bound by loan-to-value requirements experiencing a relaxation of this requirement of 18% (when mortgage financing becomes more widely available at 95% rather than 80% LTV), would only increase their housing demand by $0.18 * 2.5\% \approx 0.5\%$.[5]

As a final exercise, I test whether the estimate varies in the distribution of incomes. I divide households into groups by whether their annual income at the time of origination was in $[\$0k, \$40k)$, $[\$40k, \$60k)$, $[\$60k, \$80k)$, or $[\$80k, \$\inf)$. I plot the distribution of relative down-payments in Figure (3.21). The evidence suggests some consistency with the idea that higher income households are less credit constrained. Implementing the bunching estimator

5. This latter number could change substantially if considering the quasi-income effects of additional loan balance becoming available.

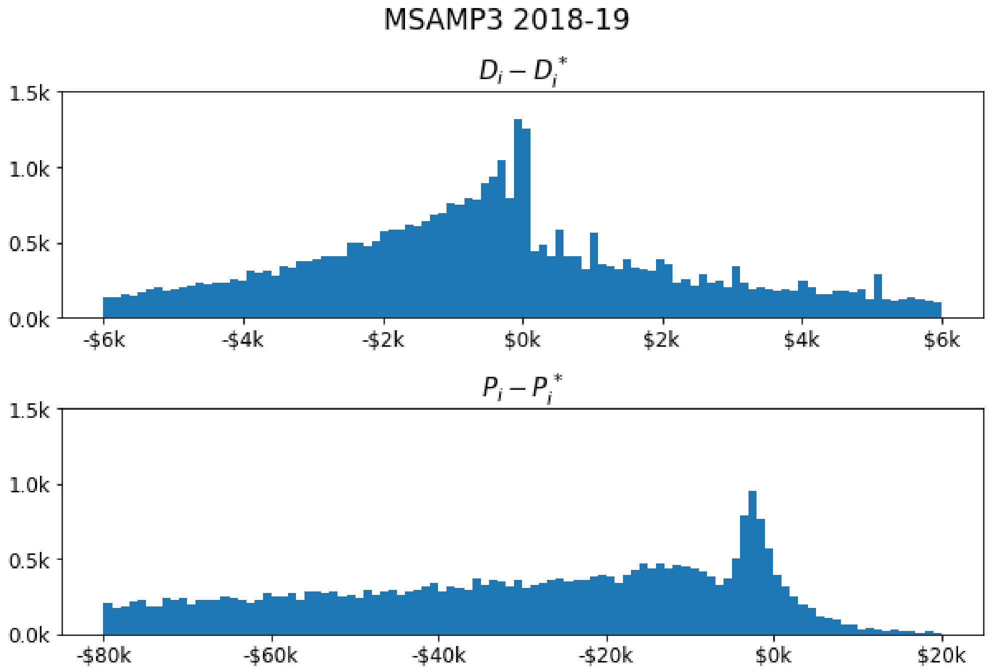

Figure 3.19: Primary evidence of bunching in both downpayment and house prices. Note that both measures map to housing quantity according to $P_i = p_H H_i$ and $D_i = (1 - \overline{L}_i)P_i$.

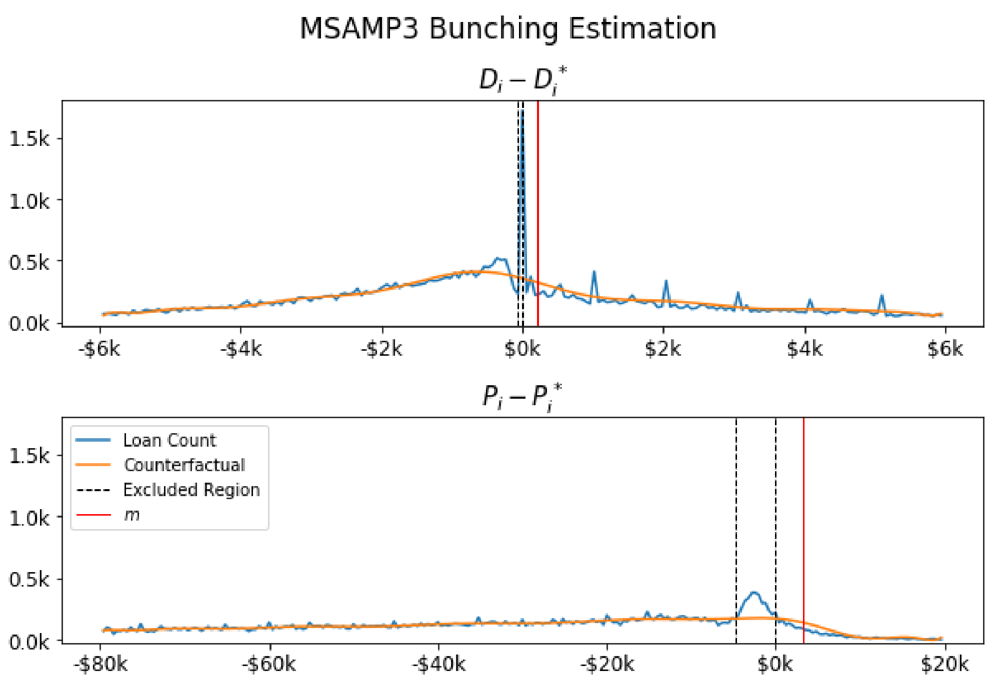

Figure 3.20: The bunching estimator implemented on the data in both downpayments and house prices. The black dashed lines designate the area in which the bunching is determined to be present. The red line designates the counterfactual behavior of the bunching agent, as measured by the bunching estimator. Following the literature, an iterative procedure is used to back out the counterfactual density, depicted in orange.

within groups confirms this in the estimated values, which are .031, .028, .022, and .022, in respective income bins. However, the standard errors on these values are too large to distinguish them statistically. Moreover, the variation in the magnitude is does not appear to be sufficiently large to explain differing housing returns in different markets.

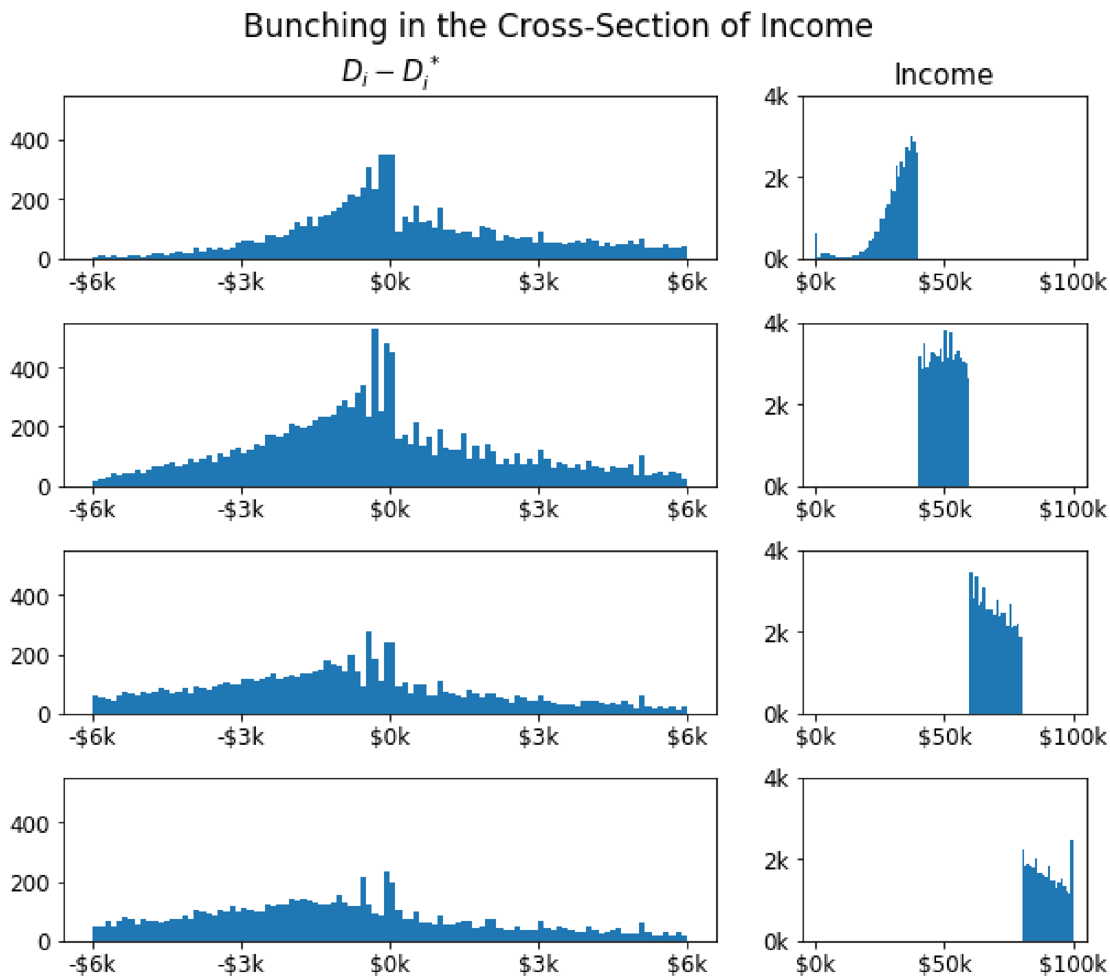

Figure 3.21: This figure depicts the pattern bunching relative to the reference house price in the cross section of household income. The figures on the left depict the bunching and the figures on the right depict the annual income for the households on the left. I consider income buckets \$0-\$40k, \$40k-\$60k, \$60k-\$80k, and \$80k-\$100k. There is evidence that the bunching is declining in household income.

3.7 Conclusion

In this paper, I develop a new estimator of the loan-to-value limit elasticity of housing demand. This estimator has the benefit that it captures only partial equilibrium effects and is therefore more plausibly excludes the confounding effects of changing house price expectations. In addition to providing extensive analysis on how the estimator is constructed and interpreted, I apply it to a sample of FHA loans from recent years. While I find evidence that credit constraints play a role in limiting these households' access to housing on the intensive margin, the magnitude appears to be somewhat limited.

CHAPTER 4

HOUSING WEALTH MANAGEMENT AT RETIREMENT

4.1 Introduction

American households (HHs) tend to have inadequate savings to finance consumption in retirement. The median net worth of HHs ages 51-56 in 2004 was only $152k. Moreover, retirement consumption is financed out of private pensions, housing equity, and social security (SS) entitlements; were it not for SS benefits, 39pp rather than 9pp of elderly HHs would have lived in poverty as of 2017 [Sommer, 2019]. This may be due to misperceptions about working ability in retirement as well as myopia plays a dominant role in HH retirement-saving decision-making [Thaler and Benartzi, 2004].

Ensuring adequate financing for retirement consumption is becoming both more important and more difficult as the baby-boomer generation ages. It is projected that by 2035, the number of Americans 65 or older will rise to 79m from 49m [Sommer, 2019] and that one in three households will be headed by someone sixty-five or older [Joint Center for Housing Studies, 2014]. Unless modifications to the program are made, SS will have to begin drawing down assets to pay promised benefits in the year 2020. If SS funds were to be depleted, benefits would be cut by 20% [Sommer, 2019].

Housing equity comprises a significant fraction of American HH's retirement savings. Among HHs ages 70-75, about 80pp own a home and 25-30pp of retired HH wealth is in housing [Davidoff, 2009]. To understand how HHs finance their retirement consumption requires understanding how they spend out of their housing wealth. Although there has been some work on how HHs spend out of their housing wealth during retirement, to the best of my knowledge there is no work on how HHs tap into their housing wealth at the beginning of retirement.

This paper documents that retirement from the labor force prompts HHs to liquidate housing wealth. By using SS eligibility thresholds as an instrument for retirement, I show that exiting the labor force prompts 12pp of HHs to issue originate first mortgage debt and 3pp of HHs to extract equity from their home within two years. I document this effect in the Survey of Consumer Finances and corroborate the results in the Health and Retirement Study.

I show further that liquidation of housing wealth at retirement is associated with a rise in liquid balances that is larger than liquidation not prompted by retirement. This suggests that housing wealth liquidation at the retirement threshold is qualitatively different from liquidation at other points in the life-cycle. The regression results are consistent with ordinary liquidation being prompted by adverse expense shocks and retirement liquidation being pro-active, i.e. anticipating the need for liquid assets.

I consider a variety of reasons that HHs may have for liquidating or borrowing against their housing wealth at retirement. In particular, I consider that they may be moving because they value geography differently, they may be smoothing consumption, they may have relaxed credit constraints at the threshold, and they may be consolidating debt to simplify their finances. The evidence is mixed and multiple channels may be operative.

4.1.1 Literature

There is a vast literature on the effects of Social Security in the United States. Among the outcome variable considered are private savings [Feldstein, 1974], retirement [Mastrobuoni, 2009], labor supply [Liebman et al., 2009], and mortality [Fitzpatrick and Moore, 2018]. I add a different outcome variable to this literature: the timing of new mortgage debt acquisition is driven by Social Security program parameters, both as households move and consume out of home equity.

The literature on the uses of housing wealth in retirement emphasizes that, though it is a significant component of household wealth, households tend not to consume out of it. Nakajima and Telyukova [2011] finds that renters in retirement tend to run down their wealth more quickly than do homeowners. Poterba et al. [2011] documents shocks to family structure, like divorce and death, as an important correlate of reduction in housing equity. Fisher et al. [2007] finds that older Americans are increasingly mortgage free to age 80. Davidoff [2009] suggests that housing wealth may be used as a hedge against longevity risk. These trends may be true on average but we document that at the retirement threshold, households increasingly borrow against their homes.

The literature describes several determinants of refinancing. HHs may take advantage of falling interest rates [Quigley, 1987]. Structural changes in the mortgage market over time may make it less costly to access wealth in homes [Bennett et al., 2001]. The effects documented in this paper are most similar to Hurst and Stafford [2004], which notes that households refinance to smooth consumption in periods where they suffer a large drop in income. The drops in income studied there, however, are not necessarily anticipated, whereas social-security induced retirement creates predictable drops in income.

A strand of the retirement finance literature investigates the "retirement consumption puzzle". Early work suggested that the drop in household consumption expenditure at retirement might deviate from the predictions of the permanent-income hypothesis [Bernheim, 1987, Banks et al., 1998]. More recent work documents that it is a drop in consumption expenditure rather than consumption itself, and that this drop can be attributed to bargain shopping induced by reduced opportunity cost of time at retirement [Aguiar and Hurst, 2005, 2007]. Very recent work attempts to resurrect non-rational explanations of retirement savings by identifying a drop in consumption expenditure as well as a rise in savings balances at retirement [Olafsson and Pagel, 2018].

Taken together, this literature debates whether a drop in consumption expenditure at retirement is evidence that households save too little relative to the permanent-income hypothesis, a rational benchmark. By contrast, this paper documents that households make predictable and expensive divestment decisions at the retirement threshold. This raises the possibility, which I consider briefly in closing, that households are storing too much of their retirement savings in housing relative to a rational benchmark.

4.1.2 Institutional Background

Social Security is an insurance program founded in 1935 and run by the Social Security Administration (SSA), a federal agency. The program taxes earnings of working individuals to fund benefit payments, most prominently in the form of annuity distributions to retirees. The program operates on a pay-as-you-go basis, so that current payments are used to fund current benefit payments.

The amount of a worker benefit is a function of the worker's earnings over his or her lifetime. In particular, a worker's 35 highest annual earnings (after index adjustments, including zeros for years of non-employment, and truncated at maximum taxable earnings) are averaged to compute the worker's average indexed monthly earnings (AIME). The worker's primary insurance amount (PIA) is then computed as a progressive transformation of the AIME. The worker benefit amount is then a function of the worker's PIA, depending on when the recipient chooses to receive the benefit. [Liebman et al., 2009]

The worker receives the full PIA if they claim at the full retirement age (FRA). The FRA is a program parameter that is 65 for individuals born in 1937 or before, rises two months for every year born thereafter until it reaches 66 for individuals born between 1943 and 1954, then rising two months for every year again until it reaches 67 for individuals born in or after 1960. [Social Security Administration]

Workers need not claim their SS benefits at the FRA. They may claim benefits as early as age 62. For each year before the full retirement age, they receive a permanent reduction of 6.66pp in the benefit amount, pro-rated by month. Alternately, workers may delay their claim. For each year of delay up to the age of 70, they receive a permanent increase of 5pp in the benefit amount, pro-rated by month.

In addition to worker benefits, SS also offers benefits to spouses and survivors. Spouses may claim a benefit equal to the greater of their own worker benefit and 50pp of their spouse's benefit. Widowed spouses may claim a benefit equal to the greater of their own worker benefit and 100pp of their spouse's benefit. Widowed spouses may claim as soon as age 60.

4.2 Data and Empirical Strategy

4.2.1 Data

This paper uses data from the Survey of Consumer Finances (SCF), the Health and Retirement Study (HRS), and TransUnion (TU) consumer credit records. Below, I describe each dataset and the samples I construct for the analyses conducted in this paper.

The Survey of Consumer Finances (SCF)

The SCF is a pooled survey of a random sample of American households. It contains information on demographics, finances, and financial institutional affiliation. The survey has been administered every three years since 1983. The SCF assigns the designation of HH head to the male in a different-sex couple and older partner in a same-sex couple, a convention I adopt for convenience purposes. The interviews contain information on the birth year of the HH head, the year the HH head entered retirement, and the last year at which the HH issued mortgage debt at the time of interview.

I construct a "full sample" from the SCF by pooling the 1992-2016 cross sections. Summary statistics for the HHs (after adjusting for multiple implicates) in the "full sample" are presented in Table (4.1). These statistics confirm that the sample is roughly representative of all HHs in the United States. An odd feature that bears more scrutiny is the fact that many HH heads appear to claim SS outside the permissible ages. This may be a result of misreporting or imputation error by the survey administrators.

Table 4.1: Summary Statistics for SCF "Full" Sample

	N	$\overline{x}$	s_x	p25	p50	p75
Interview Characteristics:						
Interview Date	45k	2004	7.8	1998	2004	2010
Age (Interview)	45k	51.0	16.2	39	50	63
Interviewee Characteristics:						
1{Married}	45k	0.6	0.5	0	1	1
1{Homeowner}	45k	0.7	0.5	0	1	1
Income ($)	45k	641.6k	4.7m	25.0k	55.0k	140.0k
Liq. Assets ($)	45k	191.2k	1.8m	900	5.6k	33.0k
Checking ($)	45k	78.3k	0.8m	500	2.2k	11.1k
1{Dir. Dep.}	41k	0.7	0.5	0	1	1
# Accounts	45k	2.1	1.2	1	2	3
Age (Soc. Sec.)	10k	63.2	15.1	53	67	72

I construct a "retirement sample" from the SCF by subsetting the "full sample" on those HHs who have claimed SS benefits and are retired. Only with this information can I construct the distance of latest mortgage issuance to SS benefit claim dates. Summary statistics for the "retirement sample" are reported in Table (4.2). These confirm that the sample is representative of older Americans. Of the 7.0k HHs meeting the criteria, 21% report a date of most recent first mortgage origination and 4% report a date of most recent equity extraction.

Table 4.2: Summary Statistics for SCF "Retirement" Sample

	N	$\overline{x}$	s_x	p25	p50	p75
Interview Characteristics:						
Interview Date	7.0k	2006	8	1998	2007	2013
Age (Interview)	7.0k	73	8	67	72	79
Interviewee Characteristics:						
Birth Date	7.0k	1932	11	1924	1932	1941
1{Married}	7.0k	0.5	0.5	0	1	1
1{Homeowner}	7.0k	0.8	0.4	1	1	1
# Accounts	7.0k	2.0	1.1	1	2	3
1{Dir. Deposit}	6.3k	0.9	0.3	1.0	1.0	1.0
Income ($/Yr.)	7.0k	303k	2.1m	18k	35k	80k
Retirement Characteristics:						
Age (Retire)	7.0k	63	6	59	62	66
Soc. Sec. Date	7.0k	1996	11	1989	1996	2004
Age (New 1st)	1.5k	63	9	58	64	69
Age (Extract)	307	65	9	60	64	70

The Health and Retirement Study (HRS)

The HRS is a panel survey of a random sample of American households with one member above the age of 50. It asks respondents detailed questions about demographics, life-cycle events, health status, and finances. The survey has been administered every two years since 1992; a new cohort is added approximately every six years. I use the RAND Longitudinal File 1992-2014 to construct balance sheet and income statement measures for each household. I merge additional information on mortgage debt and housing wealth management from the RAND Fat Files 1992-2014.

I construct a "crude sample" from the HRS by pooling data in every wave from 1992-2014. I subset on HHs that do not change family structure during all available interview periods, i.e. HHs with SUBHH equal to zero. Table (4.3) presents summary statistics for this sample of HHs, where demographic characteristics are assigned according to the first financial respondent on record. They are representative of older Americans above the age of 50 in the years since 1992. A low proportion have claimed SS benefits in part because many HHs are not yet of claiming age. Some HHs seem to report having claimed before the age of eligibility, but this is potentially an artifact of the fact that HH birth year is constructed according to the first financial respondent in the data but HH benefit claim date is constructed by the earliest claim date for couples.

Table 4.3: Summary Statistics for HRS "Crude" Sample

	N	$\overline{x}$	s_x	p25	p50	p75
Interviewee Characteristics:						
Birth Date	22k	1936	15.34	1925	1937	1950
1{Married}	22k	0.55	0.50	0	1	1
# Children	22k	2.69	2.14	1	2	4
1{Hispanic}	22k	0.11	0.31	0	0	0
1{Black}	22k	0.20	0.40	0	0	0
1{High Sch. Grad}	22k	0.72	0.45	0	1	1
1{College Grad}	22k	0.18	0.38	0	0	0
Retirement Characteristics:						
1{Claimed SS}	22k	0.51	0.50	0	1	1
Age (Soc. Sec.)	11.5k	58.89	7.74	56.17	61.92	63.08

I construct a "refined sample" from the HRS by again pooling data in every wave from 1992-2014. In the "refined sample", I aim to include HHs with changing family structure. Post-divorce branches of HHs are included as distinct from each other and from their antecedent pre-divorce HHs. Allowing for HHs with changing family structure increases the count of HHs somewhat, to ~23k, but does not much change the demographic composition of the sample.

