An Empirical Analysis of Government-Sponsored Enterprise Policy

Joseph Patrick Hogan

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2015

ABSTRACT

An Empirical Analysis of Government-Sponsored Enterprise Policy

Joseph Patrick Hogan

Credit markets are central to our understanding of business cycles and macroeconomic policy, and in their macroeconomic importance there is considerable evidence that credit markets are unlike markets for other products. By way of illustration, the documentation for the Federal Reserve Board's baseline macroeconomic model, FRB/US, mentions "Debt" 29 times and "Interest" 82 times, while "Auto" is mentioned only twice, both times referring to "Auto Loan", a type of credit.[44] Nonetheless, credit markets are still markets, and in many ways are particularly amenable to microeconomic analysis. My dissertation focuses on a micro-level empirical analysis of a credit market with particular macroeconomic import and policy relevance, namely, the market for government-guaranteed mortgages.

During the 2000s U.S. mortgage borrowing experienced its most volatile cycle in the postwar record, with mortgage debt more than doubling between 2000 and 2008 before declining by more than 10% over the following five years. The consequences of the boom and bust for both borrowers and the wider macroeconomy were significant, with millions losing their homes to foreclosure or their jobs to the ensuing deleveraging-driven recession.[159,160] Recent research has focused on variations in credit supply as a primary determinant of both the boom in mortgage borrowing[158] and subsequent collapse,[102] as well as the concurrent rise and fall of residential real estate prices[4,78] and employment.[166] In the wake of the Great Recession many have called for countercyclical policy intervention in the mortgage market, both to restrain over-leveraging during booms and to provide additional access to refinancing credit during busts.[122,177] Moreover some analysis has placed the blame for the volatile U.S. credit cycle on the policies of Fannie Mae and Freddie Mac, the two largest government-sponsored enterprises, which have been labeled as excessively risky,[3] actively destabilizing,[110,180] and regressive.[75,94,124] Nevertheless,

though many have called for their reform[76,108] these two agencies appear to be a continuing feature of the U.S. housing finance system[90] and are currently well-positioned to implement countercyclical credit supply policies. In my dissertation I propose a novel countercyclical policy intervention by the government-sponsored enterprises and analyze its impact on mortgage borrowers.

Chapter 1, "A Descriptive Analysis of the U.S. Mortgage Market and the Government-Sponsored Enterprises", sets the stage for my subsequent analysis. I first provide a description of the institutional arrangement of the U.S. mortgage market, with a particular focus on the goals and policies of the government-sponsored enterprises. In so doing I describe the key friction embedded in this market that necessitates policy intervention: that some fixed-rate mortgage borrowers cannot refinance into lower interest rates, hindering monetary transmission to consumption.[43] I next review the relevant literature related to borrowers and mortgage credit demand, lenders and mortgage credit supply, and the effects of realized and proposed mortgage market policies. I then lay out a model of the mortgage refinancing process to help contextualize my results. The model generates two key predictions regarding the effects of guarantee fees on credit supply and the effect of liquidity preference on both applications for refinancing credit and the realized volume of refinancing. Finally, I describe the data sources used in my analysis and provide a brief outline of my empirical strategy for analyzing the mortgage market.

In Chapter 2, "Guarantee Fee Increases and Mortgage Credit Supply", I study the effect of Fannie Mae and Freddie Mac's mortgage default insurance premiums, known as guarantee fees, on the supply of mortgage credit. I show that higher guarantee fees, which function much like a tax on mortgage origination, reduce the incentive for lenders to provide credit. In the short-term, lenders absorb some of the fee increase through lower revenues, but in the long-term there is essentially complete pass-through from guarantee fees to interest rates. The extent of pass-through to interest rates is lower in more concentrated markets and higher for riskier borrowers and for nonbank mortgage lenders, suggesting a strategic response on the part of lenders. As a result of the increase in guarantee fees, the probability of refinancing declines for borrowers with government-guaranteed mortgages, which I interpret as evidence of tightened credit constraints. My estimates suggest an elasticity of refinancing with respect to interest rates of between 1.7% and 2.8%, very much in line with prior estimates, and imply that the combined effect of a pair of large fee increases implemented in 2012 was to reduce the total volume of refinancing in the U.S. by roughly $205 billion annually. I also find that default rates for borrowers with agency mortgages increased slightly following the guarantee fee increases by between 1.9% and 4.2% of the estimated change in refinancing rates, also in line with prior estimates, suggesting that tightening credit conditions may have contributed to higher default rates. I discuss the macroeconomic implications of these findings and connect the results to the policy debate regarding the future of the U.S. housing finance system.

Chapter 3, "Unemployment and the Value of Refinancing Credit", pivots towards the demand side of the mortgage market. In this chapter I provide evidence that unemployed borrowers would like to refinance, but are unable to do so because of credit constraints. These constraints generate a crucial endogeneity problem in that unemployment affects both demand for and access to credit, making it difficult to estimate preferences directly. I employ three different empirical techniques to shed light on underlying preferences and the distributional consequences of increased credit supply. First, I analyze the effect of unemployment on applications for refinancing credit and realized refinancing rates, with a focus on how this effect varies with other credit risk characteristics. I find that the overall effect of unemployment on refinancing is negative, although this effect varies considerably across borrower types. Low-credit risk borrowers, especially those with significant home equity, are more likely to refinance at higher unemployment rates, and likewise for these borrowers unemployment rates do not significantly increase the risk that they are unable to obtain credit. I interpret these results as evidence that for low-risk borrowers, the "demand" effect of greater liquidity preference outweighs the "supply" effect of reduced credit access, but not so for high-risk borrowers. Second, I study the effect of unemployment rates on the refinancing behavior of borrowers who were eligible for a government-sponsored enterprise program, the Home Affordable Refinance Program, that removed or severely attenuated the relationship between access to credit and unemployment. I find that within this cohort, a 1% increase in unemployment rates leads to a 3.4% to 4.4% increase in the annualized probability of refinancing, implying that the effect of unemployment on credit demand is positive. Finally, I employ a difference-in-differences strategy to estimate how takeup under the Home Affordable Refinance Program varied across high- and low-unemployment regions. My baseline estimates indicate that a 1% increase in unemployment rates leads to a 3% to 7% increase in program takeup, with borrowers in the top quartile of local unemployment rates refinancing under the program between 40% and 50% more often than borrowers in the lowest quartile. While these results combine both demand and supply effects, they do suggest that policies aimed at improving credit access will disproportionately benefit the unemployed.

In Chapter 4, "Guarantee Fees as a Countercyclical Policy Tool", I combine the insights from the prior two chapters to analyze how to use Fannie Mae and Freddie Mac's guarantee fees to regulate mortgage credit supply over the business cycle. My earlier results suggest that policies aimed at increasing the supply of refinancing credit, including reductions in guarantee fees, will be most-beneficial for borrowers during periods of widespread unemployment, and conversely that policies which constrain credit supply will be least-costly for borrowers during periods of low unemployment. Motivated by this finding, I simulate the effects of an alternative countercyclical guarantee fee policy that shifts credit supply from periods in which unemployment is low into periods when unemployment is high, and compare the results both to the observed path of guarantee fees and to other alternative credit supply policies. In order to quantify the effects of an alternative policy of lowering guarantee fees during high-unemployment periods, I design and estimate a structural model of the mortgage market with two important features. First, the model allows credit supply, and hence the credit constraints facing households, to be endogenously determined by, among other things, guarantee fees, and for the effect of fees to vary with the competitive landscape. Second, the model permits borrower valuations for refinancing credit, and thus the shadow value of credit constraints, to vary with individual-level employment status. These two features micro-found the results from chapters two and three and permit regime-invariant calculations of welfare under the counterfactual policy. I estimate the model with a novel combination of techniques, some of which have not been used previously in this literature, and use the estimated parameters to simulate the model under alternative policy scenarios. My structural estimates suggest that unemployed borrowers derive 130% to 220% greater utility from cash-on-hand than employed borrowers, and I find that the alternative guarantee fee policy results in 2.3% greater borrower welfare over an entire business cycle, while also reducing default rates and providing 6.3% greater revenue for Fannie Mae and Freddie Mac. My results also indicate that the effects of the alternative guarantee fee policy are quantitatively similar to a combination of a stricter loan-to-value-ratio cap and an affordable refinance program, two policies that have been either implemented or strongly considered by U.S. policymakers.

My research is intended to inform the policy debate surrounding the future of the government-sponsored enterprises and U.S. housing finance reform more generally. To the extent that the government maintains some role in the mortgage market, the question of how best to regulate credit provision over the business cycle will remain an important one. While other policy instruments are also no doubt necessary, countercyclical guarantee fees are a potentially valuable tool, and to my knowledge the question of how best to set them across the business cycle has not been considered previously in either the economics or policy literature. I hope that this research will prove valuable in the design of future housing finance policies.

Contents

List of Tables

List of Figures

Acknowledgements

Arrest this man, he talks in maths.

-Radiohead

I thank Saki Bigio, Jason Blevins, Gregory Brozak, Francesco Decarolis, Wendy Edelberg, Andreas Fuster, John Geanakoplos, Erik Hurst, Eric Johnson, Philip Kalikman, Steven Laufer, Damien Moore, Emi Nakamura, Michael Riordan, Miikka Rokkanen, Paul Scott, James Vickery, and Michael Woodford for helpful advice. I also thank seminar participants at Columbia University and conference participants at EconCon 2015 and the Southern Economics Association 2015 Annual Meetings for their comments. I am very grateful to my committee members, Chris Conlon, Tomek Piskorski, Jose Scheinkman, and Jon Steinsson, for their assistance, and particularly to my advisor, Kate Ho, for her constant support. I also thank Sandesh Dhungana, my collaborator on parts of this research, for all of his contributions. I especially thank Chris Mayer and the Paul Milstein Center for Real Estate at Columbia Business School for providing access to their data. All errors and omissions are my own.

Chapter 1

A Descriptive Analysis of the U.S. Mortgage Market and the Government-Sponsored Enterprises

1.1 Introduction

The economic experience of the U.S. and other advanced countries during the 2000s and early 2010s has highlighted the importance of credit markets to both individuals and the macroeconomy. For a variety of reasons, the relative contributions of which are still subject to considerable debate in the economics literature, credit supply expanded dramatically in the period roughly from 2003 to 2007 in the U.S. Partly as a result, household debt as a fraction of GDP increased by nearly 30%, driven largely by home-equity borrowing, and the U.S. underwent a simultaneous boom in house prices,[4,78] employment[166] and consumption.[161] Prompted by increasing losses on high-risk consumer lending, eventually resulting in a financial crisis, the supply of credit tightened beginning in 2007, and the attendant decline in borrowing and spending led to the deepest postwar U.S. economic contraction, now known as the Great Recession. The recovery from the Great Recession was considerably slower than from past U.S. recessions, and foreclosure and unemployment rates remained elevated for several years following the official start of the recovery in late 2009. Many have attributed the slow recovery to frictions in U.S. credit markets,[73,104,160,161,173] and in particular the refinancing friction inherent in fixed-rate mortgage (FRM) contracts.[34,65,136,165] Based on this assessment many have therefore called for policy intervention both to restrain the supply of credit during credit market booms[137] and to expland the availability of refinancing credit during recessions[1]. My dissertation analyzes the market for FRM credit in the U.S., with a particular focus on policy interventions designed to mitigate credit frictions and lean against destabilizing cycles in this large and systemically important market.

There are two aspects to such countercyclical policy interventions. On one side the goal of typical stimulative policies employed during downturns is to transfer funds

[1]jaffee2011future,scharfstein2011economics

towards those with high marginal propensities to consume. This for example is the theory behind unemployment insurance,[140,144] as the unemployed are likely to spend a substantial fraction of any income and hence support current aggregate consumption. On the other side, the goal of macroprudential policies is to lean against the buildup of systemic risk, often by targeting credit risk and overall leverage.[110] For this purpose many advocates of countercylical credit policies advocate for policymakers to utilize a "second instrument" in addition to the typical policy lever used by central banks, the level of short-term interest rates, as interest rates are a blunt tool and the level appropriate for managing aggregate demand may be inappropriate for managing credit growth and vice-versa.[24] While the Bank of England has developed a macroprudential toolkit that enables it to regulate, for example, mortgage lending standards,[24] a similar set of tools has not yet been developed in the U.S. In many ways, the U.S. agencies best-positioned to regulate the supply of credit across the business cycle are Fannie Mae (FNMA) and Freddie Mac (FHLMC)[2], and in my dissertation I analyze the effects of their policies and assess the prospects for potential alternative policies.

Since the Great Recession FNMA and FHLMC, the two largest of the government-sponsored enterprises (GSEs), have become the primary conduit for a variety of mortgage-market-specific policy interventions. Though they have a variety of tools and goals, one of the GSEs primary roles is to sell insurance against defaults on mortgages meeting certain criteria. Prior to 2008, the two agencies had been nominally private, although with strong government guidance and what many assumed to be an implicit government backstop. In September 2008, following combined losses of $109 billion in that year due to increasing default rates, the implicit guarantee became explicit as the two agencies were nationalized and their insurance obligations assumed by the federal government. Since then the GSEs have operated essentially as government

[2]These are not their real names. They are, respectively, the Federal National Mortgage Association and Federal Home Loan Mortgage Corporation, although the nicknames developed by traders proved popular and today they are better known as such.

agencies and have implemented two of the signature post-recession policy interventions in the housing market, a pair of programs designed to reduce default rates and increase refinancing. At the same time, at the direction of the U.S. government the two agencies have also considerably raised their default insurance premiums, known as guarantee fees, in order to repay the treasury for funds used to honor their prior obligations. These fee increases have predictable effects on the mortgage market, leading to higher interest rates, reduced refinancing volume, and greater default rates[3]. As a result, the GSE's realized pricing policies over this period were strongly procyclical[4]: during the boom from 2003 to 2007, they pursued a low-fee policy that eased credit constraints, and from the Great Recession onward they implemented a high-fee policy that tightened credit constraints. Moreover their policies vis-a-vis their guarantee fees worked against other concurrent policy interventions, for example by raising mortgage interest rates at a time when Federal Reserve interventions were working to lower them, or by increasing defaults and reducing refinancing at a time when other GSE policies were attempting to engineer the reverse.

Thus at the core of my dissertation is a policy proposal for the GSEs to set their guarantee fees countercyclically so as to lean against credit cycles. The object of such a policy would be to tighten credit constraints during boom times at which unemployment rates are low, and to then loosen credit constraints during and after recessions when unemployment rates are elevated. In order for such an alternative policy to be both effective and revenue-neutral relative to current policy the benefits from credit shifted into high-unemployment periods must exceed the loss from reduced credit supply during low-unemployment periods. In Chapters 2 and 3 of my dissertation I document two key stylized facts in support of this premise. First, in Chapter 2 I show that increases in guarantee fees, which function much like a tax on mortgage origination,

[3] I document each of these effects empirically in Chapter 2.

[4] The term "procyclical" is used in two different ways. In the policy literature, it typically implies "exacerbating business cycle fluctuations", while in the macroeconomic literature it ususally denotes "covarying positively with GDP". Throughout, whenever I use the term "procyclical" I intend the former.

lead to higher mortgage interest rates and, concurently, declines in the propensity to refinance and increases in the propensity to default with plausible elasticities. Then in Chapter 3 I use exogenous variation in the tightness of credit constraints induced by GSE policy interventions to provide evidence that the households that value the refinancing option the most, and as a result benefit most from credit supply expansions, are those in high-unemployment areas. The former finding could potentially be driven by a desire on the part of households facing unemployment spells to refinance for consumption-smoothing purposes,[51,123] while the latter may simply indicate that the first borrowers to be cut off from from credit markets when credit conditions tighten and last to return as credit loosens are the unemployed. As a consequence, the benefit to borrowers facing income shocks during recessions from increased credit supply induced by lower guarantee fees may exceed the loss from reduced credit supply owing to higher guarantee fees during booms. To quantify this intuition, in Chapter 4 I estimate a structural model of the mortgage market and use the model to simulate the effects of various GSE policies. The results of these simulations indicate that a countercyclical guarantee fee policy would benefit borrowers overall by combining both the stimulative and macroprudential aspects of countercyclical policy, and in quantitative terms would approximate the combined effects of a refinancing stimulus program and a tighter cap on household borrowing limits. Moreover while government agencies such as the Federal Deposit Insurance Corporation (FDIC) and Federal Housing Administration (FHA) have at times endeavored to implement through-the-cycle insurance pricing policies,[100,110,172] to my knowledge the proposal that the GSEs charge countercylical default insurance premiums is novel.

In addition to the policy proposal and counterfactual analysis described above, my dissertation makes several other substantive contributions to the existing economic and policy literatures. First, in Chapters 2 and refch:chapter3 I provide new empirical evidence on the effects of two key GSE policy interventions, namely, increases in guarantee fees and targeted refinancing programs. To my knowledge there is no prior direct evidence regarding the effects of the former, while I provide the first evidence regarding heterogeneity in takeup under the latter program. Second, my results in Chapter 3 provide new empirical evidence on the relationship between credit demand and unemployment and the distributional consequences of credit supply expansions that may inform both policy work and academic research on computable general equilibrium models. These two sets of results also help motivate the structural model that I develop in Chapter 4, although the model endogenizes both findings. In so doing I extend the current literature on household financial management, and in particular structural models of refinancing and default, in order to incorporate endogenous credit constraints and latent liquidity preference. I employ a novel set of empirical techniques in order to estimate the primitives of the model structurally using a unique dataset. My bottom-up approach to analyzing the mortgage market is, for reasons I will discuss later, particularly suitable to analyzing housing finance policy interventions, and while I focus only on simulations of a few GSE policies, my microfounded model structure is capable of analyzing a range of other policy interventions, including changes to GSE credit standards and Federal Reserve asset purchase programs. In this regard I hope the model and simulations I present in this dissertation prove valuable both to academics and practitioners.

The remainder of this chapter proceeds as follows. In Section 1.2 I provide a description of key institutional details regarding the mortgage market. As my focus is on the effects of GSE policy on mortgage borrowers, I describe in turn the manner in which credit constraints impact borrowers and the manner in which GSE policy affects credit constraints, and conclude with a discussion of recent GSE and other market-specific policy interventions. Then in Section 1.3 I survey the literature as it relates to my dissertation. Widespread perceptions that the Great Recession had its roots in the mortgage market have led to a proliferation of research in that area, and I describe a variety of recent research into household financial management, the determinants of credit supply, the causes and consequences of credit constraints, and the effects of GSE policies. In Section 1.4 I sketch a simplified model of the refinancing market and use the model to generate two key predictions in order to inform the results from Chapters 2, 3, and 4. Then in Section 1.5 I review a number of data sources used throughout my dissertation, which include detailed individual-level data on both the demand for and supply of refinancing credit, and describe in brief some procedures to clean and augment the data. Section 1.6 with an overview of the remaining analysis in Chapters 2, 3, and 4 and how the results from those chapters can inform both the academic and policy literatures.

My dissertation is intended to add to the debate surrounding the future of the GSEs and future U.S. housing finance policies more broadly. All evidence suggests that the GSEs in some form will continue to play an important role in the mortgage market moving forward,[76,108,189] and as a consequence their ability to regulate the provision of mortgage credit across the business cycle will remain a valuable policy tool. While other instruments must also be part of U.S. housing finance policy, a countercyclical guarantee fee policy may prove a potentially valuable and as yet underexplored addition. I hope that this research contributes to the design of future beneficial housing finance policies.

1.2 Institutional Background on the Agency Mortgage Market

In this section I describe the institutional arrangement of the agency mortgage market. The primary focus of my dissertation is the transmission of changes in GSE policy to borrower-level outcomes and ultimately welfare. As such, the purpose of this section is to establish the institutional features that allow this transmission to take place. In explaining these mechanisms, I first outline the process of mortgage refinancing, with particular emphasis on how credit constraints can impede borrower's choices. I then turn to the supply side and explain the process by which mortgage loans are originated, the incentives facing mortgage lenders, and how GSE policy can affect the willingness of lenders to extend credit and the prices they charge. I conclude with a discussion of actual policy changes with a direct impact on the agency mortgage market implemented from the start of the Great Recession through the end 2012.

1.2.1 The Mortgage Refinancing Process

Mortgage refinancing constitutes a large market in the U.S., and choosing when to refinance is perhaps the most important financial decision a typical household will make. Through a combination of government policy and historical accident, homeownership and FRMs are much more common in the U.S. than in most other countries. The market share of FRMs usually hovers around 70%, while the total volume of refinancing can range from $1-1.5 trillion per year, as it has in the wake of the Great Recession, to as much as $2.75 trillion in 2003 during the peak of a refinancing wave. FRMs feature level payments each month regardless of changes in interest rates, so as mortgage interest rates decline, borrowers who took out mortgages previously at high interest rates can often save money on their monthly payments by taking out a new loan at a low rate and paying off the old loan. Borrowers can also further reduce their monthly payments by extending the term of their loan in order to spread their payments over more periods. Although relative to adjustable-rate mortgages (ARMs) FRMs reduce the exposure of borrowers to interest rate risk, they introduce the problem of how and when to refinance. Deciding when to refinance and shopping effectively for a competitive interest rate offer is a difficult financial problem, and some evidence suggests households behave sub-optimally in this regard[5].

[5] See Agarwal et al (2013),[7] Agarwal and Evanoff (2013),[8] Bucks and Pence (2008)[46] and Woodward and Hall (2012)[188] for a discussion of lack of borrower sophistication in the mortgage market.

Homeowners can also in many cases choose to increase the balance on their mortgage via cash-out refinance. This tool is especially valuable given the composition of middle-class U.S. household balance sheets. According to the 2013 Survey of Consumer Finances, 67% of U.S. households own a residence, with a median value of $170,000. By contrast just 16% of households own stocks or bonds, with a median value of $25,000[6]. While households can easily convert financial asset wealth to consumption spending by selling their assets in liquid markets and spending the proceeds, monetizing housing wealth is comparatively difficult. Houses are indivisible, and one cannot sell half a house. Moreover, real estate markets are highly illiquid, sellers have high transaction and search costs, and houses often have sentimental or flow consumption values that make selling a home unattractive. Cash-out refinance provides an alternative to sale that allows households to monetize their housing wealth by essentially selling a fraction of their home equity. When their home is worth more than the principal outstanding on their mortgage, households can refinance into a higher balance, giving them free cash to spend on consumption goods after paying off the existing loan. Greenspan and Kennedy (2008) [101] suggest that the four most-common uses of home equity wealth extracted via cash-out refinancing are to consolidate non-mortgage debts, to make home improvements, to spend on consumption goods, and to finance the purchase of financial assets. Three of these uses (all save consumption expenditure) are related in that they are essentially investments. Mortgage debt is particularly suitable as a means to finance such investment because by virtue of being secured by real property and having tax-deductible interest payments in the U.S., the after-tax interest cost of mortgage debt is much lower than that on other debts such as credit cards and personal loans, making it cost-efficient for borrowers to increase the size of their mortgage in order to retire other types of debt. These interest savings function equivalently to those from rate-term refinancing, or refinancing into a similar balance at lower rates in order to reduce one's

[6]The quoted medians are among those who actually hold that asset class, so across all households median residential real estate assets are much higher than median financial assets.

monthly payments, in easing household budget constraints. However, choosing how much equity to extract when refinancing is also a difficult decision, as any increase in debt reduces the borrower's equity stake in the home and thus increases the risk that their wealth will be eroded completely if house prices decline. Hence withdrawing mortgage equity can potentially increasing a household's risk of default[7], and as with refinancing overall some evidence suggests that households sub-optimally over-leverage[8].

Moreover, borrowers also face substantial fixed costs in completing a refinance application, including both monetary and non-monetary costs. In a typical timeline, a borrower will supply several potential primary market lenders (known as "originators") with their desired loan balance and estimates of certain credit risk information, such as their credit score and income. Due to the difficulty in navigating the refinancing market, in many cases borrowers actually contract with brokers or correspondents ("third-party originators" or TPOs) who shop for the best available offer in exchange for a fee, saving borrowers the time and search costs. The originator will quote a best-offer interest rate based on these estimates, subject to revision on verifying their credit risk. Using these offers, the borrower will make a formal credit application to one or more originators, typically paying an application fee, at which point the originator will verify their credit risk (often involving additional fees, such as for home appraisals or document review) and decide whether or not to grant the application. If the originator grants the loan application, they will then send out a rate offer sheet detailing the tradeoff between the interest rate the borrower pays, any upfront fees ("points") they must pay or be paid to secure the offer, and the length of time ("lock-in period") they have to decide whether or not to accept the offer. Borrowers who accept pay off the balance remaining on their existing mortgage as well as, in some cases, a prepayment fee that is often a stipulated

[7] See Laufer (2013) [145] and Mian and Sufi (2009, 2011) [158,159] for a evidence of increasing default risk from cash-out refinance.

[8] See Duca and Kumar (2014) [69] or Disney and Gathergood (2011) [67] for evidence on the link between financial literacy and mortgage equity withdrawal.

fraction of the existing balance.

In some cases, borrowers who would like to refinance may be unable to do so due to credit constraints. Borrowers with risky credit profiles, measured as either a high loan-to-value (LTV) ratio, high debt-to-income (DTI) ratio, or low credit score (FICO), are more likely to have their refinancing applications denied[9]. Moreover even in cases where a risky borrower's application is approved, the high interest rate charged by the lender to compensate for that risk may be too high for the borrower to pay, leading the borrower to turn down the offer. Some evidence suggests that mortgage borrowers may not have wanted to deleverage so rapidly following the Great Recession, but were constrained in their ability to access credit by reductions in credit supply[10]. Hence easing credit supply conditions may improve borrower welfare both by reducing the interest rates offered to borrowers and by reducing the likelihood that borrower applications will be denied. Moreover to the extent that refinancing at a lower rate or longer term can reduce a borrower's monthly payments, constraints that inhibit the capability of borrowers to refinance their mortgages during recessions may increase default rates as borrowers who are unable to afford their monthly payments turn to their next-best alternative[11].

On a macroeconomic level, the mortgage market is in many ways a key driver of business cycles and refinancing in particular is an important channel of monetary transmission in the U.S. Because the majority of U.S. home mortgages are fixed rate, in order to benefit from a decline in interest rates most mortgage borrowers must actively choose to refinance in order to lock in a new lower rate. However, the ability of fixed-rate borrowers to refinance depends on their creditworthiness, and in particular, if a borrower does not have sufficient equity it is typically very difficult for them to obtain

[9] See Goodman and Li (2014) [95] for a discussion of how these characteristics affect access to credit.

[10] See Gropp et al (2014) [102] and Bhutta (2012) [37] for a discussion of credit supply restrictions during and after the Great Recession.

[11] See Gerardi et al (2013, 2015), [92,93] Keys et al (2014), [136] and Tracy and Wright (2012) [185] for evidence of the effect of refinancing constraints on default.

credit. These constraints on refinancing due to tightened credit supply are considered by both academics and policymakers to be a contributing factor in prolonging recessions[12]. In addition, some evidence suggests that homeowners have a higher propensity to spend out of income during recessions, when they are more likely to be unemployed or otherwise income-constrained[13]. Likewise, loosening credit constraints during expansionary periods may lead to over-borrowing and contribute to unsustainable asset price appreciation and increasing economic fragility[14]. The experience of the U.S. over the most-recent business cycle highlights the importance of mortgage markets, and in particular variations in mortgage credit supply, as a driver of macroeconomic outcomes.

1.2.2 Mortgage Originators

Mortgage markets in the U.S. are highly segmented, and in most cases the primary originator who issues a mortgage loan to a household is not the ultimate investor.[181] Historically, primary originators were often small local enterprises with specialized knowledge in evaluating borrowers' creditworthiness, although this has changed in recent years as primary originators consolidated and expanded geographically following a wave of deregulation[15]. Smaller primary mortgage originators are often unable to bear the risks associated with issuing and holding FRMs. Because they mortgages only in one area, locally-concentrated originators are highly exposed to idiosyncratic local economic shocks, such as a decline in oil prices causing a rise in mortgage defaults in an oil-producing region. Many specialized nonbank mortgage lenders also lack the capital necessary to hold a large volume of loans on their balance

[12]See Mishkin (2007),[165] Bennett et al (2001),[32] and Black et al (2010)[39] for a discussion of the refinancing channel, or Guerrieri and Lorenzoni (2011),[104] Midrigan and Phillipon (2011),[173] and Tenreyro and Thwaites (2013)[184] for a discussion of the effect of credit constraints on recessions.

[13]See Cooper (2013),[59] Zhou and Carroll (2012),[191] Jappelli and Pistaferri (2014),[123] and Carroll et al (2014)[51] for evidence on income shocks, home-equity borrowing and the marginal propensity to spend.

[14]See Khandani et al (2013),[137] Duca et al (2011),[70] Chen et al (2013)[53] and Laufer (2013)[145] for a discussion of the the contribution of weakening credit standards to individual and systemic risk.

[15]See Loutskina and Strahan (2011)[150] for an overview of this trend.

sheet, and these originators account for a fairly high market share in the U.S.. Mortgage holders also face prepayment risk[16] because borrowers tend to refinance when mortgage rates are low, so from the investor's perspective refinancing exchanges a high-yielding note for cash at precisely the time when that cash cannot be reinvested at a high rate. Moreover, FRMs are long-dated assets and their value falls significantly when interest rates rise, although due to prepayment risk their value does not rise as much when interest rates fall[17][18]. Certain institutional investors, especially pension funds and insurers, also have long-maturity liabilities and as such are less susceptible to asset-liability maturity mismatch and hence better able to bear the risk associated with holding high-duration FRMs to maturity, while other investors such as hedge funds are better able to hedge away prepayment risk. Since these larger investors often lack expertise in evaluating borrower creditworthiness, a common arrangement is for primary originators to specialize in underwriting mortgages and then sell the loans to investors on the secondary mortgage market.

For the majority of secondary market sales, mortgage loans are packaged into a mortgage-backed security (MBS) through a process called securitization that pools risks across multiple borrowers prior to sale. In the simplest form of securitization, a primary originator bundles multiple previously-originated FRMs into a simple MBS known as a pass-through pool. The monthly proceeds from this pool depend on the weighted-average coupon (WAC), or average interest rate paid on the underlying mortgages. Investors in the pool receive a share of the proceeds proportional to their stake, and the price they are willing to pay for the pass-through pool depends on the

[16]Prepayment denotes any event in which the mortgage balance is paid off in full prior to maturity, including refinances and sales.

[17]The asymmetry of this relationship is known as negative convexity.

[18]Due to prepayment and default the duration of a mortgage, or the sensitivity of their price to interest rates, is actually random, and most market participants use model-derived estimates of expected duration. Absent any prepayment or default, the scheduled payments on 30-year FRMs have a duration of 7 to 12 years depending on the interest rate, so market participants often gauge mortgage yields relative to 10-year treasury bonds. The realized duration of most mortgages, factoring in prepayment, is between 5 and 10 years.

WAC, as well as other risk factors. By pooling the risk of prepayment and default across multiple mortgages MBS reduce the risks associated with investing in mortgages. Primary originators receive the cash from selling the securities upfront and a fee, known as a servicing fee, on an ongoing basis as compensation for collecting payments from borrowers, which then can choose to retain for themselves or use to hire a third-party servicer[19]. On the cost side, many primary originators must secure the cash necessary to complete the loan origination for the entirety of the lock-in period, and this money is often borrowed from larger warehouse lenders. Originators thus face the risk that a borrower will reject their credit offer prior to the end of the lock-in period, leaving them with no mortgage with which to repay the interest on their warehouse loan. Borrowers often reject offers because prevailing mortgage interest rates have declined, which tends to *ceteris parabus* increase the value of existing mortgages, so originators also face pricing risk in that borrowers will tend to disproportionately accept offers when the value of their mortgage has declined.

In order to reduce their exposure to this pricing risk, a large fraction of simple pass-through pools are sold forward in a standardized futures market known as the TBA market. As described in Vickery and Wright (2013),[186] originators can sell MBS at any time for delivery at a specified date in the future. Importantly, the actual securities to be delivered are not specified at the time of the trade, rather, only the general features such as the coupon of the security are arranged in advance[20]. Delivery dates are scheduled well in advance by the Securities Industry and Financial Markets Association, and there is usually one delivery date each month for each class of securities, although 15-year mortgages and 30-year mortgages for example will have different delivery dates. Two days prior to the delivery date, TBA sellers must inform buyers of the securities they will deliver, and failure to deliver results in harsh punitive payments that it is essentially

[19] A simplified pictorial representation of the cash flow between different participants in an agency MBS transaction is shown in Figure 1.1 and will be discussed more fully in Section 1.2.3.