TransUnion (TU)

The TU data are records of outstanding consumer debt as reported by consumer lenders. TU supplies a 5% random sample of their records to Booth. For a simple analysis of credit scores at the end of the paper, I draw a random sample of 100k consumers from the TU files in the month of June 2012.

4.2.2 Empirical Strategy

I estimate the causal effect of retirement on HH propensity to originate mortgage debt. Correlations in the timing of retirement and new mortgages may be determined endogenously, thought the direction of the bias is unclear ex-ante. In the event of an adverse medical shock, a HH may borrow against their home to pay the medical bills. If the health shock is debilitating, the HH may be forced into retirement at the same time and the OLS estimates will be biased upward; if the health shock is expensive but not debilitating, the HH may continue to work to help pay the associated bills and the OLS estimates will be biased downward. To address endogeneity concerns, my main specification instruments for retirement with SS eligibility thresholds to isolate the effects of exiting the labor force. Specifically, I run the following IV regression:

$$Y_{it} = \beta_0 + \beta_1 * X_{it} + \beta_2' Z_i + \varepsilon_{it} \tag{4.1}$$

$$X_{it} = \gamma_0 + \gamma_1 * 1_{it}\{SS\} + \gamma_2' Z_i + \eta_{it} \tag{4.2}$$

For HH i in year t, the main outcome is an indicator either that the HH originated any new first mortgage, $Y_{it} \equiv 1_{it}\{New\ 1st\}$. I also consider other housing wealth management decisions, such as the purchase or sale of the home, equity extractions, origination of a new second mortgage, and origination of a HELOC. The regressor of interest is an indicator of HH retirement, $X_{it} \equiv 1_{it}\{Retire\}$, defined as retirement of the man in a mixed-sex HH or the older individual in a same-sex HH.[1] Where available in the HRS data, I include FEs, Z_i, for gender, race, education, marital status, number of children, and birth cohort. After instrumenting for SS eligibility, described below, I predict that retirement will induce new

1. This definition is for convenience purposes and follows the SCF designation of the HH head

mortgage debt origination, $\beta_1 > 0$.

The instrument, $1_{it}\{SS\}$, is an indicator of whether the HH head crosses a significant SS eligibility threshold in that year. Specifically, the indicator takes the value of one in the year the HH head first qualifies for SS benefits, the year they reach the FRA, and the latest year at which they can claim benefits. Unless widowed, a possibility we ignore, individuals first qualify for SS at 62. The FRA is 65 for those born before 1942 and 66 for those born between 1943 and 1959.[2] Age 70 is the latest an individual can claim the benefits.

My first identifying assumptions is that SS eligibility thresholds are relevant for retirement, $Cov(1_{it}\{Retire\}, 1_{it}\{SS\}) > 0$. It is widely known that SS thresholds, especially qualification for eligibility at 62, are relevant for retirement. I confirm the age distribution of SS claims in Figure (4.1) and the age distribution of HH head retirement in the SCF data in Figure (4.2).

The second identifying assumption is that SS eligibility thresholds do not affect mortgage acquisition or equity extraction except through the effects of retirement, $Cov(\varepsilon_{it}, 1_{it}\{SS\}) = 0$. The SS program parameters are plausibly exogenous because they are legislatively determined and outside the control of the HHs considered. In principle, it is possible that exogeneity is violated because SS eligibility is correlated with other institutional affiliations that might drive housing refinance decisions. However, to the best of my knowledge, no other federal program uses age 62 as an eligibility threshold. Moreover, the use of variation in cohort FRA mitigates concerns about changes taking place at age 65, specifically. And the most important contemporaneous change at age 65 is qualification for Medicare; but to the extent that health insurance buffers against the financial effects of adverse health shocks, this should reduce the propensity of HHs to borrow against their home and work against the effects documented in this paper.

To better understand the economics of refinancing at retirement, I analyze the effect of the refinancing decision on changes in HH liquid balances. There are similar endogeneity

2. It rises to 67 for those born in 1960 onward but there are no such individuals in our sample

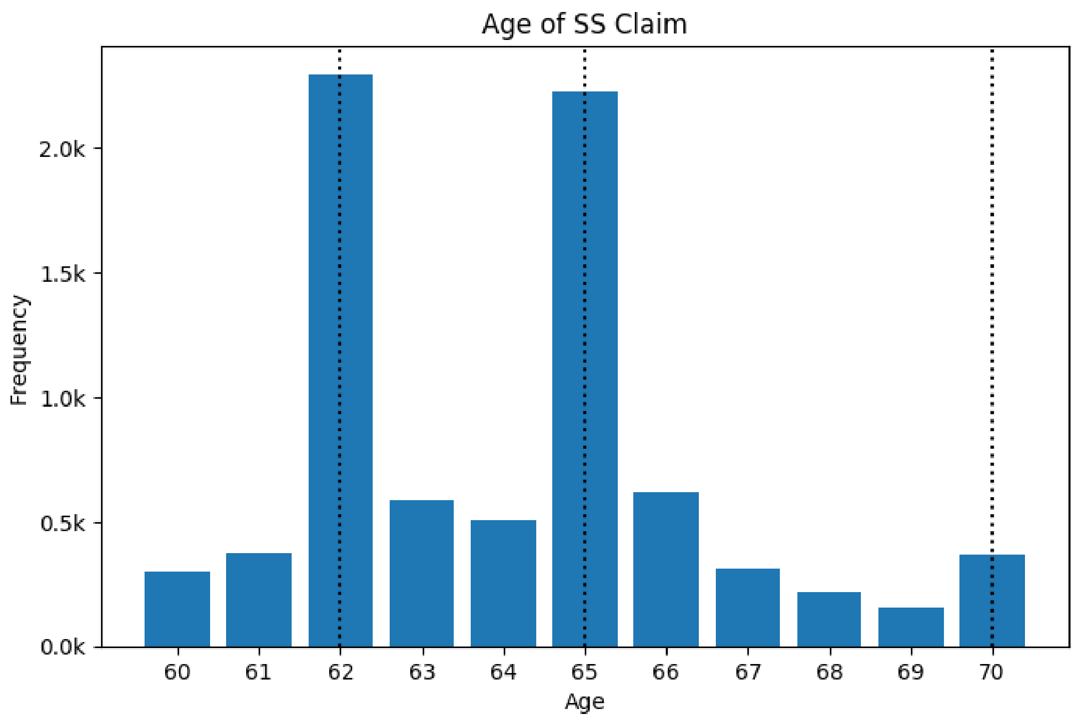

Figure 4.1: The above are counts of social security claim ages in the SCF 1992-2016 data and adjusted for implicates. These data are roughly consistent with aggregate measures of when social security is claimed.

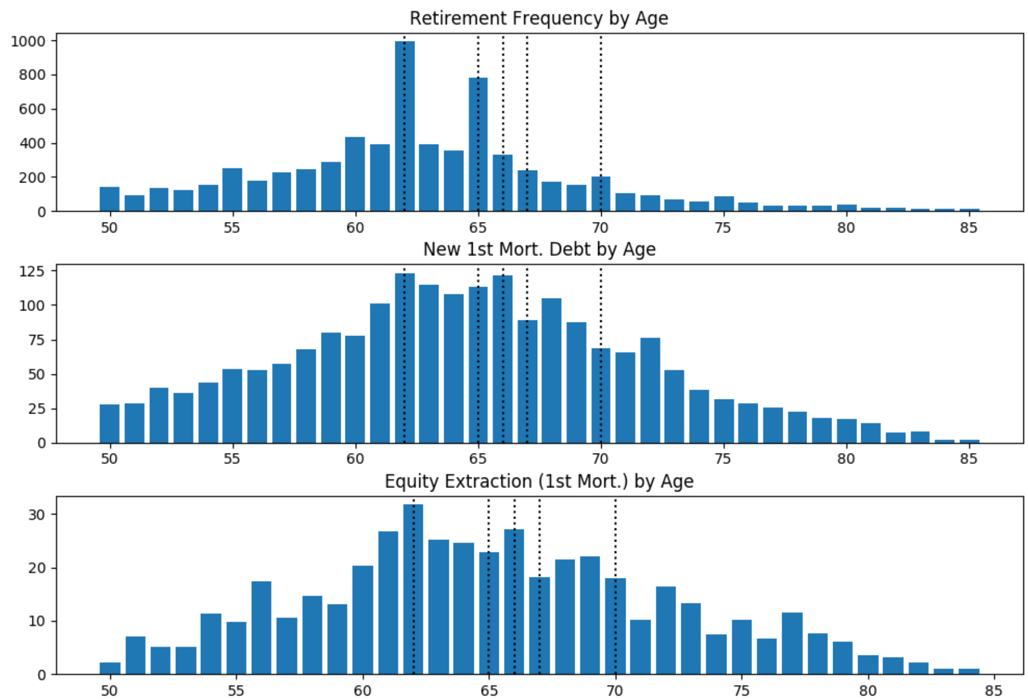

Figure 4.2: The above are counts of retirement and new mortgage obligations by age in the SCF 1992-2016 data and adjusted for implicates. These graphics depict the IV specification visually. The instrument is approximated by the dashed vertical lines denoting the ages in which the instrument applies. The uppermost panel represents the first stage and shows that retirement is induced by social security eligibility. The middle and lower panels represent the reduced form for all mortgage debt and equity extraction.

concerns, like the possibility that an adverse health shock would prompt HHs to spend down liquid savings and originate mortgage debt to spend out of housing wealth. To identify the effects on liquid balances of those HHs prompted to refinance by retirement, I use a similar IV specification.

In this latter case, for each HH i in period t, the outcome is a measure of the change in HH liquid balances, $Y_{it} \equiv \Delta Sav_{it}$. The regressor of interest is an indicator that the HH extracted equity, $X_{it} \equiv 1_{it}\{Extract\}$. I also consider sale of the home and origination of second mortgages. Again I include demographic FEs, Z_i. And I predict that after instrumenting for SS eligibility, refinancing will induce a rise in savings balances, $\beta_1 > 0$.

4.3 Results I: Retirement and New Mortgage Debt

I find evidence that retirement induces HHs to take on new mortgage debt, both to finance new home purchases and against current housing. Retirement makes a HH 2-12% more likely to issue any new first mortgage and 2.6-8.5% more likely to extract equity from the home. I begin by documenting a spike in mortgage origination activity at the same time that HHs claim SS benefits in both SCF and HRS data. Then I use SS eligibility thresholds to construct an instrument for retirement and estimate a causal effect of retirement. I include robustness checks for different regression specifications in the appendix.

4.3.1 Event Study

I use the SCF "retirement sample" to conduct an event study of housing wealth management at the time HHs first claim SS benefits. In the SCF data, I select HHs in which the HH head reports having claimed SS benefits. Then I compute the time, relative to their claim of SS benefits, at which they report last having originated mortgage debt, $t_i^{M.Debt} - t_i^{SS}$. I count the frequency of new mortgage debt origination by years from SS claims. I report these counts for first mortgages, second mortgages, and other home loans; I decompose first mortgage originations into new home purchases, equity extractions, and rate refinances.

The results are reported in Figure (4.4). In each panel, there are more new mortgage originations in the year the HH head first claimed SS benefits relative to the two adjacent years. New first mortgages are the most common type of reported loan, but the increase in mortgage debt origination at the SS claim date relative to adjacent years is ~25%-33% for all categories. Many of the first mortgage originations appear to finance the purchase of new homes, but ~20% are equity extractions.

I confirm these results in the HRS "refined sample". For comparability with the SCF, I assign a HH head as being the first reported financial respondent if male and their partner if female. Then I assign each HH the SS claim date, t^{SS}, of the HH head. Next, I merge data from the RAND Fat Files 1992-2014 to this sample. I compute the time, relative to the claim of SS benefits, at which HHs report having taken some action, A, to manage their housing wealth, $t_i^A - t_i^{SS}$. I count the frequency of housing wealth management decisions by years from SS claims. I report counts of home purchases, sales, equity extractions, second mortgage originations, and HELOC originations.[3]

The results are reported in Figure (4.5). Again, in each panel, there is an uptick of housing wealth management in the same year that the HH head first claimed SS benefits relative to the adjacent years. Home purchases, analogous to first mortgage origination for new housing purchases in the SCF, still comprise ~20% of debt issuance at the SS threshold. Borrowing against existing housing in the form of equity extractions, second mortgage originations, and HELOC originations account for more of the housing wealth management decisions at the threshold, though these measures may suffer from imputation error. Again, in all the categories, actions in the year of SS claims increase by ~10-50% relative to adjacent years.

3. Since the 1996 wave, HHs have been asked about the timing of home purchases and sales. I impute equity extractions as a first mortgage balance that has increased since the previous wave and no purchase or sale of housing in the interim. I impute second mortgage originations as a positive second mortgage balance in a wave subsequent to one with no second mortgage balance and no purchase or sale of housing in the interim. I impute a HELOC origination as the presence of a HELOC in a wave subsequent to one with no reported HELOC. I cannot date these latter actions, so I assign the second interview date as a noisy approximation, $t_i^A \approx t_i^{Int.}$.

The HRS results address concerns about selection in the SCF due to the fact that only the most recent issuance of mortgage debt at the time of interview is reported. Because the HRS is a panel, repeated issuances of debt can be observed. There may, however, still be attrition from the sample in years more distant from SS claims, as HHs are not yet eligible or may become deceased. For this reason, we focus specifically on the spike at the SS claim date relative to adjacent years.

4.3.2 IV Evidence

SCF

I formally test for the effect of retirement on origination of new mortgage debt with the instrumental variable strategy described in Section (4.2.2). Specifically, using the SCF "retirement sample", I construct a synthetic panel. For each HH, I construct a yearly time series from age 55 to age 75. For each year and for different kinds of mortgage debt, I construct indicators of whether the HH originated new mortgage debt in that or the subsequent year. I construct an indicator of retirement, $1\{Retire\}$, using the reported retirement date in the SCF. Finally, I construct the instrument as an indicator of SS eligibility thresholds.

This IV is presented visually in Figure (4.2). The uppermost panel plots the first stage for both specifications, the frequencies of retirement by age. The vertical dotted lines depict the ages at which the instrument may take a value of one. This plot demonstrates clearly that SS eligibility satisfies the relevance assumption. In general, the distribution of retirement ages is plausible, centered in the 60s but with tails above and below. The middle and lower panel depict the reduced form regressions. The most common age at which people report having last originated mortgage debt and extracted equity is 62, and there are noteworthy peaks at 66 as well. This consistent with the observed ages of SS claim and the observation of elevated mortgage debt origination frequencies at SS claim. Taken together these figures suggest that retirement is driving the origination of new mortgage debt.

These pictures are consistent with a causal link between retirement and refinancing but raise some selection concerns. In particular, the high refinancing rates observed in the early 60s relative to the 50s and 80s may be driven by attrition but may bias the estimate of SS claims upward. I truncate the sample between the ages of 55 and 75 to address this concern. In the appendix, I demonstrate that results are similar when further truncating between the ages of 59 and 71. I can also flexibly control for age as a way to parametrically address the selection concerns.

The regression results are presented in Table (4.4). Column (1) demonstrates that SS eligibility thresholds are relevant for retirement and that the instrument is strong, with a significant F-test value of over 1k. In general, the IV estimates are higher than the OLS estimates. This suggests that the endogenous relationship between housing debt and retirement is consistent with the story of working longer and consuming housing wealth to finance adverse expense shocks. Retirement induces a 12.1% increase in the likelihood that HH issues new first mortgage debt within two years. It increases the likelihood that a HH extracts equity by 2.6% and increases the likelihood that the HH adjusts their mortgage rate by 2.0%. I cannot reject the null of no effect on propensity to issue debt to purchase new housing or propensity to issue a second mortgage.

Table 4.4: Housing Wealth Management at Retirement (SCF)

	Dependent variable:										
	1{Retire}	1{New 1st}		1{Buy}		1{Extract}		1{Rate}		1{New 2nd}	
	Stage 1	*OLS*	*IV*	*OLS*	*IV*	*OLS*	*IV*	*OLS*	*IV*	*OLS*	*IV*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
1{Retire}		0.013***	0.121***	0.006***	0.014	0.001	0.026**	0.003*	0.020*	0.00003	0.002
		(0.003)	(0.035)	(0.001)	(0.013)	(0.001)	(0.011)	(0.001)	(0.012)	(0.0003)	(0.004)
1{SS}	0.040***										
	(0.003)										
1	0.03***	0.02***	0.01***	0.00***	0.00***	0.00***	0.00**	0.00***	0.00***	0.00***	0.00**
	(0.001)	(0.002)	(0.001)	(0.0002)	(0.0004)	(0.0003)	(0.0004)	(0.0005)	(0.001)	(0.0001)	(0.0002)
N	199k	199k	199k	199k	199k	199k	199k	199k	199k	199k	199k
R^2	0.006	0.0004	−0.024	0.0004	−0.002	0.00003	−0.010	0.0001	−0.002	0.00000	−0.001
Adj. R^2	0.006	0.0004	−0.024	0.0004	−0.002	0.00003	−0.010	0.0001	−0.002	−0.00000	−0.001
F Stat.	1,184***	73.0***		73.5***		6.1**		13.5***		0.02	

Note: *p<0.1; **p<0.05; ***p<0.01
SEs computed with SCF bootstrap weights, adjusted for SCF multiple implicates, and clustered by person and year. Age between 55 and 75.

HRS

I confirm the results of the IV approach with data from the HRS, which helps address potential selection concerns. Using the HRS "refined sample", I again construct a synthetic panel. For each HH, I construct a monthly time series from two years (the look back period) before their first interview to the date of their last interview. For each month and for different kinds of housing wealth management actions, I construct indicators of whether the HH took the action in a given month. I assign a HH head analogously to the SCF designation, as the financial respondent if male and the spouse if the financial respondent is female in coupled HHs. I construct an indicator of retirement, $1\{Retire\}$, as the earliest reported date of retirement by the HH head. Finally, I construct the instrument as an indicator of SS eligibility thresholds.

The regression results are presented in Table (4.5). Broadly speaking, these results corroborate the results in the SCF. In column (1), SS eligibility is shown to be relevant for retirement and the F-statistic demonstrates the instrument is strong. That it is somewhat weaker than in the case of the SCF may have to do with the very conservative monthly frequency I use for the HRS data. Again, the IV estimates of the effect of retirement are larger than the OLS estimates. Retirement increases the likelihood of extracting equity from one's home in the same month by 8.5% and the likelihood of selling a home in the same month by 7.7%. I cannot reject the null for home purchases, new second mortgages, or new HELOCs, but this may have to do with my choice to look at actions contemporaneous with retirement.

The regression results should be interpreted with caution. The synthetic panel consisted of many monthly observations and for computational reasons, I have not yet been able to cluster standard errors. Clustering did not much change the SEs in other other analyses I did with these HRS data, but the statistical significance of the results could be exaggerated.

Table 4.5: Housing Wealth Management at Retirement (HRS)

	Dependent variable:										
	1{Retire}	1{Sell}		1{Buy}		1{Extract}		1{New 2nd}		1{New HELOC}	
	Stage 1	*OLS*	*IV*	*OLS*	*IV*	*OLS*	*IV*	*OLS*	*IV*	*OLS*	*IV*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
1{Retire}		0.003^{***}	0.077^{***}	0.004^{***}	0.016	0.001^{***}	0.085^{**}	0.0004	0.026	−0.0004	0.024
		(0.0004)	(0.025)	(0.0005)	(0.029)	(0.001)	(0.033)	(0.0004)	(0.023)	(0.0005)	(0.030)
1{SS}	0.009^{***}										
	(0.0004)										
1	0.0001	0.001	0.001	0.001	0.001	−0.001	−0.001	−0.0001	−0.0001	−0.0001	−0.0001
	(0.005)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)	(0.003)	(0.003)	(0.003)	(0.003)
N	2.423m	2.423m	2.423m	2.423m	2.423m	2.423m	2.423m	2.423m	2.423m	2.423m	2.423m
R^2	0.0012	0.0005	0.0004	0.0004	0.0004	0.0007	0.0007	0.0004	0.0004	0.0004	0.0004
Adj. R^2	0.0011	0.0004	0.0004	0.0004	0.0003	0.0007	0.0007	0.0003	0.0003	0.0003	0.0003
F Stat.	24.0	9.7	9.3	8.8	8.0	14.9	14.8	7.8	7.8	7.3	7.3

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
SEs not clustered.

4.4 Results II: Retirement, New Mortgage Debt, and Liquid Balances

I study the effects of retirement-induced housing wealth management decisions (i.e. purchase, sale, equity extraction). In the HRS, I find that HH liquid balances increase by ~\$5k on average in the year that they claim SS benefits and that HH mortgage debt increases by about half this amount at the same time. I use SS eligibility thresholds as an instrument for refinancing decisions to look for an effect on liquid balances. Instrumenting for home sale, equity extraction, and second mortgage origination increases the effect on liquid balances relative to an OLS estimate. Equity extraction yields a ~\$20k increase in liquid balances when induced by social security eligibility.

4.4.1 Event Study

I use the HRS "crude sample" to conduct an event study of HH balance sheets at the time they claim SS benefits. For each HH, I define the date they first claimed SS benefits, t_i^{SS} using the earlier claim date in the case of couples (i.e. regardless of gender). For any interview in which any member of the HH participated, I record balance sheet and income statement elements, B, and date them by the distance, in years, to the time of the SS claim, $t_i^B = t_i^{Int.} - t_i^{SS}$. I group observations by years from SS claim and compute means and 95% confidence intervals within each bucket.

Figure (4.3) reports the results. The upper left panel shows SS income spiking precisely after SS claims, evidencing that the HHs have been aligned correctly. The upper right panel depicts total income across the SS claim threshold. HH total income is declining because HHs have retired and SS income does not completely replace these lost earnings. The decline in income (though this, of course, may be accompanied by a decline in consumption) makes the rise in liquid balances, especially a sudden rise, somewhat surprising.

The middle left panel depicts liquid balances across the SS claim threshold. These appear relatively stable before and after, but there is a discrete increase in liquid balances at the time the HHs are first claiming SS. [4] The middle panel on the right depicts HH total assets. The magnitudes and error bars make it difficult to interpret but there does not appear to be a discontinuous change in wealth. Instead, liquid balances appear to be the result of reorganizing the balance sheet.

The bottom two panels examine HH housing wealth and mortgage debt. In the aggregate, there is limited evidence of HH liquidation of housing wealth. The wealth measure is trends smoothly through the SS claim threshold. The mortgage debt appears to increase somewhat in the first year after HHs claim SS benefits. The increase is statistically insignificant but it can account for 50% of the increase in liquid balances. The fact that there is no similar trend reversal in housing wealth could be due to downsizing and equity extraction off-setting each other.

4. The increase in liquid balances before and after retirement has been documented in Olafsson and Pagel [2018] but they argue this is due to changing consumption patterns. In contrast to their work, this paper offers that balance sheet reorganization is driving changing liquid savings balances. This is supported both by the fact that the liquid balances appear to increase at a specific threshold rather than gradually, a fact their paper does not disentangle.

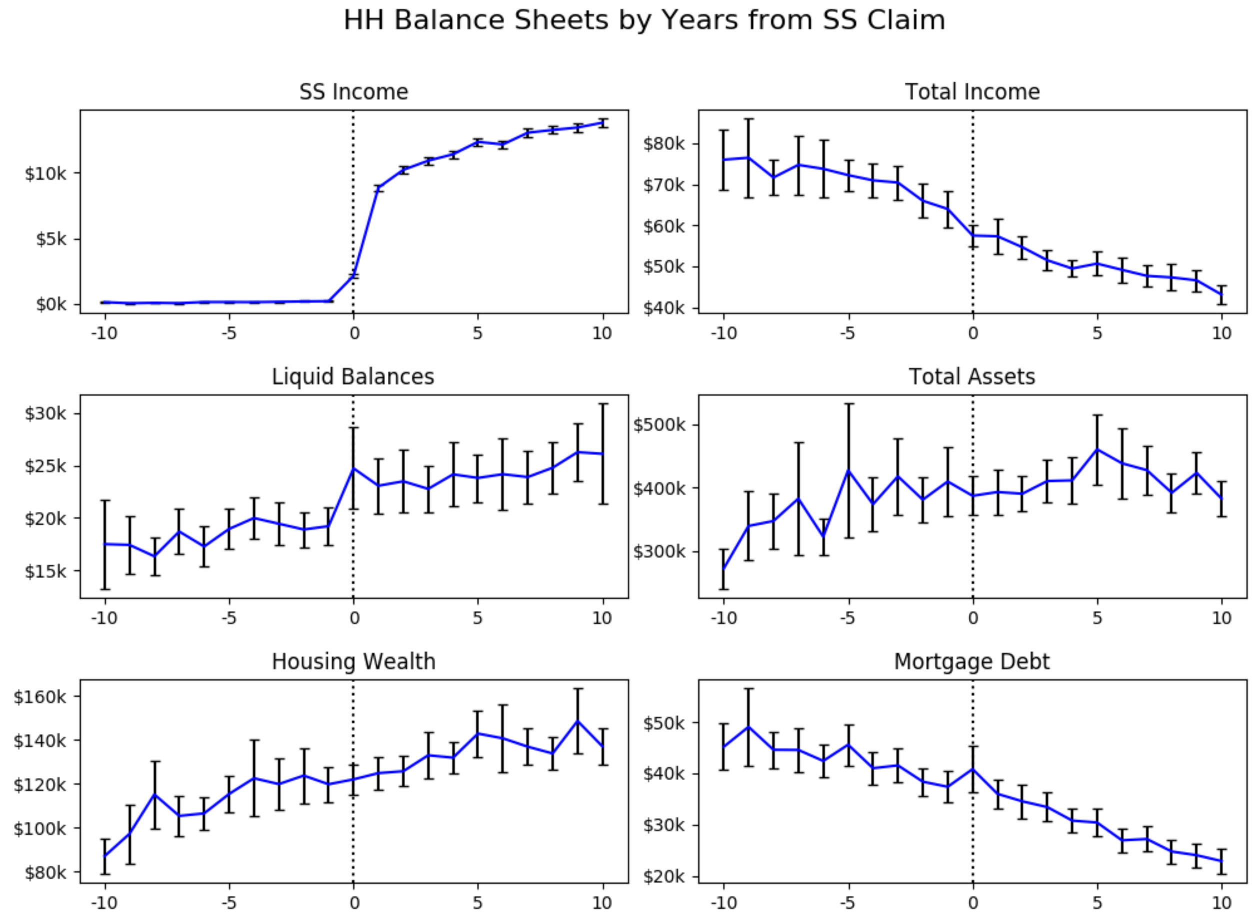

Figure 4.3: The above are average income statement and balance sheet items of HRS HHs with unchanging family structure, at least one member claiming SS benefits from 2002-2014, and responding in a given interview. Means and 95pp confidence intervals of observations are reported, dated by distance of the interview to the SS claim date, and binned by years. Liquid balances (savings, checking, and money market mutual funds) increase discontinuously at the SS claim date despite falling total income. There is statistically inconclusive but suggestive evidence that increasing mortgage debt may be driving this.

4.4.2 IV Evidence

I formally study whether the additional housing wealth management activity taking place at the SS claim threshold contributes to elevated liquid balances by using the IV specification described at the end of section (4.2.2). Specifically, using the HRS "refined sample", I construct a panel in interview-wave time. I construct a series, by interview wave, of the change in HH liquid balances from its previous interview, ΔSav_{it}. For each interview period and for different housing wealth management actions, I construct indicators of whether the HH took the action since the previous interview wave. [5] For each HH, I assign a HH head analogously to the SCF designation, as the financial respondent if male and their spouse if female for coupled HHs; I construct the SS eligibility instrument according to whether the HH head became eligible since the last interview. Finally, to ensure that outliers are not driving the result, I winsorize the top and bottom 5% of ΔSav_{it} in the panel.

The OLS and IV regression results are presented in table (4.6). Columns (1), (4), and (7) show that the SS eligibility instrument is much weaker but still somewhat relevant. It may be weakened because interviews are far apart making it difficult to pinpoint in interview-wave time when a HH crosses an eligibility threshold. It appears strong still in the case of equity extraction. The IV estimates are larger than the OLS estimates, in the case of home sales and equity extractions, significantly so. This is consistent with adverse shocks prompting HHs to spend down liquid balances and borrow against or sell their houses to spend out of housing wealth as well. Instrumenting identifies the effect on liquid balances of the additional margin of houses refinancing at retirement. These are \$60k and \$23.1k for housing sales and equity extractions, respectively.

I conclude that the additional housing wealth management activity taking place at the SS claim date can help explain, in part, the simultaneous one-time discontinuous rise in liquid balances.

5. I consider purchase, sale, equity extraction, second mortgage origination, and HELOC origination and impute these values as described above.

Table 4.6: Housing Wealth Management and Liquid Balances at Retirement (HRS)

	Dependent variable:								
	1{Sell}	ΔSav_{it}		1{Extract}	ΔSav_{it}		1{New 2nd}	ΔSav_{it}	
	1st Stage	*OLS*	*IV*	*1st Stage*	*OLS*	*IV*	*1st Stage*	*OLS*	*IV*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
1{Sell}		3.4k***	60.6k*						
		(472)	(32.3k)						
1{Extract}					−103	23.1k*			
					(180)	(12.4k)			
1{New 2nd}								−872***	140k
								(280)	(137k)
1{SS}	0.004***			0.011***			0.002		
	(0.001)			(0.002)			(0.002)		
1	−1.1k	−0.003	−923	−0.031	−1.2k	−496	−0.007	−1.1k	−138
	(1.3k)	(0.004)	(1.4k)	(0.020)	(1.3k)	(1.7k)	(0.006)	(1.3k)	(2.1k)
Demo FEs	x	x	x	x	x	x	x	x	x
Clustered SEs	x	x	x	x	x	x	x	x	x
N	93,572	94,510	93,572	92,259	93,190	92,259	93,018	93,949	93,018
R^2	0.006	0.004	−0.285	0.035	0.003	−0.161	0.013	0.003	−2.836
Adj. R^2	0.004	0.002	−0.287	0.034	0.001	−0.163	0.012	0.001	−2.842
F Stat.	3.928***	2.466***		24.824***	1.776***		9.116***	1.869***	

Note: *p<0.1; **p<0.05; ***p<0.01
Data winsorized by ΔSav_{it} at 5pp level.

4.5 Interpretation

I have documented that retirement causes the issuance of new mortgage debt, but issuance remains a choice that is mediated by the HH. In this section, I conduct an audit of the HH's decision-making to better understand their economic motives. I consider several potential explanations: HHs may be more motivated to move after retirement, they may need accessible wealth to supplement reduced income and smooth consumption, they may simplify their financial obligations, and they may enjoy better credit after claiming SS benefits.

As evidence, I consider the origination of mortgage debt for current and purchased housing, the effects of refinancing on savings balances, credit scores at SS eligibility thresholds, and survey evidence on the stated purpose of refinancing. The evidence is modest and suggests that more than one explanation may apply, but ultimately inconclusive. It is consistent with the explanation that HHs are moving but this cannot explain all of the mortgage debt originated. It appears inconsistent with expanded credit supply for HHs on SS income. It is plausible both that HHs are consumption smoothing or simplifying their financial obligations, though we document a transfer of wealth to liquid balances that appears to persist and which reduces the power of those explanations on the margin.

4.5.1 Explanations

Retirement might motivate a HH to originate new mortgage debt for several reasons. In this section, I consider the following:

1. HHs may become less attached to the geographic region of their former residence after retirement. This may be because they substitute away from working and toward leisure activities or because, even if they wanted to re-enter the labor force, they do not benefit more by being close to a former employer. HHs may pay off their former mortgage and take on a new one in order to finance the purchase of a new house in a different region.

2. It is possible that HHs face credit constraints in borrowing against their homes, in

particular if prospective lenders consider their future income to be too risky. SS benefits are guaranteed by the government, however, and lenders might be more willing to lend because of this. The rise in additional borrowing against current housing might be explained by an extension of credit to HHs who would have borrower earlier if they could have.

3. Households retiring from the labor force may be required to consider their finances as the source and amount of their income changes, as they begin to tap into their pension wealth, and as they budget for retirement and plan to unwind their estate. This required degree of scrutiny may reduce the marginal effort required to simplify finances by, say, consolidating or paying off debt. Borrowing more against current housing may be a way of doing this.

4. Income at retirement falls predictably because SS benefits do not completely replace lost employment income. The PIH predicts that a HH faced with a predictable loss in income will have saved wealth and then begin to consume out of it. Liquidating housing assets might be an instance of such consumption smoothing.

4.5.2 Evidence

Current vs. Acquired Housing

To the extent that HHs value geography differently after retirement, they will issue debt against newly acquired housing rather than lever their current housing assets. In both the SCF and HRS, it is possible to observe purchases of new housing separately from issuance of additional debt against current housing. The spikes in the upper right panel of figure (4.4) and the upper left panel of figure (4.5) both suggest that new purchases contribute to the debt issuance at retirement. I cannot reject the null that retirement does not induce home purchases in both the SCF and HRS regression analysis, but the HRS is very conservatively estimated.

Even if moving accounts for some of the issuance of new mortgage debt, it cannot account for all of it. In particular, equity extractions appear to spike at retirement. This means explanations which rely on HHs changing their housing consumption decision are incomplete. Among those who do not move but borrow more heavily, housing is being exploited as a purely financial asset. To understand their motives, it is necessary to understand what they do with the liquidated funds.

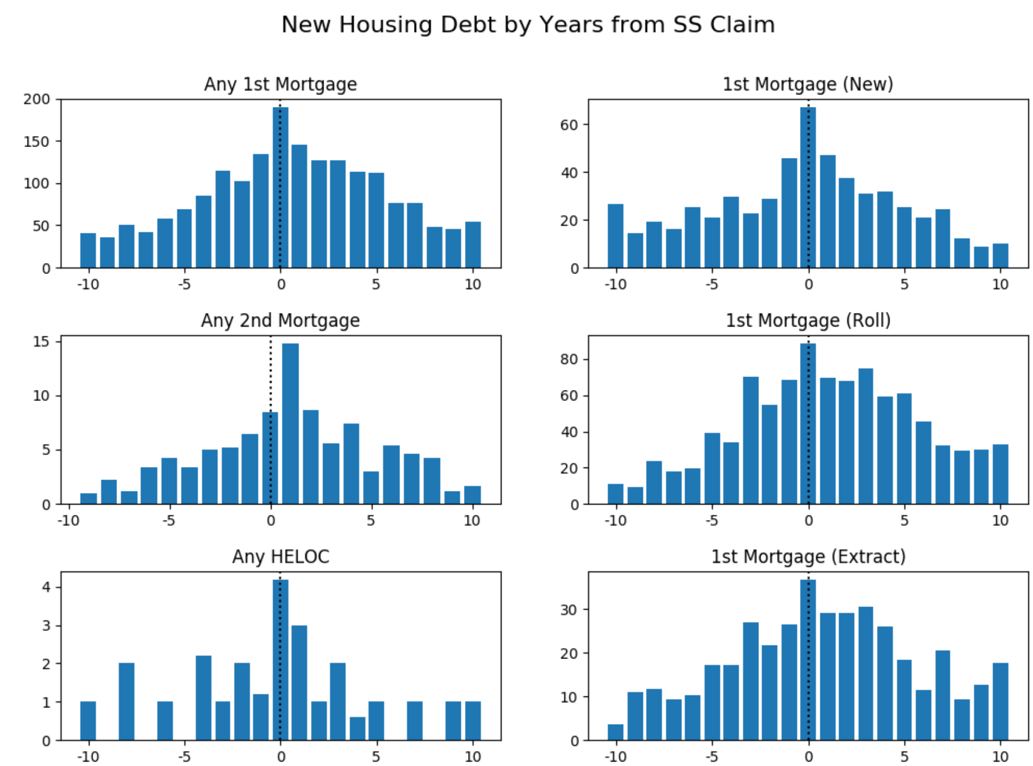

Figure 4.4: The above are counts of when HHs take out mortgage debt relative to the year they claim SS benefits. They are derived from the SCF 1992-2016 and are adjusted for multiple implicates. New mortgage obligations spike in the year or year after HHs claim SS benefits. Some is the result of new home purchases, but some is the result of borrowing against existing housing, especially 2nd mortgages and equity extractions.

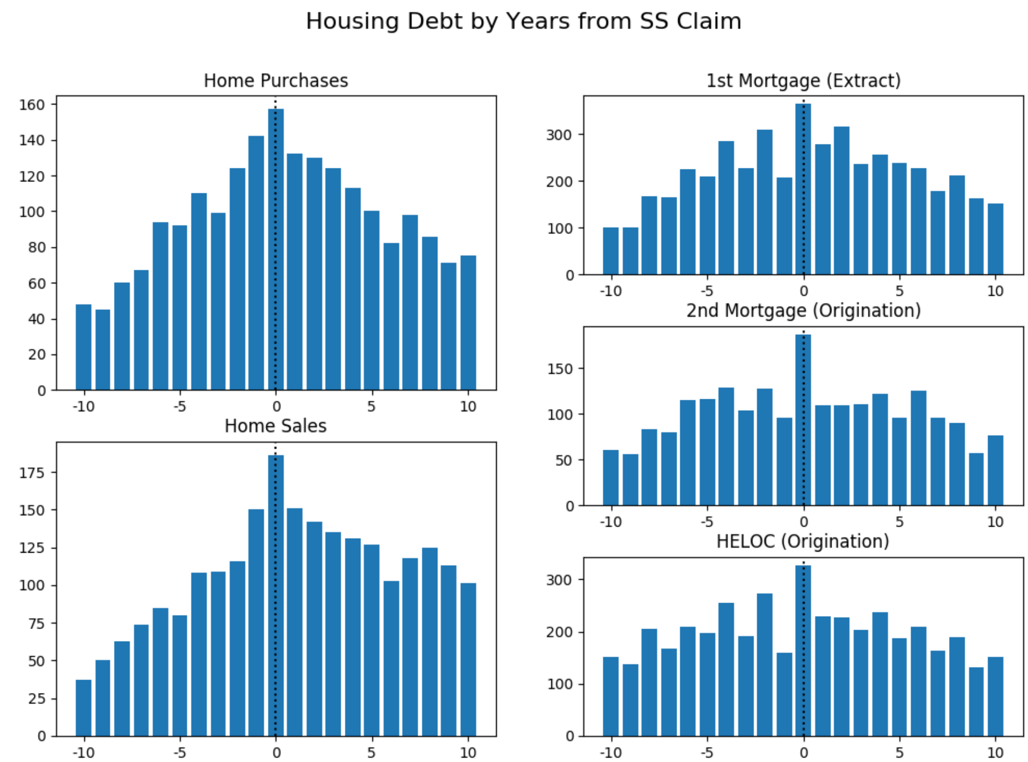

Figure 4.5: The above are counts of housing wealth management decisions by HHs relative to the year in which they first claim SS benefits. They are derived from the HRS 1992-2014 RAND Fat Files. HHs increase their purchase and sale of homes, but also take out significantly more debt against their homes in SS claim years. The role of additional debt on already owned homes is more prominent in the HRS than in the SCF.

Refinancing and Liquid Balances

In section (4.4), I document that the additional margin of sales and equity extraction taking place at retirement increases liquid balances. The fact that the sign of association between extractions and changes in liquid balances changes after instrumenting is noteworthy. It suggests that the housing wealth management activity taking place at the threshold is proactive and not a response to expense shocks.

In principle, if all of the extracted dollars wind up permanently in liquid balances, this is inconsistent with stories of debt consolidation and consumption smoothing. So the fact that some dollars are deposited in liquid balances tempers the extent of those explanations.

Of course, it is not clear that everything extracted is deposited there or that liquid balances aren't used for consumption at a later date. The fact that liquid balances appear relatively stable after the SS claim date in figure (4.3) suggests that these transfers are not being systematically consumed (even at a later date). If precautionary liquid balances became more valuable in retirement, this might help explain the persistence of the increase.

Credit Scores

I use the random sample from the TU dataset to analyze how credit scores change at SS eligibility thresholds. I subset the sample to retain individuals whose birth date and Vantage3 score is available. I bucket by age in years and compute average Vantage3 scores. Figure (4.6) plots this average as well as the SS program thresholds. There are no discontinuous increases in the Vantage3 credit score at any of the SS program thresholds, so if credit becomes more available, it does not appear to be happening through a credit score. Credit may become available to SS recipients through channels other than the credit score, but this suggests that refinancing activity is not driven by credit constrained individuals realizing additional credit supply.

Stated Purpose of Extraction

The SCF asks HHs who report extracting equity from their home their reason for doing so. Figure (4.7) plots counts of equity extractions by years from SS claim according to HHs stated purpose. The quantities of HHs are low so the results should be interpreted with caution. What is notable is that the spike in extractions contemporaneous with claiming SS benefits persists in several of the subgroups, including investment purposes.

Equity extractions for the purpose of non-housing consumption do not tick up in the year of SS claim. This suggests that perhaps consumption smoothing is not the motive. There is, however, an uptick in HHs that report extracting equity to invest in their homes. The future flow of housing services enjoyed might reasonably count toward consumption. It is worth

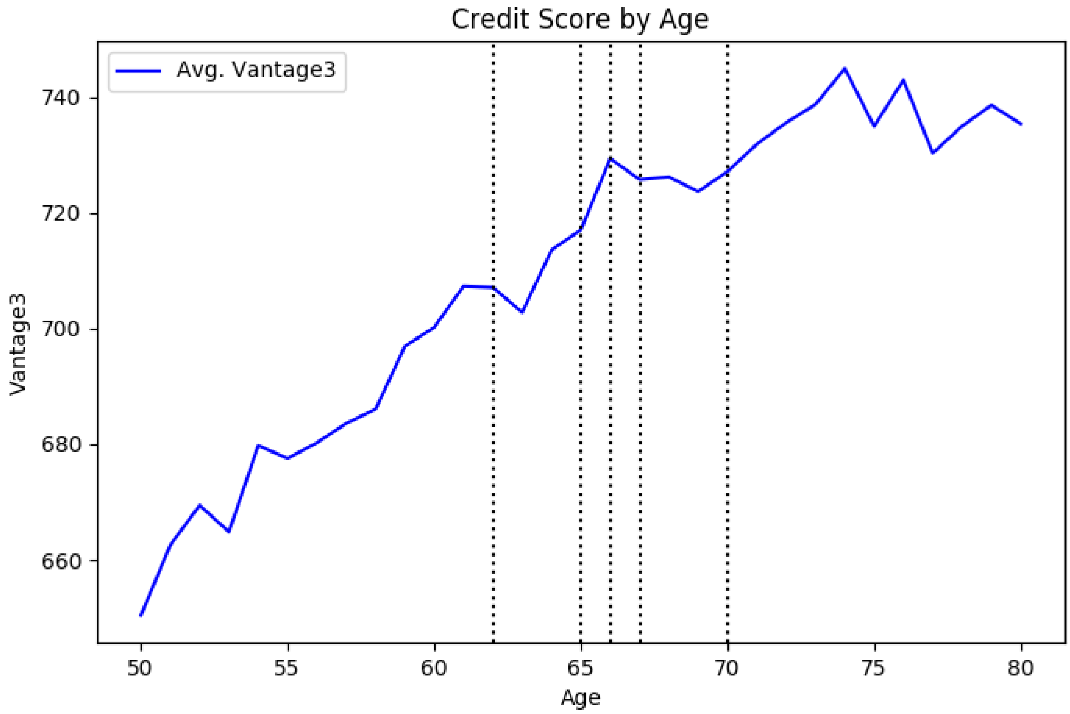

Figure 4.6: The above is a plot of average Vantage3 credit scores by age for a random sample of 100k individuals in June 2012. Various SS eligibility thresholds are plotted. There does not appear to be any discontinuous increase in credit score over these thresholds. This suggests that retirement refinancing induced by SS eligibility is not the result of changing credit supply.

noting, though, that many of these investments, anecdotally, are undertaken to allow HHs to age in place. These sorts of investments, then, represent age-differentiated consumption goods that are not captured by the simple single-good PIH model.

There is a notable increase in HHs who cite debts and gifts as their reasoning for extracting equity when claiming SS benefits. Paying off debts would be consistent with a debt consolidation story but gifts would not be and it is not possible to decompose this measure further.