[20] Only certain types of MBS, specifically those backed by the GSEs, are eligible for inclusion in a TBA deal.

always in the incentive of the seller to avoid. Because the composition of TBAs is hidden until two days prior to delivery, prices reflect market expectations of the value of delivered securities, with a downward adjustment to account for the fact that lenders will preferentially deliver their lowest-quality mortgages. Moreover because there are a relatively limited set of categories of TBAs[21], and fewer still with liquid markets, the mortgage rates that originators are willing to offer are substantially determined by a narrow set of market prices.

However, mortgage credit supply is not entirely driven by secondary market conditions, and lender strategies at the primary-market level are an important determinant. In many markets, mortgage lending is a concentrated industry[22] and lenders tend to behave as if they have a certain amount of market power[23]. An important consequence is that in concentrated markets, lenders may be able to profitably charge higher interest rates without fear that the borrower will reject the offer and apply to a different lender. In addition lenders are able to preferentially sell or securitize their riskiest loans and hold their least-risky loans on balance sheet[24]. The anonymized nature of the TBA market only increases the scope for adverse selection, although there is also a secondary market for specific MBS with superior risk features that enables these pools to sell at a premium. In addition, while only government-guaranteed mortgages can be pledged to a TBA transaction, mortgages can also be privately securitized through broker-dealers; while issuance in the private-label secondary market was significant prior to the great recession, it has since effectively ceased. Thus the extent to which credit supply varies with secondary market prices will depend on how reliant the lender is on securitization, as lenders with balance-sheet capacity can hold the mortgages they

[21] TBA coupons are traded in .5% increments.

[22] See Loutskina and Strahan (2011) [150] for a discussion of how deregulation has driven increasing expansion by large lenders and Scharfstein and Sunderam (2013) [178] for a discussion of local-market concentration.

[23] See Allen et al (2013, 2014), [14,15] Amromin and Kearns (2014), [17] Scharfstein and Sunderam (2013), [178] and Gurun et al (2013) [105] for examples of monopolistically-competitive behavior in the mortgage market.

[24] See Agarwal et al (2012), [5] Downing et al (2009), [68] and Keys et al (2012) [134] for a discussion of adverse selection.

issue until conditions improve.

1.2.3 The Government-Sponsored Enterprises

Absent government intervention, the secondary market for mortgage loans would become a lemons market. As mentioned above, primary originators have a strong incentive to use the private information gained from screening borrowers to retain on their balance sheet their least-risky mortgages and sell their most-risky mortgages. In order to support a functioning secondary mortgage market, FHLMC and FNMA offer to guarantee repayment of all principal and interest on a mortgage meeting their standards on schedule in exchange for a fee. To the extent that investors price the GSEs' implicit government guarantee, this removes all default risk on agency MBS[25], allowing primary originators to in effect pay an insurance premium, known as a guarantee fee, to convert a risky mortgage loan into a stable and relatively safe callable annuity. Note that this does not actually remove the incentive for adverse selection; primary originators still have an incentive to securitize their highest-risk loans, but the GSE guarantee makes investors indifferent to the idiosyncratic default risk associated with these mortgages, as FHLMC and FNMA bear all of the credit risk[26]. In order to mitigate adverse selection, FHLMC and FNMA specialize in rating and pricing default risk on mortgages they guarantee, and vary their fees depending on the primary originator and characteristics of the loans.

While my dissertation focuses on the policies of FNMA and FHLMC with respect to their guarantee fees and their support of secondary markets, these two agencies have other policy goals and tools, and there are also a variety of other government-sponsored enterprises[27]. One of FNMA and FHLMC's other important policy instruments is

[25]This assumes all market participants act (as it appears that they do) as though the U.S. Government will guarantee the obligations of the GSEs in case of any serious problem (as it appears that they will).

[26]Agency MBS still face prepayment risk, or the risk of lower yield due to borrowers refinancing their loans in low-interest-rate environments, but there is less scope for adverse selection with regard to prepayment risk.

[27]Throughout I use the term "GSE" to refer solely to FNMA and FHLMC.

portfolio purchases of either agency MBS[28] or privately-issued MBS[29]. The GSE's portfolio purchases make them some of the largest institutional investors in the mortgage market, and through these MBS purchases they are able to affect originator incentives to underwrite certain types of loans. Crucially, the GSEs are charged with increasing the availability and affordability of credit for homeowners in low-income areas, and a certain mandated fraction of their portfolio purchases must contribute towards this goal, typically by purchasing loans originated in low-income areas. Some commentators have suggested that efforts to meet these affordable housing goals contributed to an unstable expansion of risky credit prior to the Great Recession[30], although contemporary research does not support this proposition[31]. By their charter the GSEs are only allowed to purchase or guarantee mortgages meeting certain quality standards, and importantly there are limits on the size of loans meeting their standards. Loans for amounts greater than this threshold, the conforming loan limit, are referred to as "jumbo" loans and are ineligible for the GSEs guarantee. The GSEs have also historically varied credit risk standards for loans they underwrite, by example changing their income documentation requirements or minimum down-payment requirements, and evidence suggests both that their changes to their credit standards have important consequences for credit supply [175] and tend to be procyclical and exacerbate credit cycles [83] and that these changes. Other organizations and government-sponsored enterprises fulfill a role similar to the GSEs in other settings, as for example Farmer Mac (another GSE) supports secondary markets for agricultural mortgages in a similar fashion and GNMA (a government organization) securitizes and guarantees home mortgage loans issued by the Department of Veterans Affairs (VA) and the FHA. In the remainder of my analysis on focus on the GSEs guarantee fee business for FRMs in

[28]Those backed by FNMA, FHLMC, or another non-GSE company, Ginnie Mae (GNMA).

[29]Those arranged by a private financial institution with no government backing.

[30]See Acharya et al (2011) [3] or Calomiris and Haber (2015) [48] for variations on this argument.

[31]See Avery and Brevoort (2015) [21] or Bolotnyy (2014) [41] for evidence on the role of affordable housing goals in the late-2000s mortgage crisis.

particular, although it is worth noting that other aspects of their policy stance also have important macroeconomic consequences.

Within their guarantee fee business, the GSEs indirectly influence the provision of credit in the primary market by directly altering originator incentives to underwrite new mortgages. Recall from Section 1.2.2 that the prices of MBS depend on the WAC, as shown in the simplified schematic of a securitization transaction in Figure 1.1. As shown in the diagram, any increase in guarantee fees leaves less of the WAC available to divide between the coupon and the originator's excess servicing spread, both of which determine the originator's profit on the transaction. For each security an originator offers to FNMA or FHLMC, the GSE quotes two types of fees, an upfront "delivery fee" assessed immediately as a proportion of the principal balance and an ongoing "guarantee fee" assessed on a continual basis. Due to differences in prepayment speeds and hence duration across mortgages, it is difficult to make apples-to-apples comparisons of the cost to originators of delivery fees relative to guarantee fees. However, the GSEs permit originators to freely convert upfront delivery fees to annualized guarantee fees and vice versa by dividing them by a present value multiple. This "buy-down" (if reducing the ongoing fee) or "buy-up" (if increasing the ongoing fee) multiple, typically a number between 4 and 7[32], varies negatively with expected prepayment speed, so at higher expected refinancing rates the upfront fees translate into larger ongoing fees. Hence in this analysis I treat guarantee fees and annualized delivery fees equivalently and refer to the combined annualized figure as the guarantee fee. After choosing the pool coupon and paying any fees, the primary originator then sells the MBS on the secondary market, in many cases to one of the GSEs themselves[33], and in some cases also sells the servicing rights, which trade at present-value multiples similar

[32] See Fuster et al (2013) [84] for more details on converting between ongoing and up-front fees.

[33] There are many additional options available to originators securitizing via the GSEs. In particular, securitization is often locked-in through futures contracts before the mortgage loans are actually issued, as described in Section 1.2.2, and there are separate products for other types of loans such as ARMs and 15- or 40-year mortgages. This exposition describes only the most basic securitization of 30-year FRMs for the purposes of making clear how a primary originator would respond to GSE policy.

Figure 1.1: Simplified Schematic of an Agency MBS Deal

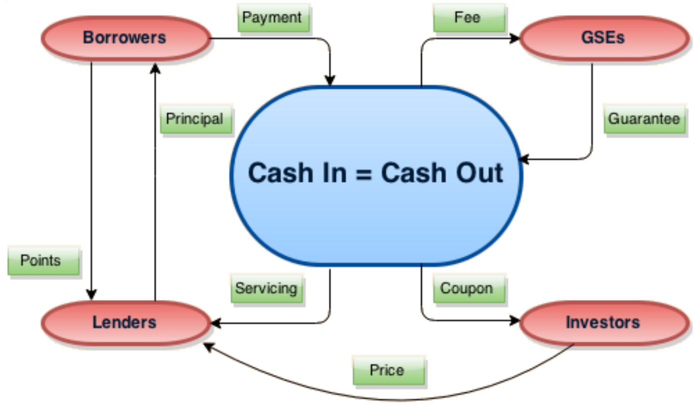

Notes: Red boxes represent agents, green boxes represent cash flows, arrows represent payer and recipient of cash flow. MBS shown in circle. Borrowers pay monthly payments and points at origination. GSEs recieve fee and guarantee payments in case of default. Investors pay MBS price and receive coupon payments. Lenders receive servicing payments monthly and points and price of MBS at origination.

to those for up-front fees. In the case that a borrower defaults on their mortgage, the GSE will continue making regular payments for a specified length of time, typically 120 days, and if the borrower has not resumed making payments, the GSE will then pay off the remaining balance. Hence while GSE securitization essentially removes all risk of default, all else equal, higher guarantee fees lead to some combination of lower MBS coupons and lower servicing spreads, reducing the profitability of a securitized mortgage.

In determining the disposition of an originated mortgage or "funding" channel, primary originators typically compare multiple options[34]. Originators usually decide how to fund a loan after deciding upon the interest rate and whether the loan application will be approved, and subsequently compare whether to securitize the loan via GNMA, the GSEs, or private-label, and whether to retain the loan on balance sheet for their own portfolio. GSE securitization and portfolio placement are typically reserved for the highest-quality credits, while private-label and GNMA securitization

[34]The remarks in this paragraph on how lenders make mortgage funding decisions are based on discussions with industry participants at the Federal Reserve Bank of New York's Mortgage Contract Design Conference in May of 2015. Owing to Chatham House Rules in effect for the conference, I am unable to attribute these remarks to particular individuals.

are only used for lower-quality credits. For mortgages where GSE securitization and portfolio placement are the viable options, primary originators will be more willing on the margin to originate mortgages and to securitize when the prices those mortgages command on the secondary market are higher, or equally when guarantee fees and upfront fees are low. Prior studies have demonstrated that active secondary markets increase the supply of mortgage credit in the primary market[35]. If the primary originator intends to hold a mortgage on balance sheet, their risk and expected profit depend on factors affecting default and prepayment risk, such as the borrower's income and credit score. If the primary originator intends to sell a mortgage loan, their decision whether or not to grant credit also depends on the price the loan will command on the secondary market, and thus the guarantee fee. For the conforming 30-year FRMs on which this paper will focus, GSE securitization has retained a dominant position since the start of the Great Recession, and what little private exposure has persisted in this market segment has been mostly originators holding loans in portfolio.[82,96]

Between March 2008 and April 2012, the GSEs increased their fees seven times on various mortgage products. As shown in Figure 1.2, the result was that by 2013 average guarantee fees were more than double their 2007 levels[36]. The GSEs have always varied their guarantee fees based on experience with particular mortgage originators, but an important change in policy over this period was to introduce considerably more risk-based pricing of delivery fees at the loan level. Prompted by increasing default rates on guaranteed mortgages, in March 2008 FHLMC and FNMA raised their delivery fees by 25 basis points (bp) on all newly-securitized loans and began varying their delivery fees based on FICO scores and LTV ratios for all loans[37]. These fees were raised again for

[35]See Gabriel and Rosenthal (2007)[86] and Loutskina and Strahan (2009)[149] for examples.

[36]Figure 1.2 shows the average charged fee, which is unadjusted for sample composition. The decline in average fees in early 209 reflects tightened credit conditions that locked higher-risk, higher-fee borrowers out of the market. Hence while average fees were lower in 2009, the actual fee schedule was greater.

[37]Prior to March 2008 guarantee fees varied depending on some loan characteristics, such as whether the borrower was an investor and whether the property was an apartment or single-family home, but only certain types of loans ("A-Minus" Mortgages) were priced differentially depending on FICO and LTV.

some loan types, including cash-out refinances, high-LTV loans, and low-FICO loans in June 2008. In late July 2008 the Housing and Economic Recovery Act significantly increased government oversight of the ostensibly private GSEs, and in September 2008 FHLMC and FNMA were officially placed into conservatorship under the newly-created Federal Housing Finance Agency (FHFA), becoming wards of the state. Subsequently delivery fees were raised three more times, in November 2008, April 2009 and March 2011, on high-LTV and low-FICO loans. In April 2012 and again in December 2012 guarantee fees were raised by 10 basis points across the board, the former mandated by congress to fund the payroll tax cut and the latter mandated by FHFA [80][38]. On a mortgage with a $200,000 balance, a 10 bp increase in guarantee fees will translate to roughly $4,000 in additional payments over the life of the mortgage, so each of these fee increases was economically meaningful. With the exception of the congressionally-mandated increase, all of these fee increases were intended to correct for perceived mispricing of default risk. As shown in Figure 1.3, which is based on FHFA's annual report to congress, [80] the GSEs did not merely adjust their fees upward in order to make them actuarially fair; rather, they set their fees so as to make positive profit in 2009-2012 in order to make up for losses incurred in 2008 and previously[39]. The extent of this intertemporal cross-subsidization was higher for refinance loans, and cash-out refinances in particular, than it was for home purchase loans. Despite the stated rationale for the other fee increases, the congressionally-mandated fee increase highlights the extent to which the GSEs have become an explicitly public enterprise, and it would not be a novel development for the government to treat guarantee fees as an instrument of public policy.

[38]A typical rule-of-thumb conversion from delivery fees to guarantee fees uses a present-value multiple of 5, so a 10 bp increase in guarantee fees corresponds to a 50 bp increase in delivery fees, or double the increase in March 2008

[39]Note that both the GSEs themselves and their conservator acknowledge both this cross-subsidization and its rationale.

Figure 1.2: Average Charged Guarantee Fee

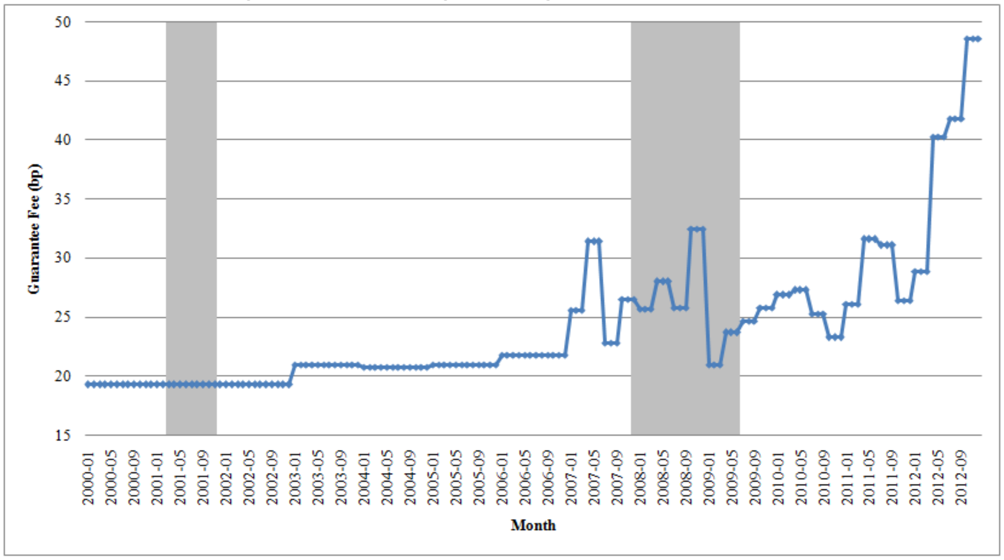

Notes: Average guarantee fees charged on FNMA MBS. Grey shading indicates NBER recession dates.

Figure 1.3: Intertemporal Guarantee Fee Cross-Subsidization

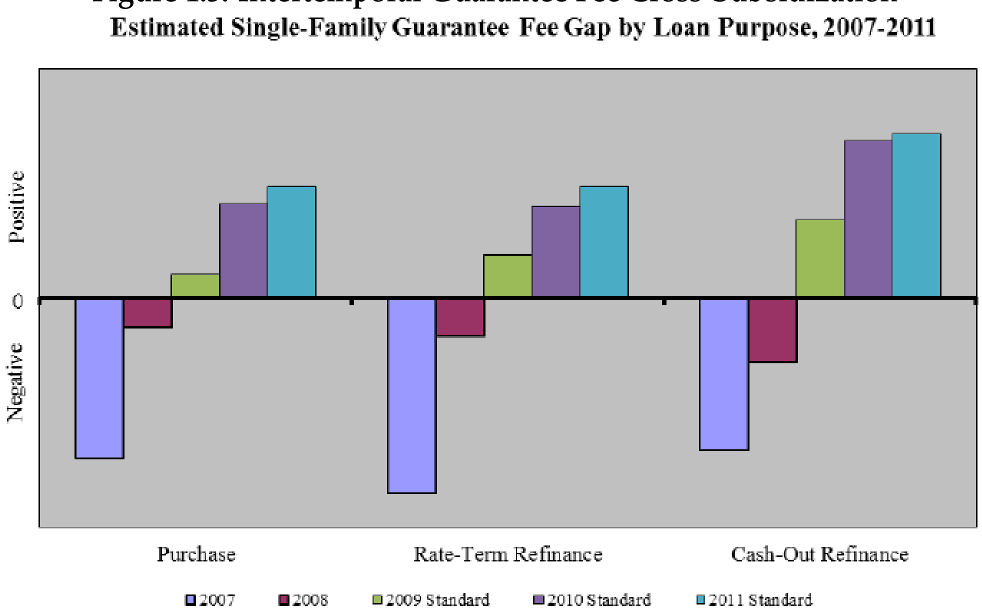

Notes: Estimated gain/loss on GSE mortgage guarantee business.

1.2.4 Policy Interventions in the Mortgage Market

In addition to changes in guarantee fees, since start of the Great Recession there have been a number of new policy measures implemented in the U.S. directed specifically at the U.S. mortgage market. The most relevant intervention for my analysis, as mentioned in Section 1.2.3, was the bailout of the GSEs themselves. Starting in early 2007, defaults began rising among mortgages guaranteed by the GSEs, and in 2008 the GSEs lost a combined $109 billion on their guarantee business. In order to forestall a loss of investor confidence in the GSE's guarantee that could disrupt secondary markets, the federal government committed in September 2008 to honor all of their obligations in exchange for direct oversight and an 80% equity stake in the enterprises, effectively nationalizing them. The GSEs lost nearly as much combined in 2009, and over time more equity injections were required, although since 2012 the GSEs have been profitable. Since that time the FHFA has exercised significant authority in directing the GSEs business operations, beginning with the dismissal of their top executives and continuing with restrictions on their political activities and dividend payments, among others. The agencies have also been required to remit a substantial fraction of their earnings to the treasury, and as of 2015 have repaid all of the bailout funds they received.

Under the direction of the FHFA, the GSEs have also been the conduit for two key pieces of mortgage-market-specific policies. As shown in Figure 1.4, in January 2009 the GSEs introduced the Home Affordable Modification Program (HAMP), a program designed to help borrowers avoid default. HAMP offered mortgage servicers incentive payments to modify the terms of delinquent or at-risk mortgages in order to reduce their monthly payments for a specified length of time. The goal of the program was to shift servicer incentives away from foreclosure and towards forbearance (delayed payments) via a subsidy. The program was generally considered to be unsuccessful,[113] as relatively few borrowers received modifications and many program participants re-defaulted as soon as the period of lower payments expired. As shown in Figure 1.4, in June 2009 the FHFA also created the Home Affordable Refinance Program (HARP) to streamline the refinancing process for borrowers with little home equity. As discussed in Section 1.2.1, borrowers with little or no home equity are often unable refinance, and the decline in home prices starting in 2007 left many borrowers in this position and thus unable to benefit from sizable reductions in both short-term interest rates and mortgage rates. The HARP program allowed these borrowers to refinance, provided they met a number of eligibility requirements, including that their original loan be guaranteed by FHLMC or FNMA, that the loan be originated prior to June 2009, and that they be current on their mortgage with a current LTV above 80%. Under the direction of the FHFA, FNMA and FHLMC created special programs[40] that allowed lenders to securitize with or sell to the GSEs on very favorable terms loans meeting these stipulations, again with the goal of incentivizing lenders to grant HARP refinances. The initial program guidelines had very strict requirements for the guarantees lenders made regarding the quality of their loans, placing lenders at risk that the GSEs would find fault with their underwriting standards, cancel the guarantee, and transfer the defaulted loan back to the lender, and also originally featured a 105% LTV ratio cap for eligibility. Because lender participation and borrower takeup were so low relative to predictions, both of these features were relaxed in several stages starting in late 2009. In response to what was widely deemed a failure, the FHFA created HARP 2.0 in January 2012, removing the LTV caps entirely and reducing the legal risks facing lenders.

In Chapter 3 I exploit the design of the HARP program extensively in my empirical design, hence a more detailed discussion of these design features is warranted. As mentioned above, the HARP program only allowed borrowers to refinance if they met a number of eligibility requirements. First, the loan must have been guaranteed by FNMA or FHLMC, as the program was run through the GSEs. Second, the loan must have been

[40] Respectively, DU Refi + and Relief Refinance.

Figure 1.4: Timeline of Mortgage-Market Interventions

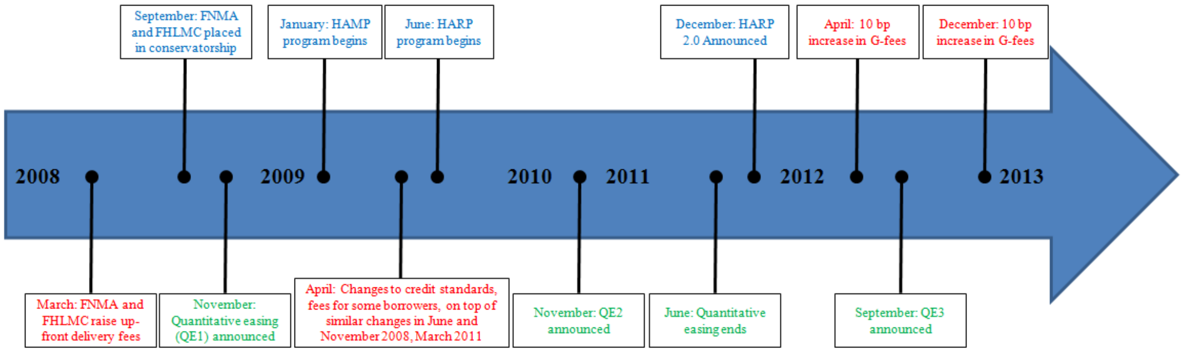

Notes: Changes to GSE fees shown in red, other GSE policy interventions shown in blue, Federal Reserve interventions shown in green.

sold to or securitized by one of the GSEs prior to June 2009, the official start date of the program[41]. Third, the borrower must have had a current LTV ratio of at least 80%, and, at various times in the program's history, an LTV ratio below some cutoff. Fourth, the loan size cannot have been above the GSE's maximum size (the aforementioned conforming loan limit) and cannot have been originated as part of certain special GSE programs[42], although in most cases loans other than 30-year fixed-rate mortgages were permitted[43]. Finally, the borrower must have been current on their mortgage at the time of the refinance, with no missed payments in the past six months and no more than one missed payment in the past twelve months. There was an additional requirement placed on lenders that the refinance loan provided tangible benefit to the borrower, although lenders could meet this requirement either by offering a lower interest rate or by moving from an adjustable-rate to fixed-rate mortgage, and as such almost all refinances qualified, as average interest rates declined sharply prior to program implementation.

Despite the fact that these HARP-mandated programs provided strong incentives for lenders to refinance eligible borrowers regardless of their creditworthiness, lender discretion still played some role in determining access to credit. Typically, by selling to

[41]This requirement had the effect of preventing borrowers from using HARP twice.

[42]In particular, Alt-A mortgages were excluded. Combined with the first requirement, this essentially restricted the HARP programs to only prime borrowers.

[43]For example, 15- and 40-year mortgages qualified, as did most varieties of ARMs.

or securitizing with one of the GSEs the lender transfers all credit risk to the GSEs, and if this can be done cheaply then the lender can often issue mortgages profitably with very little risk. In choosing whether or not to grant refinancing credit to a HARP-eligible borrower, lenders had to account for two principal risk factors. First, if the lender did not provide credit, the borrower could simply apply to another lender, costing them lost business[44]. Second, if the lender did provide credit, they faced GSE putback risk, or the risk that the GSEs (or an investor) would find flaws in the underwriting or documentation and transfer all credit risk back to the lender. Under the original HARP guidelines, the lender on the borrower's existing mortgage[45] had a dramatic advantage over other lenders in offering a HARP refinance because of differences in documentation requirements and putback risks. These guidelines had the effect of muting competition for HARP refinances by conferring a pricing advantage to one lender relative to all others[46], simultaneously reducing both risk factors. Nonetheless, at all times in the program's history borrowers could and did go to lenders other than their original lender. Moreover, while lender participation in the HARP program was totally voluntary, most lenders did in fact participate, signaling that they believed loans issued under the HARP guidelines to be profitable. However as a result of this program design, borrowers with impaired credit were still required to find a lender willing to overlook their credit risk in order to refinance.

All the same, the HARP program was developed in order to help high-credit-risk borrowers refinance, and several features of the program attenuated the effect of borrower credit risk on access to credit. The original guidelines left scope for lenders to deny eligible applicants on the basis of their increased credit risk, and some borrowers were denied because of, among other things, low FICO scores or high LTV ratios[47].

[44]If the original lender actually held the mortgage note or securities in which the note was packaged, they also lost in this scenario from unprofitable prepayments on their loan portfolio.

[45]Technically the servicer, or the entity that administers the loan, but in many cases these were the same bank.

[46]See Amromim and Kearns (2014) [17] for a more detailed discussion of this asymmetry

[47]The changes to HARP introduced in January 2012 severely reduced the scope for denying applicants.

However, several factors mitigated the otherwise highly negative effect of non-employment on credit access. First, under the program guidelines, in cases where the borrower refinanced with their original lender and their payments did not increase by 20% or more lenders were not required to re-verify either their income or employment status. Since in almost all cases borrower payments declined[48], and in most cases borrowers refinanced with their original lender[49], a large fraction of HARP borrowers were not required to verify their income or employment. Hence if these borrowers had been employed when they received their original mortgage, for the purposes of underwriting their HARP refinance their credit risk would still be evaluated based on those prior earnings. Second, even in cases where borrowers were required to provide documentation, they could either provide verification of their assets as an alternative to income or submit documentation for their unemployment insurance benefits. While the requirements for documenting assets rather than income were fairly strict, unemployed borrowers could use their unemployment benefits to document their income provided that the new mortgage payment did not exceed the rather high threshold of 45% of their monthly income. As a result, while access to HARP refinancing credit still depended on the borrower's employment status, the importance of documented employment was attenuated relative to a typical refinance.

The initial HARP program was announced in March 2009, and FHFA directed the GSEs to implement the program as soon as possible. FNMA began its DU Refi + program in March[50], while FHLMC rolled out its Relief Refinance mortgages in June. The initial program guidelines had very strict requirements for the guarantees lenders made regarding the quality of their loans (and hence high putback risks) and there was initially a 105% LTV ratio cap for eligibility. Due to unusually low program takeup, at the end of 2009 the LTV cap was raised to 125%. Nonetheless by September 2011 fewer

[48] Tracy and Wright (2012) [185] estimate an average payment reduction of 26% for HARP refinances.

[49] Agarwal et al (2015) [12] estimate that 54% of HARP refinances were via the original lender.

[50] HARP loans issued by FNMA from March through May 2009 functioned somewhat differently from other HARP loans.

than 1 million HARP refinances had been completed, well short of initial estimates of 4-5 million. As a result the FHFA changed the program guidelines dramatically, creating HARP 2.0 in January 2012. This iteration of the program removed all LTV ratio caps, waived certain fees and underwriting requirements[51], and severely reduced the strength of guarantees and putback risks facing HARP lenders. Following the implementation of HARP 2.0, takeup under the program increased dramatically, and more than 2 million borrowers refinanced through HARP in 2012. Following several extensions the program is currently scheduled to end at the end of 2016.

While my dissertation focuses primarily on the GSE policies outlined above, certain other mortgage-market trends and policy interventions are crucial to understanding the context of these policies. Between mid-2007 and late-2008, mounting losses on private-label mortgage securities caused considerable financial-market turmoil. In response, the Federal Reserve intervened on several occasions between March 2008 and November 2008 to either backstop financial institutions or facilitate their sale or liquidation[52], and the U.S. Department of the Treasury developed the Troubled Asset Relief Program (TARP) in October 2008[53] to purchase distressed securities from financial institutions in order to lean against fire-sale dynamics and put more cash on their balance sheets. Over the same time period, the Federal Reserve also lowered its Federal Funds Rate target in successive increments, finally reaching the zero-lower-bound for short-term interest rates in December 2008. In order to provide additional accommodation by lowering longer-term interest rates, the Federal Reserve in November announced a program of asset purchases referred to as quantitative easing (QE), as shown in Figure 1.4. The initial program aimed only to purchase treasury securities, but in March 2009 the Federal Reserve amended their mandate and began purchasing agency MBS as well. Evidence suggests that these purchases had a

[51] In particular, property appraisal requirements.
[52] These interventions, known as the Maiden Lane transactions, are not shown in Figure 1.4.
[53] Also not pictured in Figure 1.4.

substantial effect on MBS prices[54], leading to a reduction in secondary-market MBS yields and subsequently conforming mortgage interest rates, which fell from a peak of 6.48% in August 2008 to 4.23% shortly after the conclusion of QE in October 2010. As shown in Figure 1.4, the Federal Reserve subsequently implemented two additional rounds of MBS purchases (QE2 and QE3)[55], in each case contributing to substantial declines in mortgage interest rates. Conforming mortgage rates reached their lowest point ever in December 2012, shortly after the announcement of QE3, at 3.35%; as a result of the combined effects of the HARP program and low mortgage interest rates, refinancing activity ticked upward in 2012 and 2013.

1.3 Related Literature

My dissertation relates to many different strands of the existing economic literature, and some discussion of each is warranted. First, Chapters 3 and 4 are primarily an empirical analysis of individual default and refinancing decisions, and as such these chapters are closely related to a rich literature on household financial decision-making. Second, Chapters 2 and 4 provide an analysis of the effects of government policy on mortgage credit supply, and hence are related to a series of papers which analyze the determinants of credit supply, and in particular the effects of secondary market conditions on the primary mortgage market. Third, my research is connected to a number of papers which analyze the effects of credit constraints, both at the individual and macroeconomic level. Fourth, as an analysis of the effects of GSE policies, this paper is also related to a series of prior research papers analyzing the effects of GSE policy and policy pieces discussing options for housing finance reform. Finally, the analysis in Chapter 4 in particular is related to several other strands of the economic literature, including papers analyzing the use of credit supply policies as proxies for traditional

[54] See Cahill et al (2013) [47] and Hancock and Passmore (2014) [109] for evidence of these effects.
[55] As well as a similar program known as Operation Twist in September 2011, not pictured.

stabilization policy, papers related to insurance cycles in other contexts, and the broader empirical industrial organization literature concerning the estimation of dynamic models. In this section, I review the relevant literature in each of these areas.

1.3.1 Household Financial Decision-Making

At its core Chapter 4 is a structural model of borrowers, and in particular how they make choices and the costs of constraints on their choice set. This model builds upon several pre-existing structural models of borrower default and prepayment behavior. As a starting point, Hurst and Stafford (2004) [120] estimate a structural model of refinancing and emphasize the consumption-smoothing motive as a rationale for refinancing in high-interest-rate environments[56]. Bajari et al (2013) [22] and Ma (2014) [154] study refinancing and default behavior with a single-agent dynamic discrete choice model. Chen et al (2014) [53] and Laufer (2013) [145] enrich this type of model with credit constraints, continuous borrowing decisions and idiosyncratic income risk, although they estimate the model by matching aggregate moments rather than individual choices as in Bajari et al (2013) [22] and Ma (2014). [154] Hembre (2014) [113] similarly focuses on policy simulation by estimating a structural model of default behavior similar to that of Laufer (2013) [145] and using the model to evaluate the Home Affordable Refinance Program. My model builds on this prior work by more fully developing the structural aspects of endogenous credit constraints and latent employment status, allowing me to conduct additional counterfactual analyses on these margins, while continuing to leverage the richness of individual-level data.