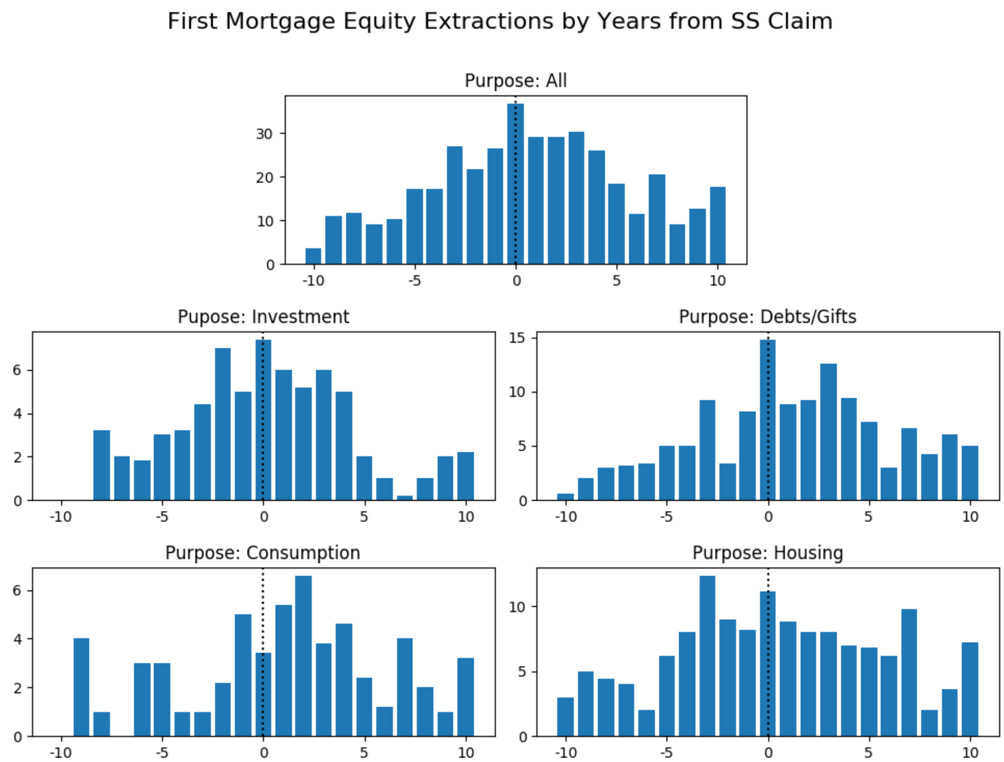

Figure 4.7: The above are counts of when HHs extract equity from their homes for different purposes relative to the year they claim SS benefits. They are derived from the SCF 1992-2016 and are adjusted for multiple implicates. HHs appear to be paying off outstanding debts and investing especially in their homes.

4.6 Discussion

I have documented that retirement prompts HHs to issue new mortgage debt, both to finance the purchase of new homes and to borrow against current housing. I have further documented that this retirement-induced refinancing appears to increase the liquid balances of HHs and can help explain the discontinuous increase in liquid balances around the time HHs claim SS benefits despite no corresponding discontinuous increase in total assets.

The fact that housing sales and equity extraction prompt larger increases in liquid balances at this threshold suggests that these housing wealth management decisions have a different character than other occasions on which HHs liquidate housing wealth. One characterization consistent with the data are that ordinarily housing sales and equity extractions are reactive, in the sense that HHs liquidate in response to an adverse expense shock; at retirement, liquidation is pro-active, in the sense that HHs may not have an immediate need for the funds but plan for future needs.

Further evidence on the motives for refinancing at the retirement threshold is somewhat inconclusive. HHs may be moving in response to changing valuation of geography, but this cannot explain all of the financial activity. It does not appear to be the case that credit constraints are relaxed at the threshold. HHs may be smoothing consumption or consolidating debt, but the fact that some extracted funds appear to be stored as liquid balances and the fact that liquid balances appear to remain high and stable after SS claims limits these margins.

Conceptually, what is oddest about the HH decision to extract equity is that the drop in income at retirement is predictable and refinancing a home is expensive, on the order of $2-5k. It would seem in the interest of HHs to borrow more money when purchasing a home or select a mortgage contract that pays down more slowly. Investing the funds elsewhere would allow HHs to access them at retirement without incurring transaction costs associated with refinancing. This would constitute a rational improvement in cases where the HH changes only financing and not consumption of housing at retirement.

There are several reasons why the proposed financing changes may not, in fact, be optimal in this case. There may be a wedge between borrowing and lending rates due to financial frictions. HHs may already be up against binding leverage constraints when borrowing. And housing liquidations may be driven by shocks to home values considered at the retirement threshold rather than predictable wealth. The extent of equity extractions due to retirement is slim in my estimates. That said, predictable and expensive divestment may be a useful approach for rejecting a rational benchmark in other settings. I leave this as a line of inquiry for future research.

CHAPTER 5

CONCLUSION

This thesis re-examines the micro-economics of households and housing in light of the lessons of the Great Financial Crisis of 2007-8. The crisis made clear the interdependence of the mortgage and housing markets, the bounded rationality and limited credit access of households, and the importance of the institutions facilitating mortgage lending in the United States. I incorporate these themes into several essays that sketch the relationship between households and housing. These essays revisit a variety of topics in housing finance such as residential sorting, investment demand for housing, and consumption-smoothing over the life-cycle.

Chapter (2) examines the household's joint problem of housing and mortgage choice at the time of home purchase. I argue that for constrained households, hedonic estimates understate the household's valuation of marginal services. I also argue that household's willingness-to-pay for credit at the time of home purchase captures the size of this bias. I develop an approach to estimate the extent of household credit constraints at the time of home purchase from the household's mortgage choice. In an application to valuations of school quality, I find that the size of the bias may be as large as 50%.

Chapter (3) identifies a kinked choice set in the housing finance choice due to the combination of loan-to-value and debt-to-income limits. I adapt the bunching framework to the household's joint problem of housing and first mortgage choice. I argue that a bunching estimator can be used to estimate the loan-to-value elasticity of housing demand that, relative to the literature, better shuts down concerns about beliefs over housing returns driving demand. I also argue that this estimator provides a direct test of whether credit constraints are binding at home purchase. I find evidence of statistically significant bunching, although within the subpopulation I examine, the magnitude is small, with an estimated elasticity of ~14-25bp.

Chapter (4) considers how households consume out of housing wealth in retirement. In contrast to retirement literature that emphasizes consumption of housing wealth only in response to shocks (e.g. divorce or spousal death), it finds evidence of predictable liquidation of housing assets at the social security threshold. It further argues that this represents a demand for liquid assets in retirement and that, because this demand is predictable and refinancing is expensive, households may be saving too much in their houses relative to a rational benchmark. This essay raises questions for further research into mortgage commitments as savings devices.

REFERENCES

Daniel Aaronson, Daniel Hartley, and Bhashkar Mazumder. The effects of the 1930s holc "redlining" maps. *Federal Reserve Bank of Chicago*, 2019.

Arthur Acolin, Jesse Bricker, Paul Calem, and Susan Wachter. Borrowing constraints and homeownership. *American Economic Review: Papers and Proceedings*, 106:625–629, 2016. doi: 10.1257/aer.p20161084. URL http://repository.upenn.edu/fnce_papershttp://dx.doi.org/10.1257/aer.p20161084.

Manuel Adelino, Kristopher Gerardi, and Paul S. Willen. Why don't lenders renegotiate more home mortgages? redefaults, self-cures and securitization. *Journal of Monetary Economics*, 60:835–853, 10 2013. ISSN 03043932. doi: 10.1016/j.jmoneco.2013.08.002.

Sumit Agarwal, John C Driscoll, Xavier Gabaix, and David Laibson. The age of reason: Financial decisions over the life-cycle with implications for regulation, 2009. URL https://dash.harvard.edu/bitstream/handle/1/4554335/laibson_ageofreason.pdf?sequence=2.

Sumit Agarwal, John C. Driscoll, and David I. Laibson. Optimal mortgage refinancing: A closed-form solution. *Journal of Money, Credit and Banking*, 45:591–622, 2013. ISSN 00222879. doi: 10.1111/jmcb.12017.

Sumit Agarwal, Jessica Pan, and Wenlan Qian. Age of decision: Pension savings withdrawal and consumption and debt response*, 2018. URL https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/301563/Pensions_fact_sheet_v8.pdf.

Mark Aguiar and Erik Hurst. Consumption versus expenditure. *Journal of Political Economy*, 113:919–948, 2005. doi: 10.1086/491590. URL https://www.jstor.org/stable/pdf/10.1086/491590.pdf?refreqid=excelsior%3A2fbdd52207b3bfd95882ccabd7dec829.

Mark Aguiar and Erik Hurst. Life-cycle prices and production. *American Economic Review*, 97:1533–1559, 11 2007. ISSN 0002-8282. doi: 10.1257/aer.97.5.1533. URL http://pubs.aeaweb.org/doi/10.1257/aer.97.5.1533.

Mark Aguiar and Erik Hurst. Deconstructing life cycle expenditure. *Journal of Political Economy*, 121:437–492, 2013. doi: 10.1086/670740. URL https://scholar.princeton.edu/sites/default/files/670740_0.pdf.

Gene Amromin, Jennifer Huang, and Clemens Sialm. The tradeoff between mortgage prepayments and tax-deferred retirement savings. *Journal of Public Economics*, 91:2014–2040, 2007. doi: 10.1016/j.jpubeco.2007.03.011. URL www.elsevier.com/locate/econbase.

Xudong An and Larry Cordell. Mortgage loss severities: What keeps them so high? *Real Estate Economics*, 49:809–842, 2021. ISSN 15406229. doi: 10.1111/1540-6229.12334.

Elliot Anenberg, Aurel Hizmo, Edward Kung, and Raven Molloy. Measuring mortgage credit availability: A frontier estimation approach measuring mortgage credit availability: A frontier estimation approach *, 2017. URL https://doi.org/10.17016/FEDS.2017.101.

Robert Argento, Victoria L. Bryant, and John Sabelhaus. Early withdrawals from retirement accounts during the great recession. *Contemporary Economic Policy*, 33:1–16, 1 2015. ISSN 10743529. doi: 10.1111/coep.12064. URL http://doi.wiley.com/10.1111/coep.12064.

Natalie Bachas, Olivia S. Kim, and Constantine Yannelis. Loan guarantees and credit supply. *Journal of Financial Economics*, 8 2020. ISSN 0304405X. doi: 10.1016/j.jfineco.2020.08.008.

Michael Bailey, Eduardo Dávila, Theresa Kuchler, and Johannes Stroebel. House price beliefs and mortgage leverage choice. *Review of Economic Studies*, 86:2403–2452, 11 2019. ISSN 0034-6527. doi: 10.1093/restud/rdy068. URL https://academic.oup.com/restud/article/86/6/2403/5194341.

James Banks, Richard Blundell, and Sarah Tanner. Is there a retirement-savings puzzle? *American Economic Review*, 88:769–788, 1998. URL https://www.jstor.org/stable/pdf/117005.pdf?refreqid=excelsior%3A95a40009231633f0b551f00886a3ec68.

Keith Barnatchez, Leland D Crane, and Ryan A Decker. An assessment of the national establishment time series (nets) database, 2017. URL https://doi.org/10.17016/FEDS.2017.110.%0Ahttps://doi.org/10.17016/FEDS.2017.110.

Nicholas Barr and Peter Diamond. The economics of pensions. *Oxford Review of Economic Policy*, 22:15–39, 2006. doi: 10.1093/oxrep/grj002. URL http://eprints.lse.ac.uk/2630/.

Robert Bartlett, Adair Morse, Richard Stanton, and Nancy Wallace. Consumer-lending discrimination in the fintech era. *Journal of Financial Economics*, 2021. ISSN 0304405X. doi: 10.1016/j.jfineco.2021.05.047. URL https://doi.org/10.1016/j.jfineco.2021.05.047.

Patrick Bayer, Fernando Ferreira, and Robert Mcmillan. A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, 115, 2007. URL http://www.journals.uchicago.edu/t-and-c.

Matteo Benetton. Leverage regulation and market structure: A structural model of the uk mortgage market. *The Journal of Finance*, page jofi.13072, 8 2021. ISSN 1540-6261. doi: 10.1111/JOFI.13072. URL https://onlinelibrary.wiley.com/doi/full/10.1111/jofi.13072.

Paul Bennett, Richard Peach, and Stavros Peristiani. Structural change in the mortgage market and the propensity to refinance. *Journal of Money, Credit and Banking*, 33:955–975, 2001. URL https://www.jstor.org/stable/pdf/2673930.pdf?refreqid=excelsior%3A393f155cbb722dce8fbdd4642f92b0c9.

Jonathan Berk and Richard Roll. Adjustable rate mortgages: Valuation. *Journal of Real Estate Finance and Economics*, 1:163–184, 1988.

Susanna B Berkouwer and Joshua T Dean. Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households. 2021.

B. Douglas Bernheim. Dissaving after retirement: Testing the pure life cycle hypothesis, 1987. URL https://www.nber.org/chapters/c6861.pdf.

B. Douglas Bernheim, Jonathan Skinner, and Steven Weinberg. What accounts for the variation in retirement wealth among u.s. households? *American Economic Review*, 91: 832–857, 2001. URL https://pubs.aeaweb.org/doi/pdfplus/10.1257/aer.91.4.832.

John Beshears, James J Choi, David Laibson, and Brigitte C Madrian. The availability and utilization of 401(k) loans, 2011. URL http://www.nber.org/papers/w17118.

Michael Carlos Best, James Cloyne, Ethan Ilzetzki, and Henrik Kleven. Estimating the elasticity of intertemporal substitution using mortgage notches, 2018. URL http://www.nber.org/papers/w24948.

Neil Bhutta and Daniel Ringo. The effect of interest rates on home buying: Evidence from a discontinuity in mortgage insurance premiums, 2017. URL https://doi.org/10.17016/FEDS.2017.086.

Xun Bian, Zhenguo Lin, and Yingchun Liu. House price, loan-to-value ratio and credit risk. *Journal of Banking and Finance*, 92:1–12, 7 2018. ISSN 03784266. doi: 10.1016/j.jbankfin.2018.04.006. URL https://doi.org/10.1016/j.jbankfin.2018.04.006.

Sandra E Black. Do better schools matter? parental valuation of elementary education. *Quarterly Journal of Economics*, 114:577–599, 1999. URL https://academic.oup.com/qje/article-abstract/114/2/577/1844232.

Richard Blundell and Hilary W Hoynes. Has 'in-work' benefit reform helped the labor market?, 2004. URL http://www.nber.org/books/card04-1.

Richard Blundell and Andrew Shephard. Employment, hours of work and the optimal taxation of low-income families. *Review of Economic Studies*, 79:481–510, 4 2012. ISSN 00346527. doi: 10.1093/restud/rdr034. URL https://academic.oup.com/restud/article/79/2/481/1532543.

Jesse Bricker, Arthur B Kennickell, Kevin B Moore, and John Sabelhaus. Changes in u.s. family finances from 2007 to 2010: Evidence from the survey of consumer finances, 2012. URL www.federalreserve.gov/pubs/bulletin/default.htm.

Kristine M Brown. The link between pensions and retirement timing: Lessons from california teachers. *Journal of Public Economics*, 98:1–14, 2013. doi: 10.1016/j.jpubeco.2012.10.007. URL http://dx.doi.org/10.1016/j.jpubeco.2012.10.007.

Meta Brown. The graying of american debt, 2016. URL https://www.newyorkfed.org/medialibrary/media/newsevents/mediaadvisory/2016/Graying-of-American-Debt-02122016.pdf.

Greg Buchak and Adam Jørring. Do mortgage lenders compete locally? implications for credit access, 2021.

Barbara A Butrica and Nadia S Karamcheva. Does household debt influence the labor supply and benefit claiming decisions of older americans? *CRR Working Paper Series*, 2013. URL http://ssrn.com/abstract=2368389.

Barbara A Butrica and Stipica Mudrazija. Home equity patterns among older american households, 2016. URL https://www.urban.org/sites/default/files/publication/85326/home-equity-patterns-among-older-american-households_0.pdf.

Kenneth Y Chay and Michael Greenstone. Does air quality matter? evidence from the housing market. *Journal of Political Economy*, 113:376–424, 2005. doi: 10.1086/427462.

Raj Chetty, John N Friedman, Tore Olsen, and Luigi Pistaferri. Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from danish tax records *. *Quarterly Journal of Economics*, pages 749–804, 2011. doi: 10.1093/qje/qjr013. URL https://academic.oup.com/qje/article/126/2/749/1870394.

Raj Chetty, Nathaniel Hendren, and Lawrence F Katz. The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment *, 2015.

Eric Chyn. Moved to opportunity: The long-run effects of public housing demolition on children †. *American Economic Review*, 108:3028–3056, 2018. doi: 10.1257/aer.20161352. URL https://doi.org/10.1257/aer.20161352.

Susan Clampet-Lundquist and Douglas S Massey. Neighborhood effects on economic self-sufficiency: A reconsideration of the moving to opportunity experiment. *American Journal of Sociology*, 114:107–143, 2008.

Courtney C. Coile. Economic determinants of workers' retirement decisions. *Journal of Economic Surveys*, 29:830–853, 9 2015. ISSN 09500804. doi: 10.1111/joes.12115. URL http://doi.wiley.com/10.1111/joes.12115.

Julia L. Coronado, Dean Maki, and Ben Weitzer. Retiring on the house? cross-cohort differences in housing wealth, 2007. URL http://pensionresearchcouncil.wharton.upenn.edu/wp-content/uploads/2015/09/0-19-923077-3-14.pdf.

Stephanie E Curcuru. The changing investment behavior of u.s. seniors board of governors of the federal reserve system, 2005.

Janet Currie, Lucas Davis, Michael Greenstone, and Reed Walker. Environmental health risks and housing values: Evidence from 1,600 toxic plant openings and closings †. *American Economic Review*, 105:678–709, 2015. doi: 10.1257/aer.20121656. URL http://dx.doi.org/10.1257/aer.20121656.

Ramona Dagostino. The impact of bank financing on municipalities' bond issuance and the real economy, 2019.

Thomas Davidoff. Housing, health, and annuities. *Journal of Risk and Insurance*, 76:31–52, 2009.

Anthony A. DeFusco and John Mondragon. No job, no money, no refi: Frictions to refinancing in a recession. *The Journal of Finance*, 75:2327–2376, 2020. ISSN 15406261. doi: 10.1111/jofi.12952.

Anthony A. DeFusco and Andrew Paciorek. The interest rate elasticity of mortgage demand: Evidence from bunching at the conforming loan limit. *American Economic Journal: Economic Policy*, 9:210–240, 2 2017. ISSN 1945774X. doi: 10.1257/pol.20140108. URL https://doi.org/10.1257/pol.20140108.

Anthony A. DeFusco, Stephanie Johnson, and John Mondragon. Regulating household leverage. *Review of Economic Studies*, 87:914–958, 3 2020. ISSN 0034-6527. doi: 10.1093/RESTUD/RDZ040. URL https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdz040/5542954https://academic.oup.com/restud/article/87/2/914/5542954.

Peter Diamond and Jonathan Gruber. Social security and retirement in the united states, 1999. URL https://www.nber.org/chapters/c7258.pdf.

Rebecca Diamond and Tim Mcquade. Who wants affordable housing in their backyard? an equilibrium analysis of low-income property development. *Journal of Political Economy*, 127:1063–1117, 2019. URL http://www.journals.uchicago.edu/t-and-c.

Antonia Díaz and María José Luengo-Prado. On the user cost and homeownership. *Review of Economic Dynamics*, 11:584–613, 2008. ISSN 10942025. doi: 10.1016/j.red.2007.12.002.

Mehran Ebrahimian. Student loans and social mobility. *SSRN Electronic Journal*, 2020. doi: 10.2139/ssrn.3680159.

Mark L Egan, Alexander MacKay, and Hanbin Yang. Recovering investor expectations from demand for index funds, 2020. ISSN 0898-2937. URL http://www.nber.org/papers/w26608.

Giovanni Favara and Jean Imbs. Credit supply and the price of housing †. *American Economic Review*, 105:958–992, 2015. doi: 10.1257/aer.20121416. URL http://dx.doi.org/10.1257/aer.20121416958.

Jack Favilukis, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh. The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy*, 125:140–223, 2 2017. ISSN 1537534X. doi: 10.1086/689606. URL https://www.journals.uchicago.edu/doi/abs/10.1086/689606.

FDIC. Affordable mortgage lending guide a resource for community bankers part i: Federal agencies and government sponsored enterprises, 2018. URL https://www.fdic.gov/mortgagelending.

Martin Feldstein. Social security, induced retirement, and aggregate capital accumulation. *Journal of Political Economy*, 82:905–926, 1974. URL http://www.journals.uchicago.edu/t-and-c.

Daniel K Fetter and Lee M Lockwood. Government old-age support and labor supply: Evidence from the old age assistance program †. *American Economic Review*, 108:2174–2211, 2018. doi: 10.1257/aer.20160921. URL https://doi.org/10.1257/aer.20160921.

Jonathan D Fisher, David S Johnson, Joseph T Marchand, Timothy M Smeeding, and Barbara Boyle Torrey. No place like home: Older adults and their housing. *Journal of Gerontology: Social Sciences*, 62:120–128, 2007. URL https://academic.oup.com/psychsocgerontology/article-abstract/62/2/S120/548979.

Maria D. Fitzpatrick and Timothy J. Moore. The mortality effects of retirement: Evidence from social security eligibility at age 62. *Journal of Public Economics*, 157:121–137, 1 2018. ISSN 00472727. doi: 10.1016/j.jpubeco.2017.12.001. URL https://linkinghub.elsevier.com/retrieve/pii/S0047272717302037.

Andreas Fuster, Laurie Goodman, David Lucca, Laurel Madar, Linsey Molloy, and Paul Willen. The rising gap between primary and secondary mortgage rates. *FRBNY Economic Policy Review*, pages 17–39, 2013. URL http://www.newyorkfed.org/research/epr/2013/1113fust.pdf.

Peter Ganong and Pascal Noel. Why do borrowers default on mortgages? a new method for causal attribution, 2020. URL http://www.nber.org/papers/w27585.

Xuwen Gao, Ran Song, and Christopher Timmins. The role of information in the rosen-roback framework, 2021. URL http://www.nber.org/papers/w28943.

Stefano Giglio, Matteo Maggiori, and Johannes Stroebel. Very long-run discount rates. *Quarterly Journal of Economics*, 130:1–53, 2015. doi: 10.1093/qje/qju036.Advance.