While the papers cited above model borrower behavior as forward-looking and rational, prior evidence suggests that this is at best an approximation. Agarwal et al (2013) [7] show that borrowers fail to refinance optimally, leaving considerable money on

[56] While direct evidence of a consumption-smoothing motive for household refinancing is scarce, Sullivan (2008) [183] and Bethune (2014) [35] provide some evidence in other credit markets.

the table, and in particular are too insensitive to changes in interest rates. This finding that borrowers are inattentive or inertial in their decisions confirms work by Andersen et al (2014) [20] and Lee and Ko (2011) [146] showing that borrower refinance decisions are subject to a variety of choice frictions. However, even when borrowers are explicitly made aware of the benefits of refinancing, they still do not tend to refinance optimally, as demonstrated via an experiment in Johnson et al (2015). [128] This sub-optimal may be due to the fact, as suggested by other studies, that borrowers may not even understand either their current mortgage product or the set of available products. Bucks and Pence (2008) [46] show using discrepancies between survey and administrative data that borrowers do not know their own mortgage terms, while Agarwal et al (2013) [8] show that borrowers typically have an inaccurate conception of their own credit risk, leading them to in many cases accept higher-priced loan offers. Troublingly, Lacko and Pappalardo (2010) [142] report that simply providing borrowers with more information does not seem to improve their ability to procure better mortgage loans. This may be because, as noted in Woodward and Hall (2012), [188] borrowers do not shop around for mortgage loans very intensively, or, as noted in Disney and Gathergood (2011) [67] and Duca and Kumar (2014), [69] borrowers with low financial or mathematical acumen are simply incapable of navigating the market effectively. While my model is largely agnostic with respect to most of these issues[57], I do incorporate some insights from this literature in the structural model I present in Chapter 4.

There is also considerable disagreement regarding the importance of unemployment status for borrower refinance and default decisions, especially vis a vis their home equity. Using a hazard-rate analysis, An et al (2010) [18] show that borrowers default more often and refinance less often when unemployment rates in their area are high or they have low home equity. Other papers explore these relationships in more detail. Elul

[57] The model presented in Section 4.2 features no definition of "optimal" outside of that suggested by observed choices, hence borrower choices are tautologically optimal. Moreover the extent to which choices are sensitive to particular key variables is not imposed but rather left as a function of parameters estimated from borrower behavior.

et al (2010)[77] suggest that both unemployment and negative equity are necessary preconditions for default, and that negative equity alone is not enough, while Bhutta et al (2010)[36] argue that borrowers with extremely negative home equity may default even while liquid for strategic reasons. In contrast Gerardi et al (2013, 2015)[92,93] suggest in a series of papers that purely negative-equity-driven default is quite rare and that individual employment status is the best single predictor of mortgage default, irrespective of the borrower's equity position. Quercia et al (2012)[174] document that these effects vary by income: high-income borrowers default strategically and their prepayment decisions weight highly on the refinancing option value, while low-income borrowers tend to default for illiquidity-driven reasons and are actually more likely to refinance when unemployment rates are elevated. While these findings seem to confirm the consumption-smoothing model of Hurst and Stafford (2004),[120] in general the relationship between unemployment and refinancing is complicated by the presence of credit constraints. Indeed, Bethune (2014)[35] finds that unemployed borrowers are both more likely to apply for credit and more likely to be denied, so even if a consumption-smoothing motive is active the covariance between unemployment and refinance may still be negative. While I do not take a strong stand on these questions, the structural model I present in Chapter 4 accommodates both strategic and illiquidity-driven default and acknowledges that borrowers who fail to refinance may be credit constrained.

Much of Chapter 3 is intended to provide evidence in support of the microfoundations underlying many models of household financial decision-making, in particular those related to home-equity borrowing as a consumption-smoothing tool. This literature traces its roots back to Aiyagari (1994),[13] who models an economy composed of agents facing idiosyncratic labor income risk and strict borrowing constraints. The Aiyagari (1994)[13] model assumes that agent's preferences are such that the shadow value of easing credit constraints is highest for borrowers in low-income states. This very reasonable assumption also appears in several subsequent papers in the computable general equilibrium tradition, such as Krusell and Smith (1998),[141] Gourinchas and Parker (2002),[99] and Lustig and Van Nieuwerburgh (2005),[152] often via the same form of the agent's utility function and borrowing constraint. Moreover many empirical structural models of household borrowing behavior, such as Hurst and Stafford (2004),[120] Chen et al (2014),[53] and Laufer (2013),[145] parameterize agent preferences analogously and use the data to estimate only the parameters governing the labor income process. The analysis I present in Chapter 3 relates closely to these prior studies by providing empirical support for the assumptions underlying these models.

While direct evidence on use of household borrowing as a consumption-smoothing tool during periods of unemployment is relatively scarce, some previous studies do shed light on this issue. Sullivan (2008)[183] shows using U.S. micro-data that most borrowers, with the exception of the very wealthy and those with no assets, tend to run up their credit card debt when they become unemployed. He interprets this as evidence that borrowers with no assets would like to borrow when unemployed but lack access to credit. Dynarski et al (1997)[72] uses many of the same data sources to study how households smooth labor income shocks. They find that during unemployment spells, households tend to lean more heavily on government transfers and less on dissaving to make up for lost income, with the largest effects on low-income households, and interpret these findings as suggestive of binding liquidity constraints. Crossley and Low (2014)[62] use Canadian survey data to study the consumption response of workers facing unemployment. They find that only a small fraction of households are totally liquidity constrained in that they cannot borrow at all, but that consumption declines most sharply for this subset. Gropp et al (2014)[102] address the question of whether the household delveraging observed since the 2008 U.S. recession is due to demand (households wanted less credit) or supply (lenders tightened credit availability). They find that in comparison to homeowners, renters deleveraged relatively more in counties with larger home price declines, suggesting that credit availability tightened in those areas, or by implication, that borrowers in those areas did not actually wish to delever. Each of these studies attempts to infer the presence of credit constraints from household behavior, under the assumption that save for differences in credit access, different types of households should behave similarly with respect to their borrowing. By contrast, Bethune (2014) [35] directly observes credit constraints in the form of applications for unsecured loans. He finds that unemployed borrowers both apply for credit more often (or have greater demand) and are denied more often (or face tighter supply), with the supply effect outweighing the demand effect. The empirical analysis I present in Chapter 3 complements this literature by using several novel techniques to control for the effect of credit supply in order to identify the effect of unemployment on demand for refinancing credit.

Finally, a substantial prior literature relates home-equity borrowing directly to Aiyagari (1994) [13]-type models of precautionary saving by providing evidence that households do in fact build up a buffer stock of home equity to draw upon when facing income shocks. Caroll et al (2003) [50] use U.S. micro-data to show that for all save the lowest-income households, those facing higher ex-ante unemployment risk tend to build up more home equity but not other types of wealth. Benito (2009) [31] studies the mortgage equity withdrawal decisions of U.K. borrowers using survey data and shows that they are more likely to borrow when facing a negative income shock. Aggregate evidence, such as that in Lustig and Van Nieuwerburgh (2010) [153] and Zhou and Carroll (2012), [191] also supports the proposition that households tend to self-insure via housing wealth. Other studies confirm that borrowers tend to borrow against their home equity to support consumption. Cooper (2013) [59] shows using U.S. micro-data that borrowing-constrained households have a higher marginal propensity to borrow out of rising home equity and interprets this evidence that housing wealth influences consumption through collateral channels. Similarly Gan (2010) [89] shows that in Hong

Kong, households with positive equity shocks tend to consume more, but especially if they refinance. However, credit constraints can and do interfere with home-equity borrowing as a consumption smoothing tool, as demonstrated in Bhutta and Keys (2014),[38] who show that the probability of equity extraction is negatively related to local unemployment rates. My analysis in Chapter 3 contributes to this literature by providing empirical support for another aspect of Aiyagari (1994)[13]-type models of idiosyncratic labor income risk.

1.3.2 Determinants of Credit Supply

My dissertation and Chapter 2 in particular are also related to a sizable literature on mortgage credit supply, and in particular how secondary market conditions affect primary market credit supply. Several recent papers have studied the effect of upstream conditions on bank lending. Jiménez et al (2011)[127] show that the ability of banks to sell securitized mortgage loans on the secondary market leads to greater mortgage credit supply. In two similar papers, Loutskina and Strahan (2009)[149] and Loutskina (2011)[151] study the effect of secondary market conditions on the relationship between bank's liquid funds and lending volumes. The former uses exogenous variation in ease of securitization derived from GSE conforming loan standards to show that banks with more abundant or cheaper deposit funding supply relatively more credit to less liquid nonconforming loans, while the latter uses a model-constructed index of "securitizability" to show similar effects across multiple loan categories. Taken together, these three papers suggest that for highly liquid loans, such as conforming mortgages, upstream secondary market conditions are at least as important a determinant of credit supply as the lender's own financial status, a result confirmed by Gabriel and Rosenthal (2007).[86] This study corroborates many of these findings by analyzing the extent of pass-through from guarantee fees to MBS prices and subsequently mortgage interest rates. Other recent papers examine aspects of supply and demand for MBS. Merrill et al

(2014)[157] study the factors driving upstream demand for loans and find that their favorable treatment for regulatory capital requirements was an important determinant. Fuster et al (2013)[84] quantify the factors determining the spread between the the rate that borrowers pay (the primary mortgage rate) and the coupon rate that MBS investors receive (the secondary mortgage rate), finding that while some of the increase is attributable to guarantee fee increases, most is unexplainable. In studying the effect of guarantee fees on credit supply, the analysis I present in Chapter 2 is most-closely connected to these studies of the interaction between secondary and primary market credit conditions.

Other recent papers also examine the strategic incentives facing mortgage lenders. In a pair of related studies, Keys et al (2010)[133] and Keys et al (2012)[134] examine how ease of securitization affects the intensity with which originators screen borrower credit risk, and find that while lenders do exert less effort in screening borrowers whose mortgages will be easier to sell, the effect is concentrated in subprime mortgage lending. Stanton et al (2014)[181] document high levels of local-market concentration in mortgage credit supply as well as coordination between upstream and downstream participants. Scharfstein and Sunderam (2013)[178] and Allen et al (2013, 2014)[14,15] study the effects of such concentration in mortgage lending on credit supply. The former find that measures of concentration are associated with reduced pass-through from lender costs to mortgage rates, consistent with oligopoly pricing, while the latter find that borrowers and high-risk borrowers in particular receive less favorable mortgage terms in more concentrated markets. Similar evidence for the effects of concentration on pricing in mortgage markets is presented in Amromin and Kearns (2014).[17] My results in Chapter 2 support these findings by documenting how guarantee fees reduce lender incentives to originate new mortgages and quantifying how the extent of pass-through to interest rates varies by levels of market concentration and borrowers risk.

Chapter 4 develops a model of mortgage credit supply that incorporates many of the insights from these and other papers. The models of credit supply I present in Sections 1.4 and 4.2 are based on the work of Einav et al (2012),[74] who study subprime auto loans, and Crawford et al (2014),[61] who study small-business lending, with three key adaptations to leverage the data I have available and more closely match key institutional features of the agency mortgage market. First, lenders face a choice in my model between holding loans and securitizing them via the GSEs, and as documented in Agarwal et al (2012),[5] Downing et al (2009),[68] and Keys et al (2012),[134] in equilibrium they will adversely select their most-risky loans for securitization. Second, as discussed in Allen et al (2013, 2014),[14,15] Amromin and Kearns (2014),[17] Scharfstein and Sunderam (2013),[178] and Gurun et al (2013),[105] local-market competitiveness partially determines the strength of lender offers in the model, with higher-priced offers in less competitive markets. Those papers all note that mortgage lending appears to follow a model of monopolistic competition, and the evidence from Amromin and Kearns (2014)[17] and Gurun et al (2013)[105] as well as that mentioned previously suggest that sub-optimal consumer choice may play a role in generating these market inefficiencies. Finally, because I do not observe interest rate offers for loan applications that are denied, the estimate procedure I describe in Section 4.3 employs a selection-correction procedure similar to that in Jiménez et al (2012, 2014)[125,126] in order to estimate the credit supply function.

1.3.3 Effects of Credit Constraints

In its own way, each chapter of my dissertation analyzes the effects of credit constraints on household financial decisions, and as such each is closely connected to several recent papers studying similar household-level effects. Laderman (2012)[143] and Krainer and Laderman (2011)[139] demonstrate that tight credit supply conditions that prevent borrowers from refinancing lead to higher rates of mortgage default. Similarly, Keys et al (2014)[136] and Tracy and Wright (2012)[185] show that borrowers making lower monthly payments, as may result from refinancing at low interest rates, are less likely to subsequently default. These results provide some support for an option-theoretic model of the type estimated in Bajari et al (2013) [22] in which borrowers on the margin choose between refinance and default. As such, we may expect reductions in credit availability to lead to higher levels of mortgage default. The results I show in Chapters 2 and 4 provide detailed evidence in support of these effects by connecting changes in credit supply conditions to refinancing activity and subsequently to default rates.

A closely-related literature uses exogenous variation induced by credit supply policies to identify household responses to changes in credit supply. The gold standard for such variation is the randomized control trial, and Karlan and Zinman (2002) [131] provide evidence for variation in credit demand using such an empirical design. Other studies rely more heavily on natural experiments. Gross and Souleles (2002) [103] show that borrowers tend to spend more when their credit card debt limits increase, particularly if they were close to their previous limit. Hochguertel et al (2005) [116] exploit changes in usury laws to estimate credit demand elasticity with respect to interest rates, and find that demand is less price-sensitive in poorer regions. Leth-Petersen (2010) [148] studies borrowing following a change in Danish law that permitted households to borrow against home equity. He finds that those with low levels of liquid assets tend to borrow more, with the exception that those who are unemployed tend to borrow less. Abdallah and Lastrapes (2012) [1] study a similar legal change in Texas and find that, consistent with binding liquidity constraints, measures of consumption (retail sales) increase relatively more in lower-income counties when home-equity borrowing is permitted. Agarwal and Qian (2014) [10] demonstrate the converse, showing with data from Singapore that borrowers reduce their debt levels when home equity borrowing is prohibited. Finallly, Markwardt et al (2014) [155] use the Danish policy reform to show that home-equity borrowing and unemployment insurance are substitutes, suggesting that both are used for consumption-smoothing purposes. The results I provide in

Chapters 2 and 3 in particular contribute to this literature by studying household-level responses to a pair of GSE policies that materially affected credit availability.

At the macroeconomic level, there is also a substantial literature on the effects of loose or tight household credit constraints to which my dissertation is related. Tenreyo and Thwaites (2013)[184] and Mora (2014)[167] find that credit constraints can interfere with monetary transmission to consumption, a result that U.S. central bankers such as Mishkin (2007)[165] and Bernanke (2007)[34] have attributed to mortgage refinancing frictions. As a result, if as noted in Jappelli and Pistaferri (2014)[123] and Carroll et al (2014)[51] and modeled in Aiyagari (1994),[13] the marginal propensity to spend is highest for income-constrained households, then easing credit constraints countercyclically can lead to significant stimulus. Eggertson and Krugman (2012),[73] Midrigan and Philippon (2011),[173] and Guerrieri and Lorenzoni (2011)[104] study theoretically the negative impact of tightened household borrowing constraints, while Mian et al (2013)[161] and Mian and Sufi (2012, 2014)[160,162] provide evidence in support of these theories. Di Maggio et al (2014)[65] and Keys et al (2014)[136] confirm many these findings at the individual level, showing that lower monthly mortgage payments (as might result from refinancing) lead to increased household consumption spending. Other papers consider the effects of loose credit constraints. Jeske et al (2013)[124] find that the mortgage subsidy provided by the GSEs leads to above-optimal levels of borrowing, while Khandani et al (2013)[137] and Mian and Sufi (2009, 2011)[158,159] show that loosening household credit constraints can lead to an over-accumulation of debt, leading to higher subsequent default rates. The results I present in Chapter 4 complement this literature by analyzing the effects of a policy that tightens credit constraints during booms and loosens them during busts.