Gopi Shah Goda, John B Shoven, and Sita Nataraj Slavov. Does stock market performance influence retirement intentions?, 2010. URL http://www.nber.org/papers/w16211.

Gopi Shah Goda, Matthew R. Levy, Colleen Flaherty Manchester, Aaron Sojourner, and Joshua Tasoff. The role of time preferences and exponential-growth bias in retirement savings, 2015. URL http://www.nber.org/papers/w21482.

Andreas Grunewald, Jonathan Lanning, David Low, Tobias Salz, Guy Aridor, Jasper Clarkberg, Thi Mai Anh Nguyen, and Yining Zhu. Auto dealer loan intermediation: Consumer behavior and competitive effects †, 2020. URL https://www.newyorkfed.org/medialibrary/interactives/householdcredit/d.

Arpit Gupta and Christopher Hansman. Selection, leverage, and default in the mortgage market. *The Review of Financial Studies*, 00:1–51, 2021. ISSN 0893-9454. doi: 10.1093/rfs/hhab052.

John Heaton and Deborah Lucas. Portfolio choice and asset prices: The importance of entrepreneurial risk. *The Journal of Finance*, 55:1163–1198, 6 2000. ISSN 00221082. doi: 10.1111/0022-1082.00244. URL `http://doi.wiley.com/10.1111/0022-1082.00244`.

Michael D Hurd and Susann Rohwedder. Some answers to the retirement-consumption puzzle. *NBER Working Paper Series*, 2006. URL `http://www.nber.org/papers/w12057`.

Erik Hurst and Frank P Stafford. Home is where the equity is: Mortgage refinancing and household consumption. *Journal of Money, Credit and Banking*, 36:985–1014, 2004. doi: 10.1353/mcb.2005.0009. URL `https://doi.org/10.1353/mcb.2005.0009https://muse.jhu.edu/article/180472`.

Erik Hurst, Benjamin J. Keys, Amit Seru, and Joseph Vavra. Regional redistribution through the us mortgage market. *American Economic Review*, 106:2982–3028, 2016. ISSN 00028282. doi: 10.1257/aer.20151052.

Min Hwang, Chen L Miller, and Robert Van Order. Effects of fha loan limit increases by esa 2008: Housing demand and adverse selection, 2016. URL `https://ssrn.com/abstract=2789620`.

Dan Immergluck. From minor to major player: The geography of fha lending during the u.s. mortgage crisis. *Journal of Urban Affairs*, 33:1–20, 2011. ISSN 07352166. doi: 10.1111/j.1467-9906.2010.00539.x.

Erica Xuewei Jiang. Financing competitors: Shadow banks' funding and mortgage market competition *, 2021. URL `https://ssrn.com/abstract=3556917`.

Stephanie Johnson. Mortgage leverage and house prices *, 2019.

Joint Center for Housing Studies. Housing america's older adults, 2014. URL `www.jchs.harvard.edu`.

Alejandro Justiniano, Giorgio E Primiceri, and Andrea Tambalotti. Credit supply and the housing boom. *Journal of Political Economy*, 127:1317–1350, 2019.

Greg Kaplan and Giovanni L Violante. A model of the consumption response to fiscal stimulus payments, 2011. URL `http://www.nber.org/papers/w17338`.

Greg Kaplan, Kurt Mitman, and Gianluca Violante. The housing boom and bust: Frequently asked questions, 2019.

Greg Kaplan, Kurt Mitman, and Giovanni L Violante. The housing boom and bust: Model meets evidence. *Journal of Political Economy*, 128, 2020.

David A Keiser and Joseph S Shapiro. Consequences of the clean water act and the demand for water quality. *Quarterly Journal of Economics*, 134:349–396, 2 2019. ISSN 0033-5533. doi: 10.1093/QJE/QJY019. URL https://academic.oup.com/qje/article/134/1/349/5092609.

Benjamin J. Keys, Devin G. Pope, and Jaren C. Pope. Failure to refinance. *Journal of Financial Economics*, 122:482–499, 12 2016. ISSN 0304405X. doi: 10.1016/j.jfineco.2016.01.031. URL http://dx.doi.org/10.1016/j.jfineco.2016.01.031.

Henrik J Kleven and Mazhar Waseem. Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from pakistan*. *Quarterly Journal of Economics*, 128:669–723, 2013. doi: 10.1093/qje/qjt004. URL https://academic.oup.com/qje/article/128/2/669/1943151.

Henrik Jacobsen Kleven. Bunching. *Annual Review of Economics*, 8:435–464, 2016. doi: 10.1146/annurev-economics-080315-015234. URL www.annualreviews.org.

Jeffrey Kling, Jeffrey B Liebman, and Lawrence F Katz. Experimental analysis of neighborhood effects. *Econometrica*, 75:83–119, 2007.

Laurence J Kotlikoff and Lawrence H Summers. The role of intergenerational transfers in aggregate capital accumulation. *Journal of Political Economy*, 89:706–732, 1981. URL https://www.jstor.org/stable/pdf/1833031.pdf?refreqid=excelsior%3A0f75212bc0d469592c32045d37ad1936.

Amrita Kulka. Sorting into neighborhoods : The role of minimum lot sizes, 2019. URL https://drive.google.com/open?id=1Wr8T687wz-jVVMEVWoZCQXB53pxCM5hK.

Nicolai V. Kuminoff and Jaren C. Pope. Do "capitalization effects" for public goods reveal the public's willingness to pay? *International Economic Review*, 55:1227–1250, 2014. ISSN 14682354. doi: 10.1111/iere.12088.

Tim Landvoigt, Monika Piazzesi, and Martin Schneider. The housing market(s) of san diego †. *American Economic Review*, 105:1371–1407, 2015. doi: 10.1257/aer.20111662. URL http://dx.doi.org/10.1257/aer.20111662.

Kenneth Lee, Edward Miguel, and Catherine Wolfram. Experimental evidence on the economics of rural electrification. *Journal of Political Economy*, 128:1523–1565, 2020. ISSN 1537534X. doi: 10.1086/705417.

Mindy Leow and Christophe Mues. Predicting loss given default (lgd) for residential mortgage loans: A two-stage model and empirical evidence for uk bank data. *International Journal of Forecasting*, 28:183–195, 2012. ISSN 01692070. doi: 10.1016/j.ijforecast.2011.01.010. URL http://dx.doi.org/10.1016/j.ijforecast.2011.01.010.

Geng Li and Paul A Smith. Borrowing from yourself: 401(k) loans and household balance sheets. *FRB Finance and Economics Discussion Series*, 2008. URL https://www.federalreserve.gov/pubs/feds/2008/200842/200842pap.pdf.

Geng Li and Paul A Smith. New evidence on 401(k) borrowing and household balance sheets, 2009. URL https://www.federalreserve.gov/pubs/feds/2009/200919/200919pap.pdf.

Jeffrey B Liebman, F P Luttmer, and David G Seif. Labor supply responses to marginal social security benefits: Evidence from discontinuities. *Journal of Public Economics*, 93:1208–1223, 2009. doi: 10.1016/j.jpubeco.2009.07.006. URL https://users.nber.org/~luttmer/ssbenefitlink.pdf.

David A Love, Michael G Palumbo, and Paul A Smith. The trajectory of wealth in retirement, 2008. URL http://ssrn.com/abstract=1055641http://www.bc.edu/crr.

Jens Ludwig, Greg J. Duncan, Lisa A. Gennetian, Lawrence F. Katz, Ronald C. Kessler, Jeffrey R. Kling, and Lisa Sanbonmatsu. Neighborhood effects on the long-term well-being of low-income adults. *Science*, 337:1505, 9 2012. doi: 10.1126/SCIENCE.1224648. URL /pmc/articles/PMC3491569//pmc/articles/PMC3491569/?report=abstracthttps://www.ncbi.nlm.nih.gov/pmc/articles/PMC3491569/.

Annamaria Lusardi and Olivia S Mitchell. Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth. *Journal of Monetary Economics*, 54:205–224, 2007. doi: 10.1016/j.jmoneco.2006.12.001. URL www.elsevier.com/locate/jme.

Annamaria Lusardi and Peter Tufano. Debt literacy, financial experiences, and overindebtedness, 2009. URL http://hdl.handle.net/10419/43242www.econstor.eu.

Nicole Maestas, Kathleen J Mullen, and Alexander Strand. The effect of economic conditions on the disability insurance program: Evidence from the great recession, 2018. URL http://www.nber.org/papers/w25338.

Allison Mann. The effect of late-life debt use on retirement decisions. *Social Science Research*, 40:1623–1637, 11 2011. ISSN 0049-089X. doi: 10.1016/J.SSRESEARCH.2011.05.004. URL https://www.sciencedirect.com/science/article/pii/S0049089X11001025.

Day Manoli and Andrea Weber. Nonparametric evidence on the effects of financial incentives on retirement decisions †. *American Economic Journal: Economic Policy*, 8:160–182, 2016. doi: 10.1257/pol.20140209. URL http://dx.doi.org/10.1257/pol.20140209.

Giovanni Mastrobuoni. The social security earnings test removal: Money saved or money spent by the trust fund?, 2006. URL https://www.princeton.edu/gceps/workingpapers/133mastrobuoni.pdf.

Giovanni Mastrobuoni. Labor supply effects of the recent social security benefit cuts: Empirical estimates using cohort discontinuities. *Journal of Public Economics*, 93:1224–1233, 2009. doi: 10.1016/j.jpubeco.2009.07.009. URL https://pdf.sciencedirectassets.com/271705/1-s2.0-S0047272709X0006X/1-s2.0-S0047272709000875/main.pdf?x-amz-security-token=AgoJb3JpZ2luX2VjEMX%2F%2F%2F%2F%2F%2F%2F%2F%2F%2FwEaCXVzLWVhc3QtMSJIMEYCIQCodpYwbzEmf5siqFsSUltWc1VTpmSAxksSnjHs3GYuhAIhANM%2BX3NvK8.

Brooke Helppie McFall. Crash and wait? the impact of the great recession on the retirement plans of older americans. *American Economic Review*, 101:40–44, 5 2011. ISSN 0002-8282. doi: 10.1257/aer.101.3.40. URL `http://pubs.aeaweb.org/doi/10.1257/aer.101.3.40`.

Bruce D Meyer and James X Sullivan. Consumption and income inequality and the great recession. *American Economic Review*, 103:178–183, 5 2013. ISSN 0002-8282. doi: 10.1257/aer.103.3.178. URL `http://pubs.aeaweb.org/doi/10.1257/aer.103.3.178`.

Atif Mian and Amir Sufi. The consequences of mortgage credit expansion: Evidence from the u.s. mortgage default crisis *. *Quarterly Journal of Economics*, 124:1449–1496, 11 2009. ISSN 0033-5533. doi: 10.1162/qjec.2009.124.4.1449. URL `https://academic.oup.com/qje/article-lookup/doi/10.1162/qjec.2009.124.4.1449`.

Atif Mian and Amir Sufi. The great recession: Lessons from microeconomic data. *American Economic Review: Papers and Proceedings*, 100:51–56, 2010. doi: 10.1257/aer.100.2.51. URL `http://www.aeaweb.org/articles.php?doi=10.1257/aer.100.2.51`.

Atif Mian and Amir Sufi. House prices, home equity–based borrowing, and the us household leverage crisis. *American Economic Review*, 101:2132–56, 2011. doi: 10.1257/aer.101.5.2132. URL `http://www.aeaweb.org/articles.php?doi=10.1257/aer.101.5.2132`.

Atif Mian and Amir Sufi. Credit supply and housing speculation, 2018. URL `http://www.nber.org/papers/w24823`.

Alicia H. Munnell. 401(k) plans in 2010: An update from the scf, 2012. URL `http://crr.bc.edu/wp-content/uploads/2012/07/IB_12-13-508.pdf`.

Alicia H Munnell and Matthew S Rutledge. The effects of the great recession on the retirement security of older workers. *Annals of the AAPSS*, 650:124–142, 2013. doi: 10.1177/0002716213499535. URL `http://journals.sagepub.com/doi/pdf/10.1177/0002716213499535`.

Makoto Nakajima and Irina A Telyukova. Home equity in retirement *, 2011. URL `http://ssrn.com/abstract=1898330Electroniccopyavailableat:http://ssrn.com/abstract=1898330`.

Mariacristina De Nardi, Eric French, and John B Jones. Why do the elderly save? the role of medical expenses. *Journal of Political Economy*, 118:48, 2010. URL `http://www.journals.uchicago.edu/t-and-c`.

Scott T Nelson. Private information and price regulation in the us credit card market, 2018.

Robert Novy-Marx and Joshua Rauh. Public pension promises: How big are they and what are they worth? *The Journal of Finance*, 66:1211–1249, 8 2011. ISSN 00221082. doi: 10.1111/j.1540-6261.2011.01664.x. URL `http://doi.wiley.com/10.1111/j.1540-6261.2011.01664.x`.

Robert Novy-Marx and Joshua D. Rauh. The liabilities and risks of state-sponsored pension plans. *Journal of Economic Perspectives*, 23:191–210, 11 2009. ISSN 0895-3309. doi: 10.1257/jep.23.4.191. URL http://pubs.aeaweb.org/doi/10.1257/jep.23.4.191.

Arna Olafsson and Michaela Pagel. The retirement-consumption puzzle: New evidence from personal finances, 2018. URL http://www.nber.org/papers/w24405.

Jan Ondrich and Alexander Falevich. The great recession, housing wealth, and the retirement decisions of older workers. *Public Finance Review*, 44:109–131, 2016. doi: 10.1177/1091142114551600. URL http://journals.sagepub.com/doi/pdf/10.1177/1091142114551600.

Amine Ouazad and Romain Rancière. Market frictions, arbitrage, and the capitalization of amenities, 2019. URL http://www.nber.org/papers/w25701.

Kevin A. Park. Fha loan performance and adverse selection in mortgage insurance. *Journal of Housing Economics*, 34:82–97, 12 2016. ISSN 10960791. doi: 10.1016/j.jhe.2016.07.004. URL http://dx.doi.org/10.1016/j.jhe.2016.07.004.

Kevin A. Park. Temporary loan limits as a natural experiment in federal housing administration insurance. *Housing Policy Debate*, 27:449–466, 5 2017. ISSN 1051-1482. doi: 10.1080/10511482.2016.1234501. URL https://www.tandfonline.com/doi/full/10.1080/10511482.2016.1234501.

Wayne Passmore and Shane M Sherlund. Government-backed mortgage insurance, financial crisis, and the recovery from the great recession, 2016. URL http://dx.doi.org/10.17016/FEDS.2016.031.

Anthony Pennington-Cross. Subprime and prime mortgages: Loss distributions. *Office of Federal Housing Enterprise Oversight*, 2003.

James Poterba, Steven Venti, and David Wise. The composition and drawdown of wealth in retirement. *Journal of Economic Perspectives*, 25:95–118, 2011. doi: 10.1257/jep.25.4.95. URL https://pubs.aeaweb.org/doi/pdfplus/10.1257/jep.25.4.95.

Jim Puzzanghera. Fha to get $1.7 billion in its first taxpayer-funded bailout - los angeles times, 9 2013. URL https://www.latimes.com/business/la-xpm-2013-sep-28-la-fi-0928-fha-bailout-20130928-story.html.

Min Qi and Xiaolong Yang. Loss given default of high loan-to-value residential mortgages. *Journal of Banking and Finance*, 33:788–799, 2009. ISSN 03784266. doi: 10.1016/j.jbankfin.2008.09.010. URL http://dx.doi.org/10.1016/j.jbankfin.2008.09.010.

John M Quigley. Interest rate variations, mortgage prepayments and household mobility. *Review of Economics and Statistics*, 69:636–643, 1987. URL https://www.jstor.org/stable/pdf/1935958.pdf?refreqid=excelsior%3Afbeebb55afdcf54ae858efb4be23170e.

Shanthi Ramnath. Taxpayers' responses to tax-based incentives for retirement savings: Evidence from the saver's credit notch. *Journal of Public Economics*, 101:77–93, 5 2013. ISSN 00472727. doi: 10.1016/j.jpubeco.2013.02.010.

Joshua D Rauh. Investment and financing constraints: Evidence from the funding of corporate pension plans. *The Journal of Finance*, LXI:33–71, 2006. URL https://web.stanford.edu/~rauh/research/rauh_jf_2006.pdf.

Joshua D. Rauh. Risk shifting versus risk management: Investment policy in corporate pension plans. *Review of Financial Studies*, 22:2687–2733, 7 2009. ISSN 0893-9454. doi: 10.1093/rfs/hhn068. URL https://academic.oup.com/rfs/article-lookup/doi/10.1093/rfs/hhn068.

Jennifer Roback. Wages, rents, and the quality of life. *Journal of Political Economy*, 90: 1257–1278, 1982.

Claudia Robles-Garcia. Competition and incentives in mortgage markets: the role of brokers, 2018. URL https://www.fca.org.uk/.

Harvey S Rosen and Stephen Wu. Portfolio choice and health status. *Journal of Financial Economics*, 72:457–484, 6 2004. ISSN 0304-405X. doi: 10.1016/S0304-405X(03)00178-8. URL https://www.sciencedirect.com/science/article/pii/S0304405X03001788.

Sherwin Rosen. Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82:34–55, 1974.

Andrea Ryan, Gunnar Trumbull, and Peter Tufano. A brief postwar history of u.s. consumer finance. *Business History Review*, 85:461–498, 9 2011. ISSN 0007-6805. doi: 10.1017/S0007680511000778. URL http://www.journals.cambridge.org/abstract_S0007680511000778.

Emmanuel Saez. Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy*, 2:180–202, 2010. doi: 10.1257/pol.2.3.180. URL http://www.aeaweb.org/articles.php?doi=10.1257/pol.2.3.180.

James M. Sallee and Joel Slemrod. Car notches: Strategic automaker responses to fuel economy policy. *Journal of Public Economics*, 96:981–999, 12 2012. ISSN 00472727. doi: 10.1016/j.jpubeco.2012.06.005.

David S. Scharfstein. Presidential address: Pension policy and the financial system. *The Journal of Finance*, 73:1463–1512, 8 2018. ISSN 00221082. doi: 10.1111/jofi.12710. URL http://doi.wiley.com/10.1111/jofi.12710.

John B Shoven and Sita Nataraj Slavov. The decision to delay social security benefits: Theory and evidence, 2012. URL http://www.nber.org/papers/w17866.

Todd Sinai and Nicholas S Souleles. Net worth and housing equity in retirement, 2007. URL http://www.nber.org/papers/w13693.

Jonathan Skinner. Housing and saving in the united states, 1991. URL https://www.nber.org/papers/w3874.pdf.

Hyrum L Smith, Michael S Finke, and Sandra J Huston. Financial sophistication and housing leverage among older households. *Journal of Family and Economic Issues*, 33:315–327, 2012. doi: 10.1007/s10834-012-9293-4. URL https://link.springer.com/content/pdf/10.1007%2Fs10834-012-9293-4.pdf.

Sarah Smith. Can the retirement-consumption puzzle be resolved? evidence from the british household panel survey, 2005. URL https://pdfs.semanticscholar.org/e18f/b10a88a6c4ac1dc7c192c701045a7a3ece7b.pdf.

Social Security Administration. Retirement benefits: Starting your retirement benefits early. URL https://www.ssa.gov/planners/retire/agereduction.html.

Jeff Sommer. Social security is staring at its first real shortfall in decades, 6 2019.

Richard H Thaler and Shlomo Benartzi. Save more tomorrow: Using behavioral economics to increase employee saving. *Journal of Political Economy*, 112, 2004. URL http://www.journals.uchicago.edu/t-and-c.

Steven F. Venti and David A. Wise. Aging and the income value of housing wealth. *Journal of Public Economics*, 44:371–397, 4 1991. ISSN 0047-2727. doi: 10.1016/0047-2727(91)90020-3. URL https://www.sciencedirect.com/science/article/pii/0047272791900203.

Aaron S Yelowitz. The medicaid notch, labor supply, and welfare participation: Evidence from eligibility expansions, 1995. URL https://www.jstor.org/stable/2946644.

APPENDIX A

SUPPLEMENTAL MATERIAL TO CHAPTER 2

A.1 Derivations

A.1.1 Solution to the HH Problem

To solve the HH problem, I begin by forming the Lagrangian:

$$\begin{aligned}\mathcal{L}\big(\{c_t\}_t, \{a_t\}_t, s, B^o\big) \quad = \quad & u(c_0, s_0) + \sum_{t=1}^{\infty} \beta^t u(c_t, s) \\ & + \lambda_0 \Big[a_0 + w_0 - c_0 - [P(s) - B^o] - \frac{a_1}{1+r} \Big] \\ & + \sum_{t=1}^{\infty} \lambda_t \Big[a_t + w_t - c_t - B^o r^m (B^o) - \frac{a_{t+1}}{1+r} \Big] \\ & + \sum_{t=1}^{\infty} \mu_t a_t \end{aligned} \tag{A.1}$$

By the Kuhn-Tucker theorem, the solution to the HH problem is:

$$\Big\{ B^{o*}, s^*, \{c_t^*\}_t, \{a_t^*\}_t, \{\lambda_t\}_t, \{\mu_t\}_t \Big\}$$

where the household choice variables satisfy the following first-order conditions:

$$\begin{aligned} &\text{FOC}[c_t] && \beta^t u_c^t = \lambda_t && \forall t \geq 0 \\ &\text{FOC}[a_t] && \lambda_t + \mu_t = \frac{\lambda_{t-1}}{1+r} && \forall t > 0 \\ &\text{FOC}[B^o] && \lambda_0 = \sum_{t=1}^{\infty} \lambda_t [B^o r^m (B^o)]' \Big|_{B^{o*}} && \\ &\text{FOC}[s] && \lambda_0 P'(s^*) = \sum_{t=1}^{\infty} \beta^t u_s && \end{aligned} \tag{A.2}$$

And the Lagrange multipliers satisfy:

$$\begin{aligned} \lambda_t > 0 \quad \forall t \geq 0 \\ a_t > 0 \quad \text{or} \quad \mu_t > 0 \quad \forall t > 0 \end{aligned} \tag{A.3}$$

I further assume that (i) the period utility function, u, is separable in the consumption good and amenity (ii) $\beta(1+r) = 1$ and (iii) the credit constraints are non-binding after the initial savings decision, $\mu_t = 0 \; \forall \; t > 1$.