1.3.4 Effects of Government-Sponsored Enterprise Policies

As an analysis of GSE policy, my dissertation is closely related to the literature studying the impact of GSE policies, especially with regards to credit supply. DeFusco and Paciorek (2014) [64] estimate the elasticity of mortgage demand with respect to interest rates and use the model to simulate the effect of guarantee fee increases, assuming a given level of pass-through to interest rates. The results from Chapter 2 build on their work by estimating the extent of pass-through using loan-level data, and my results are broadly consistent with theirs. Ambrose et al (2004) [16] and Kaufman (2014) [132] estimate the effect on mortgage interest rates of GSE securitization, and find that GSE-eligible loans enjoy interest rates roughly 10 to 25 bp below otherwise similar mortgages. Passmore et al (2005) [170] invert this question by studying the determinants of the conforming loan spread, or difference in interest rates between conforming and non-conforming loans, and estimate that just less than half of the spread is derived from the funding advantage the GSEs have over private lenders due to their implicit government backing. My paper follows these studies in examining how exogenous changes in guarantee fees affect the conforming loan spread. Bostic and Gabriel (2006) [42] and Lehnert et al (2008) [147] both study the effect of GSE portfolio purchases on credit supply and both find that they have no measurable positive impact on credit provision. Gabriel and Rosenthal (2010) [87] investigate the extent to which GSE securitization crowds out private mortgage credit and find that the extent of crowd-out varies over the business cycle. In particular, during expansionary period such as from 2003 to 2006 there is significant crowd-out, although the GSEs still increase mortgage credit supply; by contrast during the recession from 2007 to 2009 there was essentially zero crowd-out as non-agency mortgage credit dried up. On a related note, Peek and Wilcox [171] demonstrate that historically the GSEs served to reduce intertemporal variation in mortgage credit supply by increasing their market presence during periods in which private credit supply tightened. More recently, Fuster and Vickery (2013) [85] show that following the breakdown of the non-agency mortgage market in 2008, FRM origination became rarer relative to ARM origination, suggesting that securitization is virtually a requirement to manage the inherent risks of FRM origination. Hurst et al (2014) [121] quantify the redistributive impact of the GSEs' policy of not varying guarantee fees across U.S. regions and find that the effects are quite large in comparison with other redistributive programs. My results contribute directly to this literature by analyzing the effects of changes in GSE policy on credit availability, both for observed policies in Chapters 2 and 3 and counterfactual policies in Chapter 4.

Moreover, following the GSE's entry into government conservatorship in September 2008, a number of policy papers have proposed reforms to the US housing finance system. Proposals have tended to advocate for either a largely private system in which insurance premiums are set on the open market or a public system retaining many of the features of the GSEs are currently constituted. Jeske et al (2013), [124] Acharya et al (2011), [3] Scharfstein and Sunderam (2011) [177] and Elenev et al (2015) [75] fall into the former camp, while Mosser et al (2013) [168] and Dechario et al (2011) [63] expand on the privatization proposals by advocating for a cooperative model. Hancock and Passmore (2011) [108] and Goodman (2014) [97] note that any privatized system would likely be subject to the same destabilizing dynamics witnessed with private-label securitization, and hence, at a minimum, the government would be needed as a catastrophe re-insurer. Goodman et al (2014) [96] note that pricing such catastrophe insurance is a difficult problem, and Jaffee and Quigley (2011) [122] discuss some of the political pressures to under-price such insurance. Smith and Weiher (2012) [180] and Dynan and Gayer (2011) [71] discuss the need for countercyclical GSE policy, but focus respectively on capital requirements and market shares[58]. Frame et al (2015) [83] analyze how GSE policy relates to the tightening of mortgage credit supply post-recession, and while they do

[58] Although they do not discuss this issue, these could potentially be achieved with countercylical guarantee fees.

discuss the conflict between their increased profitability and macroeconomic policy goals they focus primarily on the GSE's efforts to better control lender fraud[59]. My dissertation contributes to this policy debate by proposing and analyzing the effects of a change to GSE policy. However, to my knowledge, none of the proposals mentioned above, nor any of the private proposals surveyed in Ellen et al (2010) [76] discusses the effects of a countercylical guarantee fee policy.

Finally, several prior papers have analyzed the effects of other GSE policies at both the household and aggregate levels. Zhu (2012) [192] studies the future default behavior of HARP-eligible borrowers using administrative data. After applying a selection correction to control for program takeup, he finds that eligible borrowers who refinanced under HARP are about 50% as likely to subsequently default, indicating that the program achieved its objective. Remy et al (2011) [176] perform a cost-benefit analysis of a simulated HARP-type program and find that relatively few defaults, roughly 2.8% of the total, would actually be averted, and that as a result the program would actually cost the government a small amount of money on net. Tracy and Wright (2012) [185] estimate a model of mortgage termination using loan-level data and simulate the effects of a payment reduction equivalent to what borrowers would receive from a HARP refinance. Their estimates suggest a reduction in default on the order of 25%, or only half as large as what Zhu (2012) [192] estimates actually occurred. Amromim and Kearns (2014) [17] show that borrowers who refinanced via HARP tended not to actively shop for better deals with different mortgage lenders, in part because the design of the program itself incentivized borrowers to refinance with their original lender. They find that lenders responded to this behavior by pricing against their captive market, leading HARP refinancers to receive interest rates .15-.2% higher than comparable non-HARP borrowers. At the macro-level, Agarwal et al (2015) [12] examine both the direct and

[59]Specifically, in the wake of the recession the GSEs strengthened the representations and warranties required of mortgage lenders, allowing them to put-back defaulted loans, or to effectively cancel their guarantee agreement, if the loan was fraudulent or if the representations were otherwise violated.

indirect effects of HARP and find that borrowers tended to benefit from refinances under the program and that the increase in refinancing stimulated household consumption. The results I show in Chapter 3 complement this literature by analyzing the factors predicting take-up under HARP.

1.3.5 Other Related Literature

Several recent papers have analyzed, as my dissertation does, how credit supply policies can proxy for elements of traditional stabilization policy. Hurst et al (2014) [121] study the redistributive impact of a guarantee fee policy that is constant across space, and find that by transferring resources towards default-prone areas, this policy provides aggregate insurance comparable to the effects of fiscal stimulus. Similarly, Kimball (2012) [138] argues that providing access to credit during recessions is likely to be a cost-effective form of stabilizing transfers. Regarding home lending specifically, Krainer and Laderman (2011) [139] and Tracy and Wright (2012) [185] argue that easier refinancing credit conditions can help prevent foreclosure, while Markwardt et al (2014) [155] shows that home-equity borrowing can be a substitute to traditional unemployment insurance. Hurst and Stafford (2004), [120] as mentioned above, explain this effect through consumption-smoothing refinance, and Carroll et al (2003) [50] among others have shown that moderate-income households tend to insure against unemployment through home-equity-based precautionary saving. Kroft and Notowidigdo (2011) [140] and Landais et al (2014) [144] study the optimal policy for a more traditional stabilizing tool, unemployment insurance, and find that the optimal policy varies with the business cycle. Chapter 4 contributes to this literature by analyzing how to use credit supply policy to lean against the business cycle.

In analyzing optimal insurance pricing policies for the GSEs, the analysis I present in Chapter 4 is also connected to several papers that study insurance cycles in other settings. As discussed originally in Winter (1991), [187] insurance cycles refer to the tendency of insurance premiums to follow a catastrophe dynamic of steeply rising premiums following periods of large losses. In many cases, including arguable the case of the GSEs themselves, this catastrophe dynamic is regulation-driven as insurers seek to recapitalize. Abruzzo and Park (2014)[2] and Murphy et al (2014)[169] study the destabilizing effects of procylical futures exchange margins, or cases in which low losses lead to low margins and hence excessive borrowing and vice versa. The issues surrounding futures margin cyclicality are in many respects directly analogous to GSE guarantee fees, in that low premiums during boom times lead to increased leverage while increases in premiums following crises choke off credit supply. Similar dynamics occur in capital regulation, as Gordy and Howells (2014)[98] and Balla and McKenna (2009)[23] discuss with regards to, respectively, the Basel II capital requirements and loan-loss provisioning. Despite the presence of a deep and forward-looking public backstop and the absence of regulatory pressures, most government programs also face insurance cycles. Pennacchi (2009)[172] note that the 2005 Federal Deposit Insurance Reform Act literally requires the FDIC to set procyclical insurance premiums, and Michel-Kerjan and Kunreuther (2011)[163] discuss the disastrous performs of the National Flood Insurance Program. Moreover as discussed in Hill (2010),[114] the same boom-bust-bailout dynamic observed with FHLMC and FNMA actually occurred previously with a smaller GSE, Farmer Mac, in the mid 1980s. Farmer Mac, as with its larger cousins, pursued a procyclical lending policy, despite the fact that the presence of the bailouts themselves suggest that policymakers actively desire the GSEs to lend countercyclically. In calling for through-the-cycle pricing on behalf of FNMA and FHLMC, Hancock and Passmore (2014)[110] note that the FHA attempts to set its mortgage insurance premiums across the cycle, although in fact much as the GSEs did the FHA raises its premiums six times between 2008 and 2013.[100] As such, any countercyclicality in market shares for either the FHA or the GSEs is in spite of their pricing policies, as both were subject to the same insurance cycle dynamics.

Finally, the model I develop and estimate in Chapter 4 follows in the tradition of structural industrial organization and as such is related to several recent papers in this field. My estimation technique follows the semi-parametric approach of Hotz and Miller (1993)[117] as applied to mortgage markets by Bajari et al (2013),[22] with first-stage policy functions estimated using the methodology from Barwick and Pathak (2011)[26] and Bazdresch et al (2014).[29] To this general framework I add three key ingredients drawn from other areas of industrial organization. First, my model includes both discrete and continuous choices, and for identification I rely on the results from Blevins (2010).[40] Second, the model features persistent unobserved heterogeneity. The technique for estimating the process for the latent unobservables is drawn from Diebold et al (1999),[66] with a modification from Chung et al (2004)[55] introduced in order to address the label-switching problem, while identification relies on the results from Connault (2014).[58] Third, the model incorporates endogenous credit constraints; the framework for modeling these constraints is drawn from Einav et al (2012)[74] and Crawford et al (2014),[61] while the multi-stage estimation procedure is adapted from Ho (2009).[115]

1.4 A Model of the Agency Refinancing Process

In this section I outline a model of the agency refinancing market featuring both borrower demand for credit and lender supply decisions. Lenders in the model offer contracts and set contract terms to maximize their profit, as in Einav et al (2012),[74] while borrowers solve a dynamic value-maximization problem in making their refinance decisions as in Bajari et al (2013).[22] The general framework follows these earlier papers, although several features are adapted to more closely match the institutional details of the agency refinancing market. In several ways this model provides context for the remainder of the paper. First, the choice variables in the model, both discrete and continuous, correspond to the data I will describe in Section 1.5, and the model

Figure 1.5: Extensive Form of Model

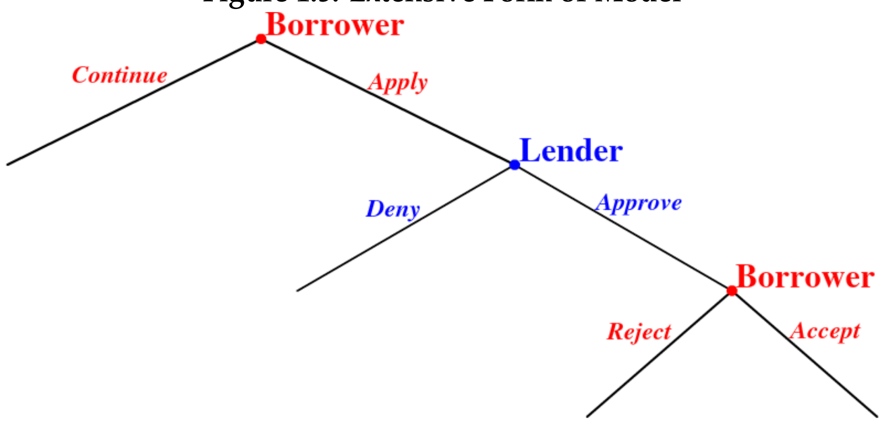

Notes: Timing of actions in stage game. Red/blue nodes denote borrower/lender decisions.

demonstrates their importance to understanding the mortgage refinancing process. Second, the model generates two key predictions regarding the effect of guarantee fees on credit supply and the effect of liquidity preference on applications and denials that serve as a framework for interpreting the results I will present in Chapter 3. Finally, the model forms the basis for the stage game played between borrowers and lenders in the full structural model I develop in Chapter 4.

The extensive form of the stage game played by borrowers and lenders is shown in Figure 1.5. Borrowers initially choose whether to submit a refinancing application, and if so, they choose how much to borrow. If the borrower submits an application, she pays a known fixed cost τ_B, and the lender then decides whether to approve the application, and if so, what interest rate to offer. Borrowers then have an opportunity to reject high-interest-rate loan offers. If the borrower accepts the offer, the lender receives a payoff based on the interest rate on the loan, the LTV ratio, and prevailing guarantee fees, while the borrower obtains value V^R from refinancing as a function of the interest rate and LTV ratio. If instead the borrower either chooses not to apply, is denied, or rejects a credit offer, she continues with her current mortgage and obtains a value V^C[60]. At the conclusion of the stage game, borrowers move on to the next period and play the

[60]Consistent with the evidence from Woodward and Hall (2012),[188] I assume that borrowers do not submit multiple applications.

stage game again next month.

Borrower i decides to Apply (A) based on a private-information type $(\eta_{i,t}, V^C_{i,t})$ which indexes, respectively, her preference for liquidity (cash) and her continuation value[61]. The former can be thought of as measuring the slope of the utility function in the borrower's current state[62], while the latter encompasses a set of beliefs about the future. Lender j decides whether to Deny (D) applicants based on a private-information type k_j, which measures funding costs, as well observed fixed underwriting costs τ_L and guarantee fees p_t. Borrower i chooses a LTV ratio $c_{i,t} \in [0, \bar{c}]$ when submitting her application, which in turn determines the amount borrowed, while lenders approving applicants extend interest rate offers $r_{i,j,t}$. Borrowers internalize the risk of being denied in their choices, while lenders internalize the risk that their high-interest-rate offers may be Rejected (R) by the borrower. I assume that lenders expect a distribution of borrower types $(\eta_{i,t}, V^C_{i,t}) \sim F(\eta, V^C)$, and likewise that borrowers expect a distribution of lender types $k_j \sim G(k)$ each month.

I solve the model via backwards induction. Equilibrium is described as a set of policies for the borrowers rejection decision (R), application decision (A), and LTV choice (c), and the lender's denial decision (D) and interest-rate offer (r) such that borrowers maximize their expected value and lenders maximize their profit. If the borrower has received a credit offer for a loan with LTV ratio $c_{i,t}$ and interest rate $r_{i,j,t}$, the borrower's rejection policy is given by:

$$R(c_{i,t}, r_{i,j,t}, \eta_{i,t}, V^C_{i,t}) = \mathbf{1}[V^C_{i,t} \geq V^R(c_{i,t}, r_{i,j,t}, \eta_{i,t})]$$

where V^R denotes, analogously to V^C, the value of refinancing. By inverting V^R, this policy can effectively be described as a threshold rule in the offered interest rate, above which the borrower rejects all offers. Given this policy, lenders offer interest rates to

[61]For ease of exposition, I suppress any other observable characteristics, such as the borrower's credit score.

[62]Essentially, the marginal value of cash-on-hand in an Aiyagari [13]-type model.

maximize their expected profit[63]:

$$r^*(c_{i,t}, p_t, k_j) = \underset{r}{\operatorname{argmax}}\, EP(c_{i,t}, r, p_t, k_j)$$

where the expectation EP, taken over borrower types (η, V^C), is written explicitly in Appendix A. If the expected profit from the lender's optimal interest rate does not exceed her fixed underwriting costs, the lender denies the application, leading to denial policy:

$$P_{SEC}(c_{i,t}, p_t, k_j, \tau_L) = \mathbf{1}[EP(c_{i,t}, r^*(c_{i,t}, p_t, k_j), p_t, k_j) \leq \tau_L]$$

Borrowers choose their optimal LTV ratio to maximize their expected value, where the expectation is taken over lender types, which in turn determine the risk of denial and of high interest rate offers:

$$c^*(p_t, \eta_{i,t}, V^C_{i,t}, \tau_L) = \underset{c}{\operatorname{argmax}}\, EV(c, p_t, \eta_{i,t}, V^C_{i,t}, \tau_L)$$

Finally, the borrower chooses to submit an application only if the expected value of the application at the optimal interest rate less any fixed costs exceeds her continuation value:

$$A(p_t, \eta_{i,t}, V^C_{i,t}, \tau_L, \tau_B) = \mathbf{1}[EV(c^*(p_t, \eta_{i,t}, V^C_{i,t}, \tau_L), p_t, \eta_{i,t}, V^C_{i,t}, \tau_L) - \tau_B \geq V^C_{i,t}]$$

Else, the borrower continues. I assume that the relevant types remain private information in every period, in other words, that borrowers and lenders do not learn from one stage game to the next.

[63]In solving the lender's problem I ignore issues of incentive compatibility and assume that at the unconstrained optimal interest rate, borrowers are incentivized to select the LTV ratio corresponding to the optimum for their type. This assumption makes the model more tractable for the purposes of generating the necessary predictions, although selection into LTVs by borrower type will be explicitly accounted for in the full structural model in Section 4.3.

Under certain assumptions, enumerated in Appendix A, this game features a unique equilibrium $(A^*, c^*, D^*, r^*, R^*)$ for each combination of types $(\eta_{i,t}, V^C_{i,t}, k_j)$ and prevailing guarantee fees p_t. With several additional assumptions[64], one can also prove the following pair of propositions:

Proposition 1. *Interest rates and denial rates both increase with guarantee fees, or*

$$\frac{dr^*(c_{i,t}, p_t, k_j)}{dp_t} \geq 0 \text{ and } \frac{d\int_{-\infty}^{\infty} P_{SEC}(c^*(p_t, \eta_{i,t}, V^C_{i,t}, \tau), p_t, k, \tau)g(k)dk}{dp_t} \geq 0$$

Proposition 2. *Application rates and denial rates both increase with liquidity preference, or*

$$\frac{dEV(c^*(p_t, \eta_{i,t}, V^C_{i,t}, \tau), p_t, \eta_{i,t}, V^C_{i,t}, \tau)}{d\eta_{i,t}} \geq 0 \text{ and } \frac{d\int_{-\infty}^{\infty} P_{SEC}(c^*(p_t, \eta_{i,t}, V^C_{i,t}, \tau), p_t, k, \tau)g(k)dk}{d\eta_{i,t}} \geq 0$$

These propositions are closely related to the analysis I will present in Chapters 2 and 3. The first proposition states that guarantee fees lead to tighter credit constraints, both by passing through to interest rates and by increasing the probability that a loan application is denied[65]. In this sense the model is consistent with the evidence that I will present in Chapter 2 regarding guarantee fees and credit supply. The second proposition provides a framework for interpreting the evidence on unemployment and the shadow value of credit constraints that I will present in Chapter 3. It states that as liquidity preference increases, borrowers apply more frequently and for larger loan balances[66], but as a result their applications are denied more frequently. As such the relationship between liquidity preference and actual observed refinancing is ambiguous, since actual refinancing is the product of A, D and R. In this sense the model also captures the intuition from Bethune (2014),[35] who shows that borrowers who become unemployed

[64] The theoretical results rest on "assumptions" per se, but in the full model these will mostly depend on the estimated values of certain parameters.

[65] The second term is the ex-ante expected probability that a loan application is denied.

[66] The first term is the expected value of submitting a refinancing application, and the mechanism by which it increases in η is through choice of loan balance.

are both more likely to apply for credit and more likely to be denied. Hence in interpreting the evidence on unemployment and refinancing, we must bear in mind that even if, as prior reserach suggests,[123] the unemployed have high liquidity preference, one may not observe additional refinancing without controlling for credit supply.

This model also forms the basis for the structural estimates to be presented in Chapter 4. The goal of my estimation procedure is to recover estimates for V^C, V^R, and π using data on outcomes D, R, c and r. In order to do so, I will need to place additional structure on the form of borrower's and lender's preferences. To provide motivation for that additional structure, in Chapter 3 I present some reduced-form evidence on borrower and lender behavior in the data. Combined with Proposition 2, these results will provide some suggestive evidence in favor of the claim that the shadow value of credit constraints is greater for unemployed households.

1.5 Data Sources

In this section I describe the key data sources used in the following three chapters. As the goal of my analysis is to trace the impact of GSE policies on credit supply and subsequently on borrowers, I require data on both credit supply and borrower outcomes. The stylized model from Section 1.4 suggests that at a minimum, I should be able to observe the outcomes of refinancing applications, the interest rates offered to refinancing borrowers, and data on when borrowers actually choose to refinance. I refer to the latter as "demand data" and the former two as "supply data", although as noted in Section 1.4 refinancing is a product of both borrower and lender choices. My primary source of borrower-level demand data is panel data on borrower refinancing and default decisions provided by FNMA and FHLMC. To estimate the structural model I develop in Chapter 4, as well as some regression specifications in Chapter 3, I merge this data with more detailed property-level data obtained from county recorder's offices. My primary source of loan-level supply data is data on refinancing applications from the Home Mortgage Disclosure Act (HMDA) Loan Application Register, supplemented with data on interest rates from the same FNMA and FHLMC borrower-level sources. For much of the analysis to be presented in Chapter 2, I also use data on the characteristics and prices of 30-year fixed-rate MBS issued via FNMA or FHLMC obtained from Bloomberg. I then supplement these primary datasets with a variety of other data sources in order to measure additional non-individualized variables and aggregate effects. In what follows, I describe each of these data sources in turn, as well as my process for cleaning, merging, and supplementing these datasets.

1.5.1 Borrower-Level Data Sources

I assemble the main borrower-level from the FHLMC Single-Family Loan-Level Dataset and FNMA Single-Family Loan Performance Dataset. These datasets contain detailed monthly loan-level panel records of all single-family first-lien 30-year FRMs sold or securitized via FHLMC or FNMA between 1999 and 2013, excluding mortgages originated through some government programs[67], with certain special features[68] or with insufficient documentation. The data is separated into a header file and a dynamic file, both of which I use for various purposes. For each loan, the header file provides detailed credit risk information, such as the LTV ratio and FICO score for the borrower, at origination. In each subsequent month, the dynamic file provides time-variant variables such as the current balance of the loan, as well as the borrower's choice of whether to make a payment, default, or prepay. I refer to the combined header-dynamic file as the "GSE dataset", and for each record I observe borrower decisions between the origination date and either the time at which the borrower defaults or prepays or December 2013, whichever comes first[69].

[67] Including HARP, VA or FHA loans.
[68] Such as Alt-A loans or loans with prepayment penalties.
[69] Additional details on the construction of this dataset are provided in Appendix B.

My sample consists of all loans contained in the GSE dataset originated nationwide between 1999 and 2013, although for various purposes I use subsets of the data. In evaluating the effect of guarantee fees on interest rates in Chapter 2, I use the header file only, as the crucial outcome is the interest rate obtained at origination, and restrict attention solely to refinance loans and not home-purchase loans. The dataset used to study this effect consists of 35,701,979 refinance mortgages originated between Q1 2000 and Q2 2013, and in estimating the effect of policy changes I restrict attention to the 21,435,660 refinance mortgages, and in particular the 2,077,099 refinance mortgages originated in a two-year period around the time of the two 2012 guarantee fee increases. To evaluate the effects of guarantee fees on the propensity of households to refinance or default in Chapter 2, I use the full header-dynamic dataset. To avoid conflating the effects of guarantee fee increases with the effects of HARP 2.0, which was implemented concurrently with the first guarantee fee increase, I restrict attention to a sample of loans that were ineligible for HARP, specifically those that were originated after June 2009. This leaves an unbalanced panel of 150,759,569 monthly servicing records for 6,973,525 mortgages observed between July 2009 and June 2013. These two sub-samples are described in more detail in Section 2.4, while summary statistics are provided in Table A1 and Figure A1 in Appendix C. In much of the analysis presented in Chapter 3, I focus on a different subset of loans originated in a 1-year window around the cutoff for HARP eligibility of June 2009, or loans originated from June 2008 to May 2010. This subset comprises 177,175,761 monthly observations of 4,664,219 distinct loans observed from origination through December 2013 and represents a substantial fraction of all loans that were eligible for HARP. For some specifications, I further restrict attention only to HARP-eligible loans during the period following the implementation of HARP 2.0, comprising some 24,063,716 loan-month-level observations for 1,975,070 distinct loans originated in the period just prior to the introduction of the HARP program. These two subsamples are described in more detail in Sections 3.3 and 3.4 and the composition of these subsamples by month of origination is provided in Table A2 in Appendix C. To estimate the structural model I outline in Chapter 4, as well as certain specifications in Chapter 3, I use an augmented version of this dataset comprising only loans originated in California[70], which I describe in more detail in Section 1.5.4. The full details regarding the construction of the baseline GSE dataset can be found in Appendix B.

1.5.2 Credit Supply Data Sources

The primary dataset I use to estimate credit supply is the HMDA Loan Application Register, a regulatory database containing a record of every mortgage loan application at almost all U.S. mortgage lenders[71]. I observe this data annually from 2000 to 2012. Each record identifies the lender who received the application, certain credit risk information such as the borrower's income and desired balance, and the location of the property securing the loan. Crucially, the data also contain information on the disposition of the loan: whether the application was denied, whether the borrower subsequently rejected a credit offer, and if the loan was originated, whether it was then sold to FNMA or FHLMC. I use this dataset in Chapters 3 and 4 to estimate whether and on what terms lenders will extend credit, and subsequently use certain of these estimates as inputs for the borrower-level structural model I develop in Chapter 4. In order to more closely match the scope of my borrower-level data, I restrict attention to what could be considered the agency market, as it is the average disposition of loans in this market which most-closely approximates the credit supply constraints facing borrowers

[70] I restrict attention to California when estimating my structural model for several reasons. First, mortgage market institutions, in particular rules regarding foreclosure, vary considerably across states, hence it is helpful to use observations only from a single state. Second, the process of matching observations to local-area house price indices and unemployment rates involved substantial manual input, as towns have to be matched by name and naming conventions differ across data providers, so it would be infeasible to repeat this process nationwide. In light of these two considerations, I chose the state which would be most-likely to provide sufficient intertemporal and geographic variation in house price growth and unemployment rates to identify the model; California, with a highly heterogeneous population and large exposure to both the house price boom and Great Recession, was the most-natural candidate.

[71] Very small lenders and lenders who are only active in rural areas are exempt from the regulation. Estimates place the coverage at roughly 85-90% of all mortgage applications in the U.S.

comprising the GSE dataset. To that end I exclude any loan applications that would not qualify for sale to the GSEs, including those with balances above the conforming loan limit and subprime loans[72], and because I estimate the structural model in Chapter 4 using only data on loans from California, I also restrict attention only to applications from California. Certain credit risk information, in particular FICO scores and LTV ratios, are not provided in the HMDA data; following Bayer et al (2013) [28] and Goodman and Li (2014), [95] I impute these data using the observed mean within a given year, lender, postal code, income decile, and origination balance decile in the GSE dataset. Summary statistics on the credit risk characteristics and outcomes for applications in this dataset are provided in Tables A4 and A5 respectively in Appendix C, while the details of the cleaning and imputation procedure are described in Appendix B. The structural model I estimate in Chapter 4 also requires data on interest rates at origination, and for these purposes I use the same dataset on refinance originations described in Section 1.5.1.

1.5.3 Secondary Market Data Sources

My analysis makes use of secondary market data on MBS in two key ways. First, in my analysis of the effects of guarantee fee increases on agency MBS in Chapter 2, I use data on the characteristics and prices of 30-year fixed-rate MBS issued via FNMA or FHLMC. Second, my structural estimates in Chapter 4 require estimates of two types of model primitives derived from secondary-market-level data, namely, a pricing model for MBS at origination as function of their average characteristics, and the elasticity of guarantee fees with respect to the MBS-level average LTV ratio. I assemble the main secondary market dataset I use for these two purposes by merging data from FHLMC's Historical Daily New Issues PC Reports and FNMA's PoolTalk database with either

[72] I do not directly observe whether a given loan is subprime, but rather infer that all loans made to lenders on the Department of Housing and Urban Development's subprime specialist lender list were subprime. See Gerardi et al (2008) [91] for a discussion of the accuracy of this imputation.

daily or weekly MBS price information obtained from Bloomberg. The FHLMC and FNMA data contain characteristics of issued MBS at the time of issuance, such as the mean, min, max and quartiles of the interest rate, FICO and LTV of the mortgages backing the MBS, as well as the coupon on the MBS, the origination date and the initial balance. In order to measure the effects of guarantee fee increases at high frequencies in Chapter 2, I collected data on all 30-year fixed-rate MBS issued via FNMA or FHLMC in a two-week window around each of the 2012 guarantee fee increase dates. I merge this data with daily-level mid prices from Bloomberg, resulting in a sample of 26,525 MBS-trading day observations for 2,696 MBS issued in a window around April 1st, 2012 and 32,913 MBS-trading day observations for 3,646 MBS issued in a window around December 1st, 2012. To measure the effects of guarantee fee increases at longer horizons in Chapter 2, I also collected data on all FNMA 30-year fixed-rate MBS issued in a two-year period bracketing the two 2012 guarantee fee increase dates. I merged this data with weekly mid prices from Bloomberg, resulting in a sample of 680,110 MBS-week observations for 58,106 FNMA securities issued between July 2011 and June 2013. For the structural estimates in Chapter 4, I use a sample consisting of all FNMA and FHLMC 30-year fixed-rate mortgage MBS with a pool size of at least $1 million issued between February 2004 and March 2013, as the timing of origination more closely matches that of the other datasets I employ for structural estimation[73]. I observe 40,307 such MBS at origination, as well as their origination price and average characteristics. The full details regarding the assembly of these datasets can be found in Appendix B, and summary statistics are shown in Table A3 in Appendix C.

1.5.4 Dataset Assembly Procedures

I supplement the datasets described in Sections 1.5.1 and 1.5.2 by importing a number of aggregate and sub-aggregate variables and using them to impute certain data

[73]Data was not available for MBS issued between 2000 and 2003.

elements. For the GSE datasets I use in Chapters 2 and 3, I first obtain house price indices from Zillow and apply them at the Metropolitan Statistical Area (MSA)-month level, or, in cases where no MSA is listed in the GSE data, the state-month level, as these are the smallest levels of geographic identification available for all observations in the GSE dataset. I impute house prices by compounding the the observed price at origination, in turn imputed from the balance and LTV at origination, by percentage changes in local-area house price indices. Second, I merge local-area unemployment data from the Bureau of Labor Statistics (BLS)[74], as with home prices merging at the MSA-month level wherever available and the state-month level if the MSA in unobserved[75]. I also impute servicing decisions, monthly payments, home equity, and in some cases monthly balances directly from the GSE data via the appropriate formula, in the case of home equity using the imputed house price[76]. Finally I import monthly data on average mortgage interest rates from Freddie Mac and average interest rates from the Federal Reserve in order to measure how much mortgage rates have changed since the borrower took out their initial loan, a proxy for the value of refinancing.

I use a variety of other datasets in a more limited fashion for three specific purposes. First, to test the effect of guarantee fees on interest rates at higher frequencies in Chapter 2, I use daily data on state-level average jumbo and conforming refinancing rates provided by Bankrate.com. This data is comprised of 66,612 state-trading day observations for the pair of average refinance interest rates for every trading day from October 2009 through November 2014. Second, as an input to the structural model I present in Chapter 4, I also estimate borrower expectations as a model primitive using

[74] Unlike the headline figures reported each month, local-area unemployment statistics are not seasonally adjusted. Because I am more interested in a predictor of individual unemployment status than in identifying long-term trends, I feel the unadjusted figure is more useful.

[75] The MSA is not listed if the borrowers lives in a rural area. I considered using statistics on unemployment rates and house prices for rural areas within states, but those statistics are relatively noisy, and there few enough loans in rural areas to make this a minor concern.

[76] Monthly ending balances are censored in the data for the first six months the loan is active. For FRMs, monthly payments and balances are a fixed nonlinear function of the initial balance, interest rate, and term, and can thus be reconstructed manually.

vector autoregressions on county-month-level unemployment, house price, and interest rate data. Finally, in many of my baseline regression specifications in Chapters 2 and 3 I instrument for home prices using the housing-supply-elasticity instrument from Gyourko et al (2008)[106] and for unemployment using the instruments from Bartik (1991),[25] which measure the change in unemployment relative to a base year that would be predicted by national-level trends. I construct instruments for house price using the Wharton Residential Land Use Regulatory Index (WRLURI)[77] and for unemployment rates using national employment data and county-level industry-mix data from the 2000 Decennial Census. I discuss the construction and validity of these instruments in more detail in Sections 2.4, 3.2, and 3.3, and details on how each of these three datasets are assembled can be found in Appendix B.