I obtain the necessary conditions for household optimization as follows:

$$P'(s^*) \overset{\text{FOC}[s]}{=} \frac{1}{r}\frac{u_s}{\lambda_0} \underset{\beta(1+r)=1}{\overset{\text{FOC}[a_1]}{=}} \frac{1}{r}\frac{u_s}{(1+r)(\lambda_1+\mu_1)} \underset{\beta(1+r)=1}{\overset{\text{FOC}[c_0]}{=}} \frac{1}{r}\frac{u_s}{u_c^1}\frac{1}{1+\frac{\mu_1}{\lambda_1}} \tag{A.4}$$

$$1+\frac{\mu_1}{\lambda_1} \overset{\text{FOC}[a_1]}{=} \frac{1}{1+r}\frac{\lambda_0}{\lambda_1} \overset{\text{FOC}[B^o]}{=} \frac{1}{1+r}\sum_{t=1}^{\infty}\frac{\lambda_t}{\lambda_1}\left[B^o r^m(B^o)\right]'\bigg|_{B^{o*}} \overset{\text{FOC}[a_t]}{=} \left[\frac{B^o r^m(B^o)}{r}\right]'\bigg|_{B^{o*}} \tag{A.5}$$

A.1.2 Estimating HH WTP for Credit

The consumer solves the following problem:

$$
\begin{aligned}
U_i &= \max_j U_{ij} \\
&\text{Where:} \quad U_{ij} = \max_{\{c_{ijt}, a_{ijt}\}_t} \sum_{t=0}^{T} \beta^t u(c_{ijt}) \\
&\qquad\qquad \text{s.t.} \quad c_{i0} = y_{i0} + a_{i0} - P_i + B^o_{ij} - \frac{a_{i1}}{1+r} && (\lambda_0) \\
&\qquad\qquad\qquad c_{it} = y_{it} + a_{it} - B^o_{ij} r^m_{ij} - \frac{a_{i,t+1}}{1+r} \qquad \forall t > 0 && (\lambda_t) \\
&\qquad\qquad\qquad a_{it} \geq 0 \qquad \forall t > 0 && (\mu_t) \\
&\qquad\qquad = \max_j u(c^*_{ij0}) + \sum_{t=1}^{\infty} \beta^t u(c^*_{ijt}) \\
&\qquad\qquad \text{Where:} \quad \beta^t u'(c^*_{ijt}) = \lambda_t \qquad \forall t \geq 0 \\
&\qquad\qquad\qquad \lambda_t + \mu_t = \frac{\lambda_{t-1}}{1+r} \qquad \forall t > 0 \\
&\qquad\qquad \text{And assume:} \quad \mu_t = 0 \qquad \forall t > 1
\end{aligned}
\tag{A.6}
$$

Comparing two financing methods, the consumer considers:

$$
\begin{aligned}
U_{ij'} - U_{ij} \quad &= \quad u(c^*_{ij'0}) - u(c^*_{ij0}) + \sum_{t=1}^{\infty} \beta^t \left[u(c^*_{ij't}) - u(c^*_{ijt}) \right] \\
&\overset{\text{Env. Thm.}}{=} \quad u'(c^*_{ij0}) \left[c^*_{ij'0} - c^*_{ij0} \right] + \sum_{t=1}^{\infty} \beta^t u'(c^*_{ijt}) \left[c^*_{ij't} - c^*_{ijt} \right] \\
&\overset{\text{FOC}[c]}{=} \quad \lambda_0 \Delta B^o - \sum_{t=1}^{\infty} \lambda_t \left[\Delta B^o \hat{r}^m_j + B^o_j \Delta \hat{r} \right] \\
&\overset{\text{FOC}[a]}{=} \quad (1+r)(\lambda_1 + \mu_1) \Delta B - \lambda_1 \left[\Delta B^o \hat{r}^m_j + B^o_j \Delta \hat{r}^m \right] \sum_{t=1}^{\infty} \frac{1}{(1+r)^{t-1}} \\
&= \quad (1+r)(\lambda_1 + \mu_1) \Delta B - \lambda_1 \left[\Delta B^o \hat{r}^m_j + B^o_j \Delta \hat{r}^m \right] \frac{1+r}{r} \\
&\propto \quad \Delta B^o - \frac{\Delta B^o \hat{r}^m_j + B^o_j \Delta \hat{r}^m}{r} + \frac{\mu_1}{\lambda_1} \Delta B^o
\end{aligned}
\tag{A.7}
$$

The difference in utility is proportional to an expression that is intuitive. The first term, ΔB^o, is the additional funds supplied by the contract j' relative to j. The second term, $\frac{\Delta B^o \hat{r}^m_j + B^o_j \Delta \hat{r}^m}{r}$, is the objectively discounted value of the additional mortgage payment obligations demanded by the contract j' relative to j. The borrower also benefits from the third term, $\frac{\mu_1}{\lambda_1}\Delta B^o$, which makes the additional funds from j' more attractive to the extent that the borrower is borrowing constrained.

For contracts j and j', we have:

$$j' \succ j \quad \Longleftrightarrow \quad \kappa^1 > \frac{\hat{r}^m_j}{r}\left(1 + \frac{\%\Delta_j^{j'}\hat{r}^m}{\%\Delta_j^{j'}B^o}\right) \tag{A.8}$$

Where everything on the right side of the $\succ$ can be observed in the data. We can infer something about the extent of borrower's credit constraints from information about their choice set, which is made available in the GSE and private mortgage insurance pricing grids.

A.1.3 Housing Market Equilibrium

To solve for the implicitly defined equilibrium prices, I solve the HH problem over all variables but the choice of amenity level. In particular, I begin by forming the Lagrangian:

$$\begin{aligned} \mathcal{L}(c_0, c_1, a_1) &= u(c_0) + \beta c_1 + v(s) \\ &\quad + \lambda_0\Big[y_0 - c_0 - \widetilde{P}(s) - \frac{a_1}{1+r}\Big] \\ &\quad + \lambda_1\Big[y_1 + a_1 - c_1\Big] \\ &\quad + \mu_1\Big[a_1 + \phi\Big] \end{aligned} \tag{A.9}$$

By the Kuhn-Tucker theorem, the solution to the HH problem is $\left\{c_0^*, c_1^*, a_1^*, \lambda_0, \lambda_1, \mu_1\right\}$, where the household choice variables satisfy the following first-order conditions:

$$\begin{aligned} &\text{FOC}[c_0] \quad u'(c_0^*) = \lambda_0 \\ &\text{FOC}[c_1] \quad \beta = \lambda_1 \\ &\text{FOC}[a_1] \quad \lambda_1 + \mu_1 = \frac{\lambda_0}{1+r} \end{aligned} \tag{A.10}$$

And the Lagrange multipliers satisfy:

$$\begin{aligned} \lambda_t > 0 \quad t \in \{0,1\} \\ a_1 > 0 \quad \text{or} \quad \mu_1 > 0 \qquad \forall t > 0 \end{aligned} \tag{A.11}$$

I further assume that $\beta(1+r) = 1$.

I solve for the equilibrium price schedule in two regimes. In case 1, I assume that households are all unconstrained at the time of home purchase. In case 2, I assume that all households are at their borrowing limits at the time of home purchase.

- **Case 1** I use the household's first order conditions, budget constraints, and the assumption that the household's shadow price of credit is zero, $\mu_1 = 0$, to obtain expressions for optimal consumption:

$$c_0^* \overset{\text{FOC}[c_0]}{=} u'^{-1}(\lambda_0) \underset{\mu_1=0}{\overset{\text{FOC}[a_1]}{=}} u'^{-1}\big((1+r)\lambda_1\big) \overset{\text{FOC}[c_1]}{=} u'^{-1}\big((1+r)\beta\big) \overset{\beta(1+r)=1}{=} u'^{-1}(1) \tag{A.12}$$

$$c_1^* \overset{(\lambda_1)}{=} y_1 + a_1^* \overset{(\lambda_0)}{=} y_1 + (1+r)\big(y_0 - c_0^* - P(s)\big) \overset{w_0 \equiv y_0 + \frac{y_1}{1+r}}{=} (1+r)\big(w_0 - P(s) - c_0^*\big) \tag{A.13}$$

I obtain the implied equilibrium price schedule by plugging optimal consumption levels into the household's objective and rearranging:

$$\widetilde{P}(s) = w^o - u'^{-1}(1) - \Big[\overline{U} - u(u'^{-1}(1)) - v(s)\Big] \tag{A.14}$$

Taking derivatives, we obtain the relation between the slope of the price schedule and the household willingness-to-pay for the amenity:

$$\widetilde{P}'(s) = \frac{v'(s)}{1} \overset{u'(c_0^*)=1}{=} \frac{v'(s)}{u'(c_0^*)} \tag{A.15}$$

- **Case 2** I use the requirement that $a_1^* = 0$ and the household period budget constraints to obtain expressions for optimal consumption:

$$c_0^* = y_0 - \widetilde{P}(s) + \frac{\phi}{1+r} \tag{A.16}$$

$$c_1^* = y_1 - \phi \tag{A.17}$$

Again, I obtain the implied equilibrium price schedule by plugging optimal consumption into the objective function and rearranging:

$$\widetilde{P}(s) = y_0 + \frac{\phi}{1+r} - u^{-1}\Big(\overline{U} - \beta(y_1 - \phi) - v(s)\Big) \tag{A.18}$$

Now, the slope of the equilibrium price schedule reflects:

$$\widetilde{P}'(s) = \frac{v'(s)}{u'(c_0^*)} = \frac{v'(s)}{1 + \frac{\mu_1}{\lambda_1}} = \frac{v'(s)}{\kappa^1} \tag{A.19}$$

Note that although these were solved as two cases, the characterization of the equilibrium price schedule slope under Case 2 nests the result in Case 1 as a special case. In particular, an unconstrained household has $\mu_1 = 0$ and the denominator simplifies to 1.

A.1.4 The Government's Problem

To solve the government's problem, I begin by rewriting its objective function:

$$\begin{aligned}
\int_i U_i(\sigma,t)di &\overset{U_i \text{ def.}}{=} \int_i \Big[u(c_0^{i*}) + \beta(c_1^{i*} - t) + v(s^{i*} + \sigma)\Big]di \\
&\overset{\text{Taylor}}{=} \int_i \Big[u(c_0^{i*}) + \beta(c_1^{i*}) + v(s^{i*})\Big]di + \int_i \Big[\sigma v'(s^{i*}) - \beta t\Big]di \\
&\overset{\widetilde{P}(s) \text{ def.}}{=} \int_i \overline{U} di + \int_i \Big[\sigma v'(s^{i*}) - \beta t\Big]di \\
&\overset{\int_i di=1; \, \mathbb{E} \text{ def.}}{=} \overline{U} + \sigma \mathbb{E}^i\Big[v'(s^{i*})\Big] - \beta t \\
&\overset{\widetilde{P}'(s)=\frac{v'(s)}{\kappa^1}}{=} \overline{U} + \sigma \mathbb{E}^i\Big[\widetilde{P}'(s^{i*})\kappa_i^1\Big] - \beta t \\
&\overset{(\text{BC}),\, \beta(1+r)=1}{=} \overline{U} + \sigma \mathbb{E}^i\Big[\widetilde{P}'(s^{i*})\kappa_i^1\Big] - I(\sigma)
\end{aligned} \tag{A.20}$$

The first-order condition is now:

$$\text{FOC}[\sigma] \qquad \mathbb{E}^i\Big[\widetilde{P}'(s^{i*})\kappa_i^1\Big] = I'(\sigma^*) \tag{A.21}$$

The optimal amount of amenity improvement is:

$$\sigma^* = I'^{-1}\Big(\mathbb{E}^i\Big[\widetilde{P}'(s^{i*})\kappa_i^1\Big]\Big) \tag{A.22}$$

Suppose the government chooses the level of investment and amenity improvement according to traditional hedonic estimates. Then it chooses the following level of amenity improvement:

$$\sigma^g = I'^{-1}\Big(\mathbb{E}^i\Big[\widetilde{P}'(s^{i*})\Big]\Big) \tag{A.23}$$

In the case that households are unconstrained, the government investing according to traditional hedonic estimates chooses the optimal level of investment and amenity improvement:

$$\sigma^g \overset{\sigma^g \, def.}{=} I'^{-1}\Big(\mathbb{E}^i\Big[\widetilde{P}'(s^{i*})\Big]\Big) \overset{\mu_1=0}{=} I'^{-1}\Big(\mathbb{E}^i\Big[\widetilde{P}'(s^{i*})\kappa_i^1\Big]\Big) \overset{\sigma^* \, def.}{=} \sigma^* \tag{A.24}$$

In the case that households are constrained, the government investing according to traditional hedonic estimates chooses amenity improvement below the optimal level:

$$\sigma^g \overset{\sigma^g \, def.}{=} I'^{-1}\Big(\mathbb{E}^i\Big[\widetilde{P}'(s^{i*})\Big]\Big) \overset{\mu_1>0;\, I''>0}{<} I'^{-1}\Big(\mathbb{E}^i\Big[\widetilde{P}'(s^{i*})\kappa_i^1\Big]\Big) \overset{\sigma^* \, def.}{=} \sigma^* \tag{A.25}$$

A.2 "Financial" Hedonic Regression

A.2.1 Framework

Consider a financial product, f, characterized by a price, p^f, and state- and time- indexed payoffs, d_{st}^f, for dates $t \in \{1, ..., T\}$ and states $s \in \{1, ..., S\}$ with realization probabilities, $\{\pi_{st}\}_{s \times t}$. And consider a menu of available products, F. For simplicity, assume the space of available products is continuous. For a given household, i, with stochastic discount factor, M_t^i, the surplus from the choice of product f is:

$$\Pi^{if} = \sum_{t=1}^{T} \sum_{s=1}^{S} \pi_{st} E[M_t^i \mid s] d_{st}^f - p^f \tag{A.26}$$

The optimizing household will choose financial product, f^*, so that the marginal surplus from adjusting any of the payoffs is zero:

$$\frac{\partial}{\partial d_{st}^f} \Pi^{if} \bigg|_{f^*} = \frac{\partial}{\partial d_{st}^f} \Big[\sum_{t=1}^{T} \sum_{s=1}^{S} \pi_{st} E[M_t^i \mid s] d_{st}^f - p^f \Big] \bigg|_{f^*} = \pi_{st} E[M_t^i \mid s] - \frac{\partial p^f}{\partial d_{st}^f} \bigg|_{f^*} = 0 \tag{A.27}$$

The slope of the menu of financial products where the borrower chooses to locate reveals the borrower's state-price. This motivates a "financial" hedonic regression of the form:

$$p_i^f \sim \alpha + \beta' d_i^f + \varepsilon_i \tag{A.28}$$

The regression coefficients then recover the cross sectional average household state price.

$$\beta_{st} = E^i\Big[\pi_{st} E[M_t^i \mid s]\Big] \tag{A.29}$$

In the simple HH problem introduced above, the financial product is the mortgage. The price is the time-1 discounted value of mortgage obligations, $\frac{B^o r^m(B^o)}{r_f}(1+r^f)$, and the payoff is the time-0 mortgage balance, B^o. The household's discount factor from time-0 forward to time-1 is deterministic and takes the form $1/M_1^i = \lambda_0^i/\lambda_1^i = (1+r^f)[1+\mu_1^i/\lambda_1^i]$.

Discounting mortgage obligations instead to time-0 eliminates the $(1+r^f)$ term in both the price and the regression coefficient. This suggests regression specification:

$$\frac{B_i^o r^m(B_i^o)}{r_f} \sim \alpha + \beta B_i^o + \varepsilon_i \tag{A.30}$$

Where the regression coefficient recovers:

$$\beta = E^i\Big[1 + \mu_1^i/\lambda_1^i\Big] \tag{A.31}$$

A.2.2 Correcting "Traditional" Hedonic Estimates

The analytical framework in this paper suggests (i) that traditional hedonic estimates actually measure household willingness-to-pay for amenities with a bias term related to household willingness-to-pay for credit and (ii) household willingness-to-pay for credit at home purchase can be measured from information on household mortgage choice. In this section, I describe how to use the information on household mortgage choice to correct hedonic estimates to recover the estimate of policy relevance, the unbiased household willingness-to-pay for amenities.

Cross-sectional heterogeneity in households poses a challenge to correcting hedonic estimates. It is not enough to know the mean willingness-to-pay for credit in order to correct the bias. A variance decomposition shows the relationship between the biased and unbiased estimates of willingness-to-pay for amenities:

$$E^i\Big[\frac{u_s^i/u_c^i}{rf}\Big] = \Big(E^i\Big[\frac{1}{1+\mu_1^i/\lambda_1^i}\Big]\Big)^{-1}\Big(E^i\Big[\frac{u_s^i/u_c^i}{rf}\frac{1}{1+\mu_1^i/\lambda_1^i}\Big] + Cov^i\Big(\frac{u_s^i/u_c^i}{rf}, \frac{1}{1+\mu_1^i/\lambda_1^i}\Big)\Big) \tag{A.32}$$

Directly computing the covariance term, $Cov^i\big(\frac{u_s^i/u_c^i}{rf}, \frac{1}{1+\mu_1^i/\lambda_1^i}\big)$, poses a challenge in particular, because amenity price schedules are often not well-defined, making it difficult to compute borrower-level estimates of willingness-to-pay for amenities. Hedonic regressions, for instance, settle for population averages.

My approach is to exploit the fact that the menu of options for household mortgage choice is more transparent and that willingness-to-pay for credit can be estimated at the household level, per Section (A.1.2). I also use the intuition that the logic of hedonic regression, that it captures the slope of the price schedule at the place where borrowers locate on it, applies equally well to subsets of the population as it does to the population as a whole.

I consider a cross-section of households, i, located in various school districts, d. I can measure their willingness to pay for credit at the time of home purchase, $\kappa_i = 1 + \mu_1^i/\lambda_1^i$, from their mortgage menu and mortgage choice, and I bin them according to the size of this wedge, $j(i)$. I regress the price of the household's home, P_{id}, on a measure of the school quality in the district, s_d, normalized by the household's willingness to pay for credit at the time of home purchase, κ_i. I also include fixed effects for the credit wedge bin, $\alpha_{j(i)}$:

$$P_{id} \sim \alpha_{j(i)} + \beta\frac{s_d}{\kappa_i} + \varepsilon_{id} \tag{A.33}$$

Intuitively, with the presence of the fixed-effects, the identifying variation comes from within credit wedge bins. The within-bin variation is roughly captured by a within-bin regression slope coefficient and these are weighted by the proportion of the sample size in each bin. This suggests the following representation of the regression slope coefficient:

$$\beta = \sum_j Pr(\kappa_i = \overline{\kappa}_j) \frac{Cov(P_{id}, \frac{s_d}{\kappa_i} \mid \kappa_i = \overline{\kappa}_j)}{Var(\frac{s_d}{\kappa_i} \mid \kappa_i = \overline{\kappa}_j)} \tag{A.34}$$

Using the conditioning information and properties of variances and covariances, I simplify:

$$= \sum_j \overline{\kappa}_j Pr(\kappa_i = \overline{\kappa}_j) \frac{Cov(P_{id}, s_d \mid \kappa_i = \overline{\kappa}_j)}{Var(s_d \mid \kappa_i = \overline{\kappa}_j)} \tag{A.35}$$

I observe that the $\frac{Cov}{Var}$ term is the slope coefficient from an 'uncorrected' traditional hedonic regression on a subset of the population, i.e. with a credit wedge of $\overline{\kappa}_j$. I apply the interpretation of traditional hedonic regressions, amended by the analysis in the paper, and applied to a population subset, to obtain:

$$= \sum_j \overline{\kappa}_j Pr(\kappa_i = \overline{\kappa}_j) E^i \left[\frac{u_s^i / u_c^i}{r^f} \frac{1}{\kappa_i} \middle| \kappa_i = \overline{\kappa}_j \right] \tag{A.36}$$

Again, I use the conditioning information and properties of expectations to simplify algebraically:

$$= \sum_j Pr(\kappa_i = \overline{\kappa}_j) E^i \left[\frac{u_s^i / u_c^i}{r^f} \middle| \kappa_i = \overline{\kappa}_j \right] \tag{A.37}$$

And, finally, I use the definition of conditional and unconditional expectations to write:

$$= E^i \left[\frac{u_s^i / u_c^i}{r^f} \right] \tag{A.38}$$

This is the unbiased marginal willingness-to-pay for amenities, the policy-relevant estimand, and the object typically, though I argue wrongly, supposed to be estimated by traditional hedonic regressions.

A.3 Appendix Tables & Figures

Table A.1: Summary Statistics for CRISM Sample

	Purchase N=5328	Rate/Term Refi N=1799	Cash-out Refi N=1862
Loan Characteristics:			
Loan Amount ($k)	226 (123)	260 (139)	219 (112)
Property Value ($k)	292 (180)	413 (336)	362 (236)
Interest Rate (%)	4.97 (1.18)	4.28 (0.97)	5.25 (1.20)
LTV (%)	81.0 (15.1)	70.0 (19.0)	65.0 (16.0)
DTI (%)	35.3 (13.3)	35.5 (16.1)	37.0 (15.0)
FICO Score	741 (52.9)	752 (52.1)	729 (56.6)
PMI Characteristics:			
Has PMI:			
No	60.8%	84.3%	93.4%
Yes	39.2%	15.7%	6.61%
Documentation Type:			
Full Doc	75.4%	76.6%	72.7%
Low Doc	16.3%	21.4%	17.1%
No Doc	8.29%	0.34%	6.91%
MI Company:			
Arch	4.33%	3.56%	0.96%
Essent	5.71%	7.12%	0.96%
GE	14.4%	10.4%	7.69%
MGIC	18.9%	13.3%	12.5%
Radian	17.3%	9.39%	21.2%
UGIC	12.6%	6.47%	11.5%
Other	26.8%	49.8%	45.2%

The CRISM Sample consists of loans from BlackKnight mortgage data with the following properties: originated 2005-2020, first liens, single-family, owner-occupied, 1-4 unit dwelling, conventional, conforming, vanilla (no balloon or interest-only payments), fixed-rate, and 30 year terms. For computational ease, only a 7pp random sample of the data are used in analyses. For computational ease, of this, only a 1pp random sub-sample is used to compute the summary statistics above.

Table A.2: Detailed Summary Statistics for CRISM Sample

	N	$\overline{x}$	s_x	min	p25	p50	p75	max
Loan Characteristics:								
Loan Amount ($k)	9k	230	120	9	140	200	300	810
Property Value ($k)	9k	330	240	16	180	280	420	7700
Interest Rate (%)	9k	4.9	1.2	1	4	4.8	5.9	9.9
LTV (%)	9k	76	18	3	67	79	89	120
DTI (%)	7.2k	36	14	1	26	35	43	99
FICO Score	8.5k	740	54	0	700	750	780	840

The CRISM Sample consists of loans from BlackKnight mortgage data with the following properties: originated 2005-2020, first liens, single-family, owner-occupied, 1-4 unit dwelling, conventional, conforming, vanilla (no balloon or interest-only payments), fixed-rate, and 30 year terms. For computational ease, only a 7pp random sample of the data are used in analyses. For computational ease, of this, only a 1pp random sub-sample is used to compute the summary statistics above.

Figure A.1: Descriptive Statistics by FICO Score

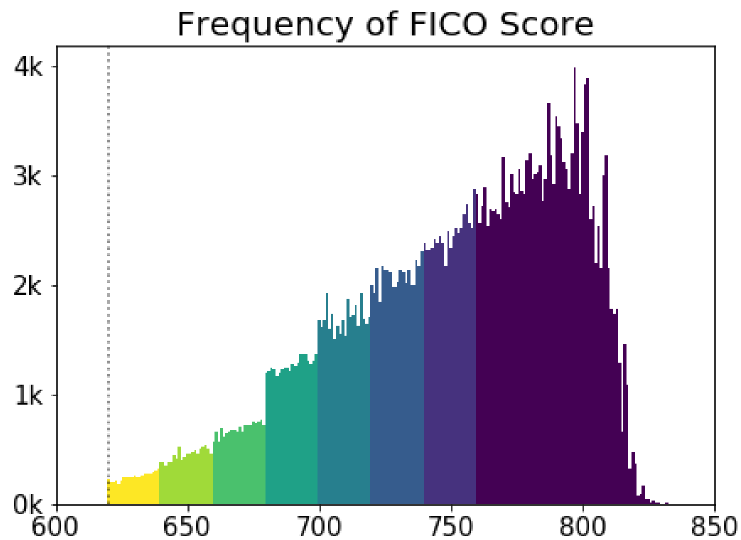

(a)

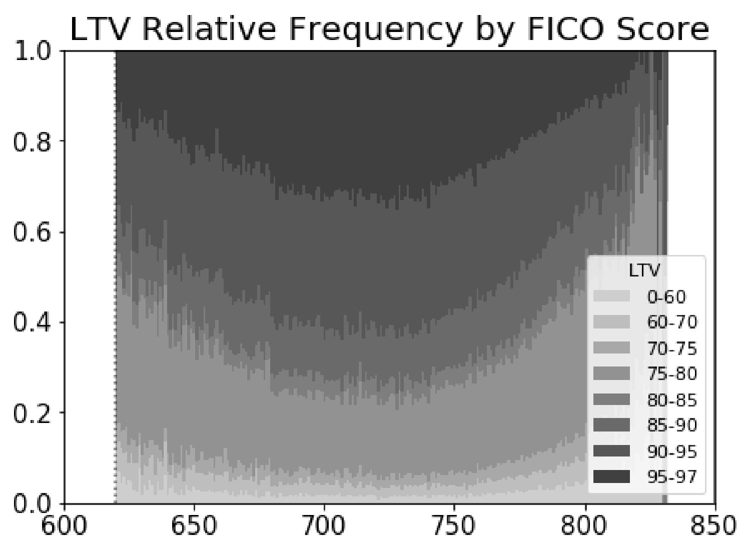

(b)

Figure (A.1a) depicts the frequency of borrowers with different FICO scores at origination in the HMDA Sample. Different colors depict the bins in the FNMA LLPA pricing grids. There are discontinuities at important credit score thresholds, potentially because of pricing benefits of being at a higher credit score, which increases the quantity of loans demanded. Figure (A.1b) depicts the relative frequency of different LTV loans within borrowers of a given credit score. Borrowers who receive loans make substantial down-payments when they have low credit scores, this tapers among borrowers of middling credit scores, and then increases at higher credit scores.

Figure A.2: HMDA-FNMA-SEDA Merge Quality

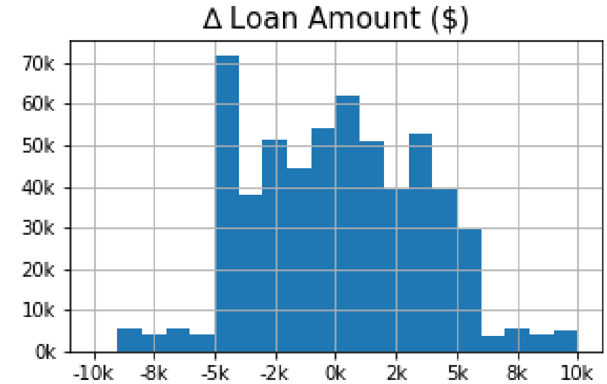

(a)

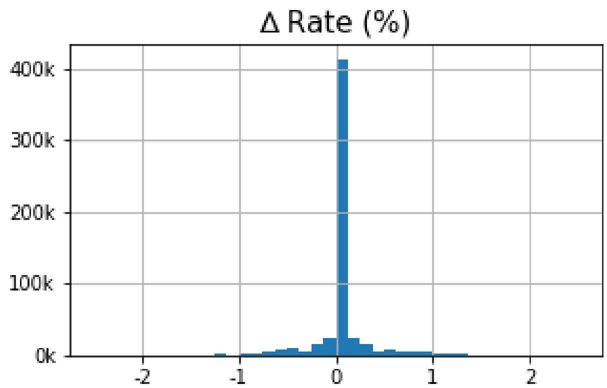

(b)

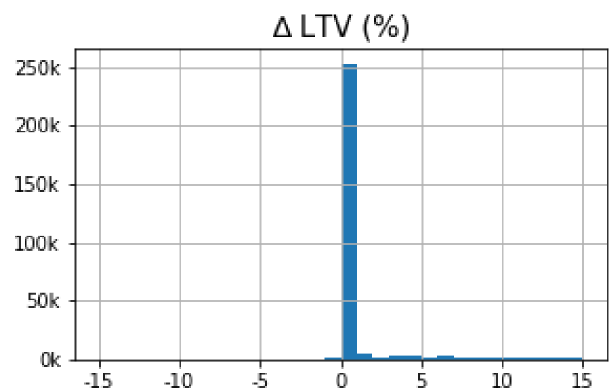

(c)

The plot above depicts the accuracy of the HMDA-FNMA-SEDA merge along the variables used for the fuzzy merge. The difference plotted denotes the Fannie Mae value minus the HMDA value. Figure (A.2a) plots the difference in loan amounts, Figure (A.2b) in interest rates, and Figure (A.2c) in loan-to-value ratios. The noise in the loan amount merge is expected because loan amounts in HMDA are redacted to the nearest $10k. Overall, merged loans appear fairly accurate.

Figure A.3: Joint pricing regimes (2011-2020)

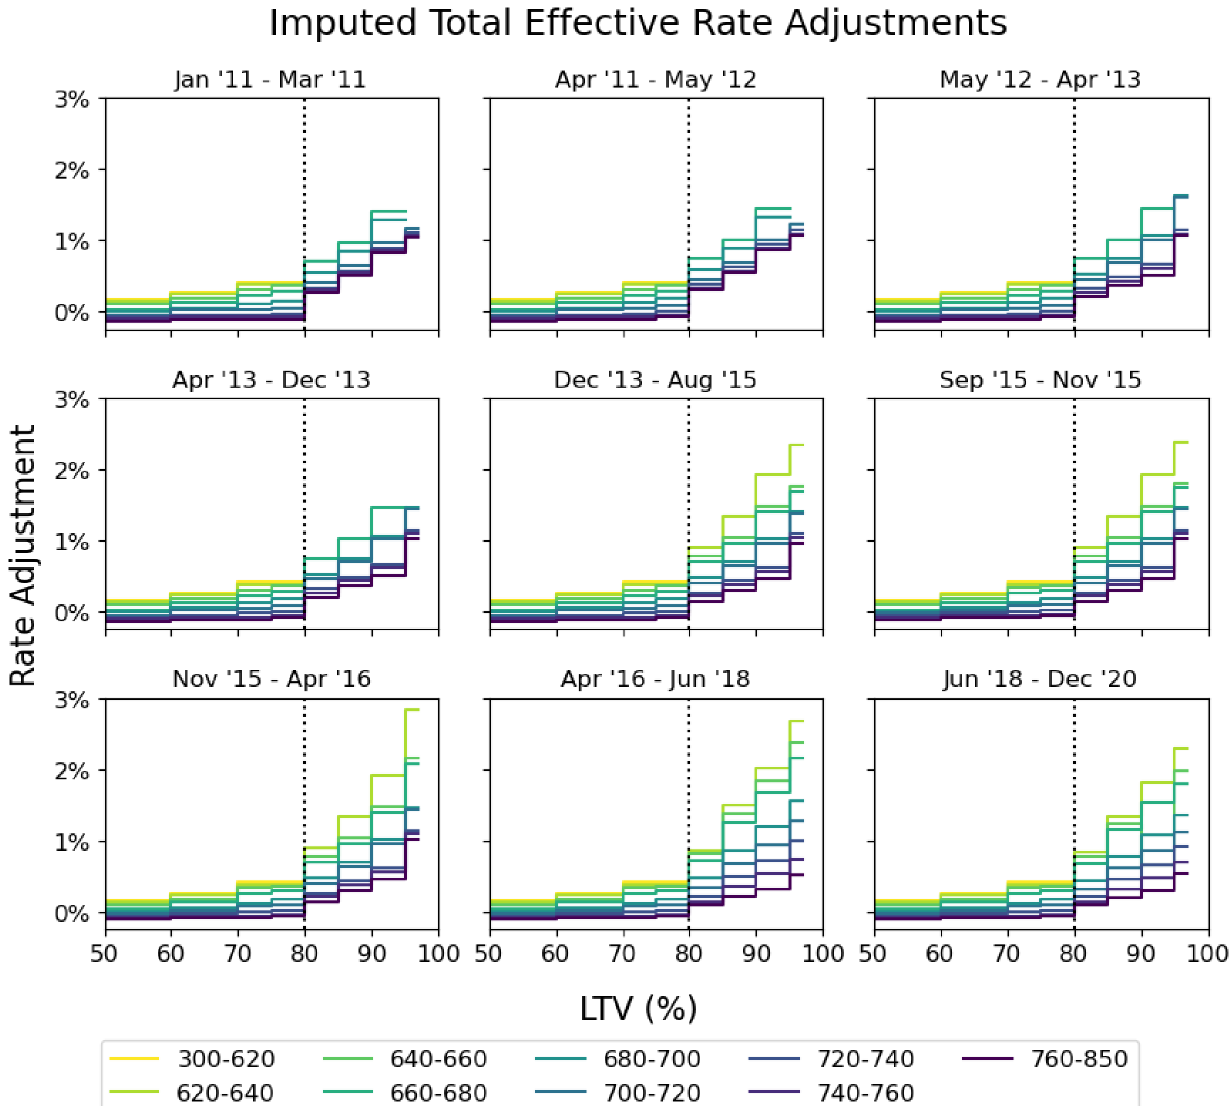

The total effective interest rate adjustment, computed as the total rate hike coming from both LLPA pass-through and PMI costs. These panels depict nine different regimes covered by the CRISM Sample.

APPENDIX B
SUPPLEMENTAL MATERIAL TO CHAPTER 3

B.1 Derivations

B.1.1 Bunching Behavior under Credit Constraints

Consider a household choosing consumption today X_0, consumption tomorrow X_1, housing assets H, a housing down-payment D, and savings A_1. The household faces a unit price of housing p_H and a liquid savings rate r_f. The household has degenerate beliefs, $\mathbb{E}_i[\tilde{r}_H]$, about the the realization of net-of-depreciation housing capital gains rate $\tilde{r}_H$. The household makes its purchases with income at time-0 and time-1, Y_0 and Y_1, and capital gains from the sale of the housing asset. It is subject to the standard budget constraints and liquid borrowing constraints.

The household is assumed to purchase housing assets outright so that the decision problem focuses on the intensive margin of housing demand. It may finance the housing purchase using a mortgage loan with an interest rate equal to the savings rate r_f. This mortgage loan must satisfy a standard loan-to-value limit $\overline{L}$ and an initial balance limit $\overline{B^o}$ which may be considered to be derived from an institutional debt-to-income constraint.

The household's problem may be solved in two steps for clarity. For a given choice of housing H and a down-payment D, it chooses the best possible allocation of time-0 and time-1 non-housing consumption, X_0 and X_1. It then chooses the best possible combination of housing H and down-payment D subject to borrowing constraints, $\overline{L}$ and $\overline{B^o}$.

Formally, we write:

$$\max_{\{H_0, D\}} \quad \Phi(H_0, D)$$

$$\text{s.t.} \qquad p_H H_0 - D \leq \overline{B^o} \qquad (\zeta)$$

$$1 - \frac{D}{p_H H_0} \leq \overline{L} \qquad (\xi)$$

Where:

$$\Phi(H_0, D) = \max_{\{X_0, X_1, A_1, H_1\}} U(X_0, H_0) + \beta \mathbb{E}\big[U(X_1, H_1) \mid \tilde{r}_H = \mathbb{E}_i[\tilde{r}_H]\big]$$

$$\text{s.t.} \qquad Y_0 = X_0 + \frac{A_1}{1 + r_f} + D \qquad (\lambda_0)$$

$$Y_1 + A_1 + p_H H_0(1 + \tilde{r}_H) = X_1 + (p_H H_0 - D)(1 + r_f) + r_H H_1 \qquad (\lambda_1)$$

$$A_1 \geq 0 \qquad (\mu_1)$$

Note that the constraints indexed by ζ and ξ may be written in (H, D)-space. ζ corresponds to Equation (3.4), which describes an initial balance limit, and ξ corresponds to Equation (3.5), which describes a loan-to-value limit. These form the boundaries of the kinked opportunity set depicted in Figure (B.1).

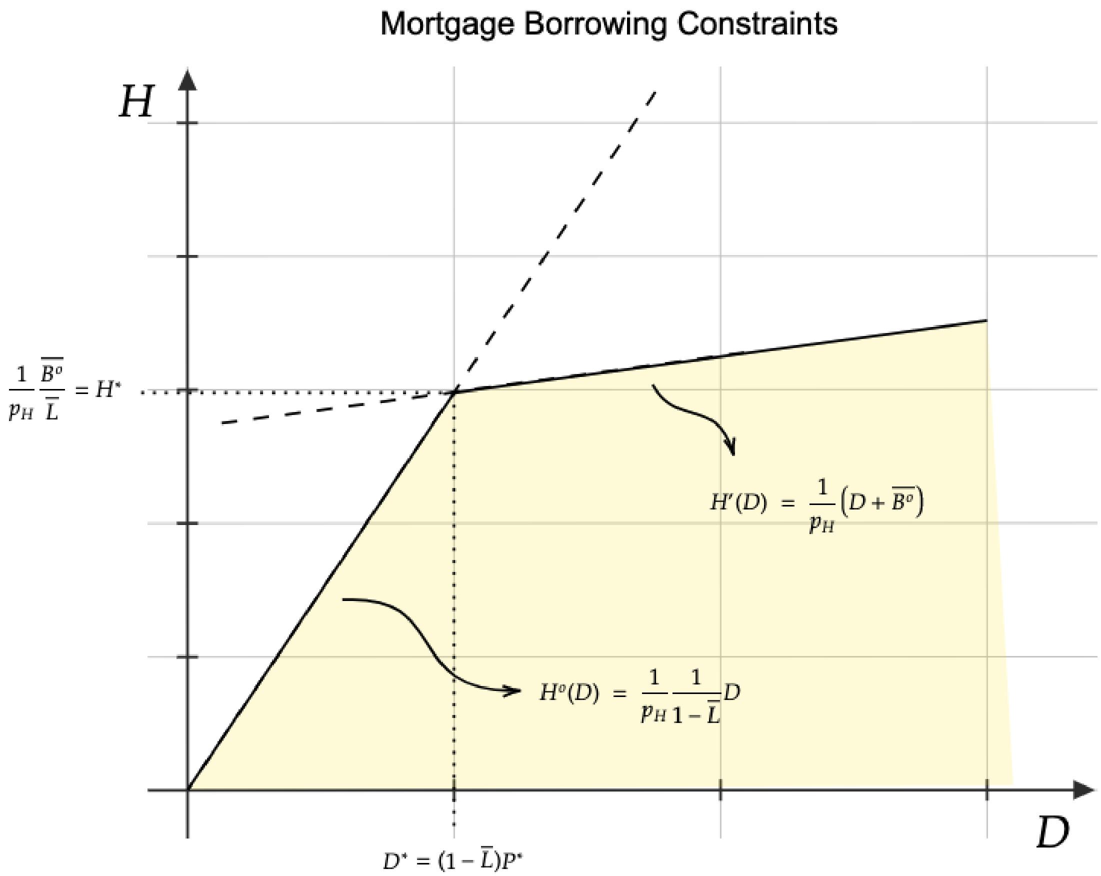

Figure B.1: This figure depicts the household's choice set of housing quantity and down-payment amount at the time of home purchase. The feasible set is restricted and, in particular, kinked, because of the two borrowing limits faced by the household, the LTV limit, and the DTI limit that functions like an initial balance limit.

To identify the household indifference curves in (H, D)-space, plug the budget constraints into the objective function and use the degenerate beliefs to simplify the expectations operator:

$$\Phi(H_0, D) = \max_{\{A_1, H_1\}} U\Big(Y_0 - \frac{A_1}{1+r_f} - D, H_0\Big) + \beta U\Big(Y_1 + A_1 + p_H H_0 \mathbb{E}_i[\tilde{r}_H - r_f] + D(1+r_f) - r_H H_1, H_1\Big)$$

$$\text{s.t.} \qquad A_1 \geq 0 \qquad (\mu_1)$$

Use the first-order condition for H_1 to obtain:

$$\frac{U_H^1}{U_X^1} = r_H \tag{B.1}$$

Use the first-order condition for A_1 to solve for the unconstrained optimal savings $\hat{A}_1$. The constrained optimal savings is given by $A_1^* = \max\{\hat{A}_1, 0\}$. For simplicity, I assume separability in U and that $\beta(1 + r_f) = 1$. Then we have:

$$\hat{A}_1 = \frac{1+r}{2+r}\Big[Y_0 - Y_1 - p_H H \mathbb{E}_i[\tilde{r}_H - r_f] + r_H H_1^*\Big] - D(1+r) \tag{B.2}$$

Note that Equation (B.2) with $\hat{A}_1 = 0$ defines a line in (H, D)-space that divides cases in which the household is constrained from those in which it is not. This line is depicted in Figure (B.2), dividing red from blue.

Plug the closed-form representation of A_1^* back into the objective function to obtain a closed-form representation of Φ. For clarity, define the household's wealth given its housing investment, $W(H_0)$, as its income and excess capital gains from housing discounted to time 0. Formally, $W(H_0) = Y_0 + \frac{1}{1+r}Y_1 + \frac{1}{1+r}p_H H_0(\mathbb{E}_i[\tilde{r}_H] - r_f) - \frac{1}{2+r}r_H H_1^*$. Then we have:

$$\Phi(H_0, D) = \begin{cases} U\Big(\frac{1+r}{2+r}W(H_0), H_0\Big) & \\ \qquad +\beta U\Big(\frac{1+r}{2+r}W(H_0), H_1^*(H_0)\Big) & \hat{A}_1 \geq 0 \\ U\Big(Y_0 - D, H_0\Big) & \\ \qquad +\beta U\Big(Y_1 + p_H H_0 \mathbb{E}_i[\tilde{r}_H - r_f] + D(1+r) - r_H H_1^*, H_1^*\Big) & \hat{A}_1 \leq 0 \end{cases}$$

Using this closed-form representation, it is possible to derive some intuitive characteristics of indifference curves in (H, D)-space. I show below that these indifference curves are convolutions of consumer preferences and (non-mortgage) borrowing constraints. I depict these indifference curves in Figure (B.2).

In the case when the household is not constrained, $\hat{A}_1 \geq 0$, the down-payment, D, does not enter its utility. Intuitively, in the unconstrained case, it is possible for the household to borrow or save around a down-payment. As a result, the borrower has horizontal indifference curves in (H, D)-space. This is represented in blue in Figure (B.2). Horizontal indifference curves will not produce bunching and so any bunching in the data constitutes evidence of financial frictions.

In the case when the borrower is constrained, making additional down-payment effectively sacrifices consumption today for consumption tomorrow, which is costly in utility terms because a high marginal utility of consumption today is what drives to consumer against her borrowing constraint. We can see this by setting $\Phi = \overline{\Phi}$ in the constrained case. Consider H_0 an implicitly defined function of D, differentiate, and solve for $H_0'(D)$. For clarity, substitute the values of consumption, $\hat{X}_0$, $\hat{X}_1$, implied by the choice of down-payment and housing. Use the results from the first-order condition in H_1, Equation (B.1), to cancel terms in $H_1^*(D)$.

$$H_0'(D) = \frac{U_X^0 - U_X^1}{U_H^0 + \beta U_X^1 p_H \mathbb{E}_i[\tilde{r}_H - r_f]}$$

To help with interpretation, consider the case in which HH do not believe that housing will have any excess capital gains relative to liquid assets. And note that the difference in marginal utility of consumption at time 0 and 1 is the Lagrange multiplier on the liquid savings constraint.

$$H_0'(D) \overset{\hat{r}=r_f}{=} \frac{U_X^0 - U_X^1}{U_H^0} = \frac{\mu_1}{U_H^0}$$

These indifference curves are represented in red in Figure (B.2).

Intuitively, additional down-payment at time 0 incurs a utility cost that is the difference between marginal utility of non-housing consumption in time-0 and time-1. The more money put into the down-payment, the more recovered from the sale of the home. Additional time-0 housing has utility value due to the housing services enjoyed at time-0 as well as potential investment value that can be enjoyed as consumption at time-1.

The numerator of this expression is positive because households run up against their liquid borrowing constraints because they are having trouble moving consumption to the present. More technically, the Kuhn-Tucker conditions require that $\mu_1 > 0$ when $A_1^* = 0$. The denominator is positive provided that housing is not expected to perform poorly because the marginal utility of additional housing services is positive. During the run-up to the crisis, housing was generally expected to do well.

Finally, it is worth noting that extreme housing optimism flattens the indifference curve.

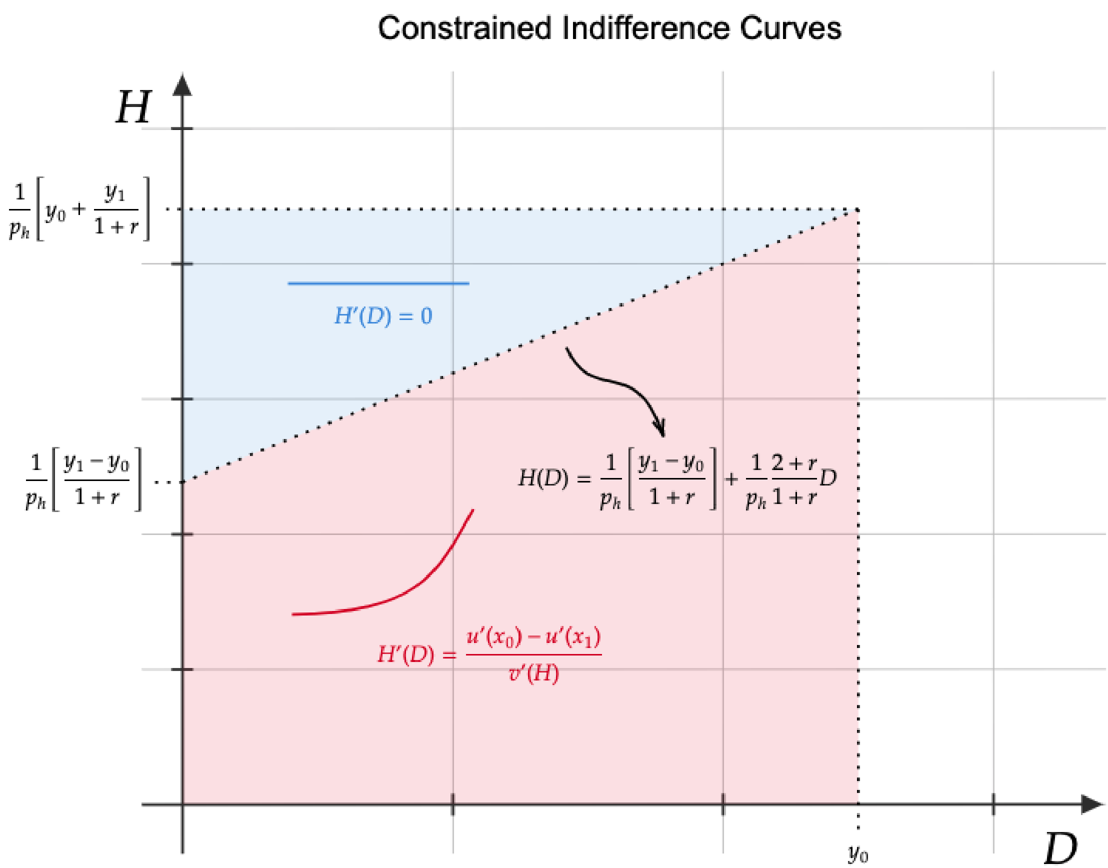

Figure B.2: This figure depicts the space of housing and down-payments from which households choose at the time of home purchase. The housing demand is limited by the household's wealth, $y_0 + \frac{y_1}{1+r}$. For sufficiently high housing demand and low down-payment (the blue region), the borrower is unconstrained and the borrower is indifferent to increasing the amount of the down-payment. Being unconstrained, they can finance around the mortgage. For low housing and a high down-payment (the red region), the borrower is constrained, and requires an increase in the quantity of housing to compensate for an increase in the required down-payment.

B.1.2 The Bunching Estimator and Curvature of the Indifference Curve

Here, I argue that the bunching estimator, which recovers an empirical measure of the downpayment adjustment of the marginal buncher, Δd, can be used to understand the curvature of the marginal buncher's (indirect) indifference curve, $\tilde{h}$.

Consider a kinked opportunity set in (d, h)-space. Denote the slope of the lower and upper portion of the kinked opportunity set as μ_0 and μ_1, respectively. Denote the downpayment and housing quantity at the kink point be denoted by $\hat{d}$ and $\hat{h}$, respectively.

The marginal buncher, m, is then defined as optimizing relative to the upper portion of the kinked opportunity set at $(\hat{d}, \hat{h})$ and optimizing relative to the lower portion of the kinked opportunity set at $(\hat{d} + \Delta d, \hat{h} + \Delta h)$. Suppose the kink is small, so that the buncher remains on the same indifference curve; we denote this indifference curve as $\tilde{h} : D \to H$. Then we have:

$$\tilde{h}(\hat{d}) = \hat{h} \tag{B.3}$$

$$\tilde{h}'(\hat{d}) = \mu_1 \tag{B.4}$$

$$\tilde{h}(\hat{d} + \Delta d) = \hat{h} + \Delta h \tag{B.5}$$

$$\tilde{h}'(\hat{d} + \Delta d) = \mu_0 \tag{B.6}$$

Begin with an approximation of $\tilde{h}''(\hat{d})$:

$$\tilde{h}''(\hat{d}) = \frac{\tilde{h}'(\hat{d} + \Delta d) - \tilde{h}'(\hat{d})}{\Delta d}$$

Multiply both sides by $\frac{\hat{d}}{\tilde{h}'(\hat{d})}$ and rearrange:

$$\frac{\frac{\tilde{h}'(\hat{d}+\Delta d) - \tilde{h}'(\hat{d})}{h'(\hat{d})}}{\frac{\Delta d}{\hat{d}}} = \frac{\hat{d}\tilde{h}''(\hat{d})}{\tilde{h}'(\hat{d})}$$

We can use (B.4) and (B.6) in the above and rewrite terms for clarity to obtain:

$$\frac{\%\Delta\mu}{\%\Delta d} = \left.\frac{d\ln\tilde{h}'(d)}{d\ln d}\right|_{d=\hat{d}} = \varepsilon_d^{\tilde{h}'}$$

Note that administrative parameters governing borrowing constraints through the FHA can be used to recover $\%\Delta\mu$ and $\hat{d}$ and the bunching estimator can be used to recover Δd, so it is possible to estimate $\varepsilon_d^{\tilde{h}'}$.

B.1.3 Compensated and Uncompensated Elasticity of Housing Demand

This proof follows the observation in Saez [2010], the formalization of this observation and sketched proof in Kleven [2016].

Consider the introduction of a kink into the loan-offer curve facing a borrower. Consider any borrower, i, who responds to the introduction of the kink and whose indifference curves in (h, d)-space are tangent to the the loan-offer curve both before and after the introduction of the kink. Note that the marginal borrower, whose behavioral change the bunching technique is designed to measure, is one such borrower.

Define borrower i's observed elasticity as $\hat{e}_i \equiv \frac{\%\Delta D_i}{\%\Delta\mu}$. Note that the observed elasticity can be decomposed into the joint effect of two elasticities, a compensated LTV elasticity of down-payment supply, e_i^c, and an (income-like) initial-loan elasticity of down-payment supply, η_i. We write:

$$\hat{e}_i = e_i^c + \frac{\%\Delta\tilde{B}}{\%\Delta\mu}\eta_i$$

Use the Slutsky-decomposition, $e_i^c = e_i^u - s_i^D\eta_i$, to rewrite the expression. Note that the down-payment share of initial-loan balance can be expressed as, $s_i^D = \frac{D_i^o\mu^o}{\tilde{B}^o}$. Simplify terms.

$$= e_i^c + \frac{\Delta\tilde{B}}{\tilde{B}^o}\frac{\mu^o}{\Delta\mu}\frac{\tilde{B}^o}{D_i^o\mu^o}(e_i^u - e_i^c) = e_i^c + \frac{\Delta\tilde{B}/D_i^o}{\Delta\mu}(e_i^u - e_i^c)$$

Note that $\Delta\tilde{B} = \Delta\mu D_i^* - \Delta\mu D_i^o$ where the first term is due to an effective increase in the initial-loan balance and the second term is the effective fall in the initial-loan balance due to the down-payment funds supplied. Substitute in this expression and simplify:

$$= e_i^c + \frac{D_i^* - D_i^o}{D_i^o}(e_i^u - e_i^c)$$

Finally, note that we can write the change in the average (cf. marginal) slope as: $\Delta a = \frac{D_i^* \mu^o + (z - z^*)\mu'}{D^o} - \mu^o = \frac{D_i^* - D_i^o}{D_i^o}\Delta\mu$. Substitute in this expression and rearrange terms:

$$= e_i^c + \frac{\Delta a}{\Delta \mu}(e_i^u - e_i^c) = \left[1 - \frac{\Delta a}{\Delta \mu}\right] * e_i^c + \left[\frac{\Delta a}{\Delta \mu}\right] * e_i^u$$

Note that when the kink is small and the change in average slope for the marginal buncher is near 0, this simplifies to the compensated elasticity, e_i^c. What is most obviously of interest in this setting, however, is the uncompensated elasticity of housing demand, e_i^u. In the run-up to the crisis, the loan-to-value on the entire loan increased, not just on the portion of the loan above some threshold.

This estimation is helped in a two regards. First, we exploit a fairly large-sized kink, which, by the above proof, will capture a weighted average of compensated and uncompensated housing demand elasticities, so that to the extent that these values differ, we are at least capturing some of the uncompensated demand. Second, in this setting, it is reasonable to suppose that, although housing demand may respond dollar-for-dollar, down-payment supply does not respond to an increase in the initial available initial loan balance. In this case the compensated and uncompensated elasticities estimated here are close. It is worth noting, finally, that when estimating structural parameters, the bunching literature generally assumes functional forms for which compensated and uncompensated elasticities of earnings supply are identical.

B.1.4 The Bunching Estimator and Loan-to-Value Elasticity of Housing Demand

The curvature of the indirect indifference curves is responsible for generating bunching. The most natural estimator in the standard bunching framework, as described in the section above, is $\frac{\%\Delta d}{\%\Delta \mu}$, or the value-to-downpayment elasticity of the down-payment. A more natural object of interest for policy purposes is understanding how housing demand responds to the loan-to-value ratio imposed by the lending program. I show how to transform the traditional bunching estimator into one better suited for policy analysis in this application. I proceed somewhat informally, in the hopes of conveying intuition:

We note the fact that $h = \frac{1}{p_h}\frac{d}{1-L_0}$ both at $(\hat{d}, \hat{h})$ and at $(\hat{d} + \Delta d, \hat{h} + \Delta h)$ and totally differentiate to derive:

$$\frac{dh}{h} = \frac{\frac{dd}{1-L_0}}{\frac{d}{1-L_0}} = \frac{dd}{d}$$

We also note the fact that $\mu = \frac{1}{1-L}$ where L is the loan-to-value ratio of the marginal dollar of housing purchased and totally differentiate to derive:

$$\frac{d\mu}{\mu} = \frac{\frac{1}{(1-L)^2}dL}{\frac{1}{1-L}} = \frac{dL}{L}\frac{L}{1-L}$$

Plugging these expressions into our elasticity of interest, we obtain:

$$\frac{d\ln h}{d\ln L} = \frac{\frac{dh}{h}}{\frac{dL}{L}} = \frac{\frac{dd}{d}}{\frac{1-L}{L}\frac{d\mu}{\mu}} = \frac{L_0}{1-L_0}\frac{d\ln d}{d\ln \mu} = \frac{L}{1-L}\varepsilon_d^{\tilde{h}'}$$

B.1.5 Analogy to Standard Estimators

Kinked budget constraints are a common feature of household choice sets; they are generated by government transfer programs, e.g. income tax schedules, and, in turn, generate bunching behavioral responses from optimizing agents. Bunching estimators commonly exploit this behavior to estimate real elasticities of fundamental interest, such as the tax elasticity of labor supply. Kinked borrowing constraints are a common feature of government lending programs, e.g. those that restrict both the LTV and loan amount, and, in principle, may generate bunching behavior as well. This behavior may be exploited to estimate financial elasticities of interest, like the elasticity of housing demand with respect to loan-to-value requirements.

APPENDIX C
SUPPLEMENTAL MATERIAL TO CHAPTER 4

C.1 Robustness

C.1.1 Retirement and New Mortgage Debt

Table (C.1) estimates the effect of retirement on new mortgage debt instrumenting with SS eligibility using the SCF data. The age bounds are larger, 50 to 85, but it includes a flexible age control to account for the sampling bias that creates a hump-shaped distribution of new mortgage debt. The results are robust to this alternative specification. Table (C.2) repeats the exercise with equity extraction rates specifically.

Table (C.3) estimates the effect of retirement on various kinds of new mortgage debt instrumenting with SS eligibility using the SCF data. The age bounds are even tighter, 59 to 71. The results for all mortgage debt and equity extractions remain robust to this very conservative specification.

Table C.1: Predictable New Mortgage Debt

	Dependent variable:					
	1{M. Debt}	1{Retire}	1{M. Debt}	1{M. Debt}	1{Retire}	1{M. Debt}
	OLS	*1st Stage*	*IV*	*OLS*	*1st Stage*	*IV*
	(1)	(2)	(3)	(4)	(5)	(6)
1{Retire}	0.012***		0.089***	0.011***		0.041*
	(0.002)		(0.023)	(0.002)		(0.022)
1{SS Threshold}		0.048***			0.038***	
		(0.003)			(0.003)	
1	0.006***	0.020***	0.004***	0.009***	0.031***	0.008***
	(0.001)	(0.001)	(0.0005)	(0.001)	(0.001)	(0.001)
Age				x	x	x
Age^2				x	x	x
Observations	315,547	315,547	315,547	315,547	315,547	315,547
R^2	0.001	0.008	−0.022	0.002	0.014	−0.002
Adjusted R^2	0.001	0.008	−0.022	0.002	0.014	−0.002
F Statistic	177.4***	2,545.0***		175.6***	1,476.3***	

Note: *p<0.1; **p<0.05; ***p<0.01
SEs computed with SCF bootstrap weights, adjusted for SCF multiple implicates, and clustered by person and year.

Table C.2: Predictable Equity Extraction

	Dependent variable:					
	1{Extract}	1{Retire}	1{Extract}	1{Extract}	1{Retire}	1{Extract}
	OLS	*1st Stage*	*IV*	*OLS*	*1st Stage*	*IV*
	(1)	(2)	(3)	(4)	(5)	(6)
1{Retire}	0.002**		0.030***	0.001		0.022*
	(0.001)		(0.010)	(0.001)		(0.012)
1{SS Threshold}		0.048***			0.038***	
		(0.003)			(0.003)	
1	0.001***	0.020***	0.001***	0.002***	0.031***	0.001***
	(0.0002)	(0.001)	(0.0002)	(0.0003)	(0.001)	(0.0005)
Age				x	x	x
Age^2				x	x	x
Observations	315,547	315,547	315,547	315,547	315,547	315,547
R^2	0.00004	0.008	−0.013	0.0003	0.014	−0.006
Adjusted R^2	0.00004	0.008	−0.013	0.0003	0.014	−0.006
F Statistic	14.2***	2,545.0***		34.6***	1,476.3***	

Note: *p<0.1; **p<0.05; ***p<0.01
SEs computed with SCF bootstrap weights, adjusted for SCF multiple implicates, and clustered by person and year.

Table C.3: Predictable New Mortgage Debt

	Dependent variable:										
	Ret1nd	M1Debt1nd		M1New1nd		M1Ext1nd		M1Refi1nd		M2Debt1nd	
	Stage 1	*OLS*	*IV*	*OLS*	*IV*	*OLS*	*IV*	*OLS*	*IV*	*OLS*	*IV*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Ret1nd		0.012^{***}	0.070^{**}	0.006^{***}	0.007	0.001	0.025^{*}	0.001	−0.001	−0.0004	−0.002
		(0.003)	(0.035)	(0.002)	(0.016)	(0.001)	(0.014)	(0.001)	(0.014)	(0.0003)	(0.005)
SSThresh1nd	0.031^{***}										
	(0.003)										
Constant	0.037^{***}	0.018^{***}	0.016^{***}	0.003^{***}	0.003^{***}	0.002^{***}	0.001^{*}	0.004^{***}	0.004^{***}	0.001^{***}	0.001^{**}
	(0.001)	(0.002)	(0.002)	(0.0003)	(0.001)	(0.0003)	(0.001)	(0.001)	(0.001)	(0.0001)	(0.0003)
Observations	124,596	124,596	124,596	124,544	124,544	124,596	124,596	124,596	124,596	124,596	124,596
R^2	0.004	0.0003	−0.008	0.001	−0.0002	0.00003	−0.010	0.00002	−0.0002	0.00001	−0.00000
Adjusted R^2	0.004	0.0003	−0.008	0.001	−0.0002	0.00002	−0.010	0.00001	−0.0002	−0.00000	−0.00001
Residual Std. Error	32.067	21.334	21.421	8.424	8.427	7.328	7.364	9.900	9.901	3.743	3.743
F Statistic	509.378^{***}	35.295^{***}		69.215^{***}		3.714^{*}		2.116		0.894	

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

SEs computed with SCF bootstrap weights, adjusted for SCF multiple implicates, and clustered by person and year.

C.1.2 New Mortgage Debt and Liquid Balances

Table (C.4) estimates the effect of various housing wealth management decisions on changes in liquid balances instrumenting with SS eligibility using the HRS data. Instead of winsorizing at 5%, this specification winsorizes at 1%. The magnitudes and significances of the estimates increases.

Table (C.5) presents summary statistics for the sample of HHs from HRS "refined sample" that are observable in the waves before and after they claim SS benefits. This subsample is used in a correlational study of changes in HH liquid balances and different housing wealth management actions. I run an OLS specification with FEs for demographics and a running variable of pre-period liquid balances and report the results in Table (C.6). On average, in the period the HHs claim SS benefits, equity extractions and second mortgages are associated with liquid balance reductions. This motivates the IV approach.

Table C.4: Housing Wealth Management and Liquid Balances at Retirement (HRS)

	Dependent variable:								
	ΔSav_{it}	1{Sell}	ΔSav_{it}	ΔSav_{it}	1{Extract}	ΔSav_{it}	ΔSav_{it}	1{New 2nd}	ΔSav_{it}
	OLS	*1st Stage*	*IV*	*OLS*	*1st Stage*	*IV*	*OLS*	*1st Stage*	*IV*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
1{Sell}	9.6k***		289k***						
	(1.1k)		(96.8k)						
1{Extract}				−153		107k***			
				(386)		(32.1k)			
1{New 2nd}							−1.4k**		695k
							(612)		(604k)
1{SS Threshold}		0.005***			0.012***			0.002	
		(0.001)			(0.002)			(0.002)	
1	−1.3k	−0.004	−19	−1.5k	−0.031	2.0k	−1.3k	−0.007	3.8k
	(1.3k)	(0.004)	(1.8k)	(1.3k)	(0.019)	(3.4k)	(1.3k)	(0.006)	(6.6k)
Demo FEs	x	x	x	x	x	x	x	x	x
Clustered SEs	x	x	x	x	x	x	x	x	x
Observations	102,868	101,760	101,760	101,474	100,373	100,373	102,280	101,179	101,179
R^2	0.003	0.006	−1.338	0.001	0.035	−0.614	0.001	0.013	−12.308
Adjusted R^2	0.001	0.005	−1.341	−0.0002	0.034	−0.616	−0.0002	0.012	−12.326
F Statistic	2.044***	4.404***		0.814	26.847***		0.862	10.026***	

Note: *p<0.1; **p<0.05; ***p<0.01
Data winsorized by ΔSav_{it} at 1pp level.

Table C.5: Summary Statistics for HRS "Refined" Sample

	N	$\overline{x}$	s_x	p25	p50	p75
Interview Characteristics:						
Wave	6.5k	6.9	3.1	4	7	9
Interview Date	6.5k	2004	6.3	1999	2004	2009
Interviewee Characteristics:						
Birth Date	6.5k	1941	7	1935	1940	1946
1{Hispanic}	6.5k	0.01	0.30	0	0	0
1{Black}	6.5k	0.177	0.38	0	0	0
# Yrs. Educ.	6.5k	12.70	3.06	12	12	15
1{Married}	6.5k	0.59	0.49	0	1	1
# Children	6.5k	2.835	1.96	2	3	4
1{Home-owner}	5.3k	0.79	0.41	1	1	1
SS Claim Date	6.5k	2003	6.2	1998	2002	2008
Savings at Retirement:						
Δ Savings$_t$ ($k)	6.5k	5.66	100.6	−2.0	0.0	5.0
Savings$_t$ ($k)	6.5k	23.8	102.5	.40	4.0	19.0
Savings$_{t-1}$ ($k)	6.5k	18.2	51.1	.20	4.0	15.0
Refinancing Activity:						
1{Buy or Sell}	4.5k	0.092	0.29	0	0	0
1{Extract Equ.}	6.4k	0.092	0.29	0	0	0
1{New 2nd Mort.}	6.5k	0.042	0.20	0	0	0
1{New HELOC}	4.7k	0.072	0.26	0	0	0

Table C.6: Liquid Balances and Housing Wealth Management at SS Claim Date

	Dependent variable:					
	$\Delta Savings_{it}$					
	(1)	(2)	(3)	(4)	(5)	(6)
1{Buy}	5,431					
	(9,139)					
1{Sell}		20,248**				
		(9,385)				
1{Buy and Sell}			6,391			
			(7,172)			
1{Equ. Extract}				−3,038		
				(4,380)		
1{New 2nd}					−10,338*	
					(6,195)	
1{New HELOC}						−6,477
						(6,582)
1	12,785	13,027	13,074	12,268	13,447	−12,012
	(73,862)	(73,837)	(73,859)	(74,486)	(74,033)	(122,583)
Demo FEs	x	x	x	x	x	x
Observations	6,446	6,446	6,446	6,335	6,413	4,570
R^2	0	0	0	0	0	0
Adjusted R^2	0	0	0	0	0	0
Residual Std. Error	98,608	98,575	98,605	99,415	98,832	113,559
F Statistic	5***	5***	5***	5***	5***	3***

Note: *p<0.1; **p<0.05; ***p<0.01