The primary dataset I use for estimating the structural model in Chapter 4 is generated in several steps by merging property-level loan data with the subsample of GSE data consisting only of loans from California. The property-level dataset (henceforth "deeds data") contains administrative data on the universe of recorded home loans in California between 2000 and 2012 collected from county recorders and clerks, containing such information as the loan principal, the borrower's name and address, the date of origination, and appraisal information such as the property characteristics and appraised value[78]. I merge these two datasets using techniques adapted from Goodman and Li (2014)[95] and Mayer et al (2014).[156] In order to merge the datasets, I first organize the deeds data into a sequence of loan observations for each property, and use the name of the borrower to classify the transaction as a refinance (if the names match), default (if the name is a bank), or continuation (if the loan is last in the sequence). After removing some observations absent in the GSE data, such as ARMs

[77] See Gyourko et al (2008)[106] for details on how this index is measured.

[78] These data were collected from county recorder's offices by an anonymous data provider and made available to me through the Paul Milstein Center for Real Estate at Columbia Business School. Because refinancing loans also generate a new deed of trust, the records include refinances in addition to sales and defaults.

or loans with balances above the conforming loan limit, I assign the mortgages into bins based on discrete variables shared with the GSE data, such as the loan purpose, lender, location, and termination reason. Within each bin, I generate a distance metric between each potential GSE-deeds match based on the origination and termination dates and origination balance, which may not match exactly due to differences in record keeping. I then use a deeds-optimal Gale-Shapley algorithm [88] to match the two record types, representing preferences using the negative of the distance metric and discarding all matches above a specified cutoff[79]. By construction, the match resulting from this algorithm tends to over-represent relatively uncommon circumstances, as mortgages issued in rural areas, for example, where there are fewer loans within each bin, are easier to match. This problem causes the weights for the sample of matched loans to be heavily biased relative the U.S. as a whole, making it difficult to extrapolate the results of my analysis. To address this concern, I resample the data in order to more closely conform to the GSE aggregate statistics. I generate origination sampling probabilities for each GSE-deeds matched pair according to the observed distribution in the GSE dataset of several key variables, including origination year, postal code, and termination reason. For each matched pair terminating in refinance, I also generate sampling probabilities based on the refinancing balance by fitting a Normal distribution to the aggregate statistics reported in the Freddie Mac Cash-Out Refinance Report. I then draw 200,000 observations at random according to these sampling probabilities. As above I supplement each observation with local-area data on unemployment rates, house price appreciation and mortgage rates in order to finalize the dataset, henceforth referred to as the "matched dataset". The full details of the matched dataset assembly procedure can be found in Appendix B, while the average characteristics of borrowers in this dataset is shown in Table A6 in Appendix C.

[79]While it is possible that this procedure may identify false matches, as in Crawford and Yurukoglu (2012) [60] the estimates should be unaffected so long as any mis-classification is random conditional on observables.

Matching these two datasets provides three key benefits relative to using solely the GSE data. First, the merge enables me to observe complementary data elements not present in the GSE dataset, including characteristics of the property itself, such as the size of the house, and far more detailed location data. I use this enhanced geographic detail, which allows me to observe the borrower's exact address rather than their MSA or state, to import more accurate local unemployment rates and home prices at the place-level rather than the MSA- or state-level. Accurate data on these figures are available at the place-level from the same sources, the BLS and Zillow respectively, and enable me to distinguish between, for example, Compton and Beverly Hills, which are within the same MSA but have different trends with respect to unemployment rates and house prices. Second, the deeds data allows me to distinguish between home sales and refinances, which the GSE dataset does not, and so I can be confident that the matched dataset consists only of refinances and not sales[80]. Finally, because I observe the loan a borrower refinanced into in the deeds data, I impute the balance and hence the LTV of their refinance application, which would be impossible using the GSE dataset alone, as loans are only observed until prepayment. As suggested in Section 1.4, this variable, the borrower's LTV choice on refinance, will be a crucial target for my structural estimation procedure in Chapter 4.

Figure 1.6 compares the distribution of originations by year in the GSE dataset, deeds dataset, and matched dataset. By construction, the three datasets match very closely. Owing to the design of the sampling scheme, several broad trends in the characteristics of originated and refinanced loans are common across the datasets. In particular, annual patterns in the percent of loan originations made as refinancing loans, the percent of refinanced loans increasing their balance, the median home price appreciation on refinance, and total home equity cash extraction as a percent of refinancing volume are

[80] In Sections 2.4, 3.2, and 3.3 I discuss the assumptions needed to interpret changes in prepayment, which includes sales and refinances and is the outcome observed in the GSE dataset, as indicative of changes in refinancing behavior.

Figure 1.6: Comparison of Origination Year Distribution Across Datasets

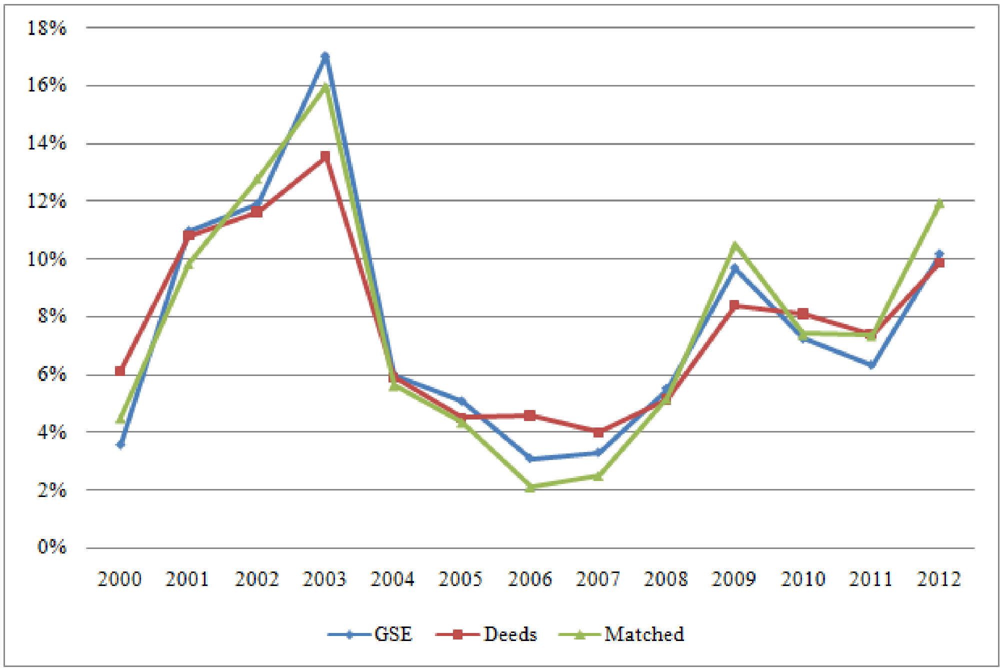

Notes: Distribution of observations by origination year in GSE, deeds, and matched datasets.

all quite similar for the different datasets. The comparisons of these statistics across datasets are shown Figures A2, A3, A4, and A5 in Appendix C.

At various points throughout my analysis I will treat imputed house prices or local-area unemployment rates as a proxy for the corresponding individual-level characteristic, and some discussion of the validity of this approach is warranted. Regarding home prices, it is standard practice in the literature[81] to impute home prices from observed purchase prices compounded at local-area aggregate appreciation rates. While this procedure may introduce some measurement error in home prices, one might reasonably suppose that individuals themselves estimate their own current house price in a similar fashion, making this imputed value is the best determinant of individual behavior, and in any event more accurate estimates of home values do not exist. Regarding unemployment, in most specifications I use local unemployment rates as a proxy for borrower-level unemployment, as individual-level measures are unavailable.

[81] See An et al (2010),[18] Bajari et al (2013),[22] or Tracy and Wright (2012)[185] for examples of this imputation.

While this may lead to systematic bias in the estimated effect of unemployment, the results from Gyourko and Tracy (2014) [107] suggest that using sub-aggregate rates as a proxy will bias any estimated coefficients towards zero, so any measured effect of unemployment will represent a lower bound on the true effect. Of greater concern is the possibility that the population of interest in most of my analysis, homeowners with GSE-guaranteed mortgages, may systematically differ from the aggregate in terms of either unemployment rates or house price appreciation. While I cannot address this concern directly, comparisons of the estimated unemployment rate among homeowners derived from Current Population Survey (CPS) microdata suggest that the unemployment rate for homeowners tends to track the overall unemployment rate closely. Moreover whenever possible I apply home price indices for the applicable type of home (e.g. single-family, condominium, etc.), and because GSE-guaranteed mortgages are a large share of the market, this figure may be quite representative of the average GSE-mortgagee's realized home price appreciation. Nevertheless, it is important to keep in mind in the analysis that follows that employment status and house prices are not observed on an individual basis.

1.6 CONCLUSION

This chapter provides an introduction to the mortgage market and the GSEs that will help frame the remainder of my analysis. The descriptive analysis presented in Section 1.2 offers some insight into how GSE policy can affect credit constraints and in turn borrower welfare. Changes in guarantee fees affect the profitability of selling or securitizing via the GSEs, which on the margin will change the profitability of originating new mortgages. As a consequence, when guarantee fees increase borrowers may receive higher interest rate offers or may have their applications denied, leaving them unable to reduce their monthly payments via refinancing and consume the resulting savings. Crucially, however, the extent of this pass-through from guarantee fees to credit constraints depends on the relative value of different loan funding channels for mortgage originators, and the cost to borrowers from reduced credit supply depends on how highly they value the refinancing option. Section 1.4 formalizes this intuition with a stylized model. I show analytically that under certain assumptions, increases in guarantee fees lead to higher interest rates and higher application denial rates, while borrowers with higher liquidity preference are both more likely to apply for refinancing credit and more likely to have their applications denied. In the following chapters, I first provide empirical support for each of these two results, and subsequently analyze via simulation the effects of a policy that reduces guarantee fees at times when liquidity preference is likely to be higher on average.

The remainder of my analysis is structured as follows. In Chapter 2, I analyze the effects on the mortgage market of two large guarantee fee increases implemented in 2012. Using an event-study framework, I provide the first empirical evidence for the effects of guarantee fee increases, and find that in the short-run lenders absorb some of the fee increase through lower revenues, while in the long-run virtually all of the fee increase is transmitted to mortgage interest rates, consequently reducing refinancing volumes and increasing the propensity to default. These results provide evidence for the mechanisms described in Section 1.2 and analyzed in Section 1.4 regarding the effects of guarantee fee increases and also offer an empirical grounding for the model of credit supply I develop in Chapter 4. In Chapter 3 I analyze the relationship between unemployment, credit constraints, and refinancing in order to address the question of how the value of refinancing varies with unemployment. I use exogenous variation in the sensitivity of credit constraints to unemployment induced by the HARP program to show that the unemployed are both more likely to refinance when credit constraints are attenuated and benefit more from credit-expanding policies such as HARP. These results, which provide the first direct evidence of heterogeneity in takeup under the

HARP program, both suggest that liquidity preference of the sort described in Section 1.4 is higher for the unemployed and help ground the structural model of credit demand I develop in Chapter 4. Finally in Chapter 4 I extend the literature on household financial management by designing a structural model of the mortgage market featuring endogenous credit constraints and latent borrower liquidity preference. I estimate the model using a novel combination of empirical techniques and use the calibrated model to simulate the effects of various credit-supply interventions on the part of the GSEs. My analysis indicates that borrowers would benefit from a countercyclical guarantee fee policy, and I conclude with a discussion of the merits of such a policy relative to other proposed or implemented alternative policies.

Chapter 2

Guarantee Fee Increases and Mortgage Credit Supply

2.1 Introduction

In this chapter I analyze the effects of changes in guarantee fees on the mortgage market, and in particular on credit constraints. As discussed in Sections 1.2.3 and 1.2.4, during the financial crisis in 2008, default rates for mortgages insured by the GSEs rose by an order of magnitude, leading to a combined $109 billion loss for the two agencies in that year. In response, the government effectively nationalized FNMA and FHLMC in September of 2008 in order to forestall a loss of investor confidence in their insurance obligations. Thereafter the two agencies, formerly nominally private, set their policy explicitly as government policy. Following the financial crisis, the policy stance of FNMA and FHLMC changed considerably, and of particular relevance for my analysis, they raised their guarantee fees substantially in order to recoup losses incurred during the crisis. My analysis uses an event study framework to analyze the effects of two large, discrete guarantee fees increases implemented in 2012, both of which can be seen in Figures 1.2 and 1.4. As the focus of my dissertation is on credit supply policies and how credit constraints affect borrowers, my analysis in this chapter traces the impact of these fee increases on credit supply in order to motivate the structural model I develop in Chapter 4.

Building on the insights from Sections 1.2.2 and 1.2.3 regarding the manner in which guarantee fees affect credit supply, I analyze the effects of the GSEs policy change in several stages. First, using data on pools of mortgages traded over the secondary market, I study the effect of the fee increases on issuance volumes, the coupons paid on new MBS, and the prices at which those MBS trade. Second, using data on refinancing loans and dynamic borrower behavior, I analyze the effect of the fee increases on mortgage interest rates and borrower behavior. I evaluate the extent to which these fees are passed through to mortgage interest rates, and subsequently what effect the resultant change in credit supply has on the propensity of borrowers to either refinance or default on their mortgages. Throughout much of this analysis I treat the change in guarantee fees as an exogenous shock to the profitability of lending and to credit supply. Because the first of these fee increases was mandated by congress as part of the Temporary Payroll Tax Cut Continuation Act of 2011, and the second was mandated by the FHFA as part of their long-term strategy for the GSEs and announced to coincide with the release of their annual report to congress,[80] I assume that the policy changes were uncorrelated with other developments in mortgage markets. I then compare outcomes for affected mortgages before and after the exogenous policy change on both the secondary and primary market.

My results suggest that guarantee fees have varying effects on the mortgage market at different horizons. In the short-term, mortgage lenders partially mitigate the effects of the fee increase by both accelerating the pace at which they package and issue new MBS in order to avoid the fee and by reducing the excess interest rate spread they retain from newly-issued securities. However, at fixed interest rates, the coupons paid on and prices commanded by MBS declines following the increase in guarantee fees, resulting in lower revenues for mortgage lenders. In response, lenders pass the vast majority of the fee increase on to borrowers in the form of higher mortgage interest rates, with a pass-through rate of between 85% and 100% depending on the specification. I find that lenders adjust to the policy change quite rapidly, with essentially all of the pass-through completed after three weeks. Consistent with earlier findings, the extent of this pass-through is higher for riskier borrowers and for nonbank lenders who are more reliant on GSE securitization, but is lower in more concentrated lending markets. As a result of this supply-side response, the probability that a borrower refinances declines and the probability of default increases, likely reflecting a tightening of household credit constraints. The magnitude of these effects is consistent with previous analyses that assume full pass-through from guarantee fees to mortgage rates. Applying my estimates to an aggregate figure for the quantity of mortgages affected yields an estimated decline of roughly \$205 billion per year in the volume of mortgage refinancing.

The key contribution of this chapter to the existing literature is to provide direct evidence of the effects of changes in guarantee fees on the mortgage market. Contemporary trade publications recognize the importance of guarantee fees for determining mortgage credit supply,[57, 189] yet to my knowledge no prior research has investigated such a connection directly. The change in credit supply induced by the GSE's policy changes is of particular relevance for macroeconomic policy as it partially determines refinancing activity and thus the extent of monetary policy transmission or other housing finance policies targeting the refinancing channel, such as HARP. While this analysis is most directly relevant for evaluating changes to GSE policies themselves, it is my hope that my estimates prove useful for a variety of macroeconomic policymakers. A secondary contribution of this analysis is to provide novel evidence on the effects of credit constraints. Changes in guarantee fees provide a very tractable setting in which to analyze the effects of credit supply, as unlike changes in interest rates they are essentially exogenous, have effects solely on certain segments of the mortgage market and not on other markets, and do not affect rates of discount or available investment returns. While my findings regarding the elasticities of refinance and default to changes in credit supply largely confirm prior research, additional evidence on this subject using a novel identification strategy is not without merit. Finally, the results from this chapter will help inform the simulation study of the effect of alternative guarantee fee policies that I conduct in Chapter 4.

The analysis in this chapter proceeds as follows. In Section 2.2 I briefly describe the secondary market data used in this study, present some suggestive patterns, and analyze the effect of guarantee fees on MBS traded on the secondary mortgage market. Section 2.3 studies how lenders respond to these changing secondary market conditions, an in particular the timing and extent of pass-through to primary market interest rates. In Section 2.4 I turn to the effects of these supply-side responses on borrowers, which is the main focus on my analysis, and estimate how the resulting change in credit supply affects household refinancing and default behavior. Section 2.5 concludes with a discussion of the relevance of these findings for macroeconomic and housing finance policy and contextualizes my results within my dissertation as a whole and the broader academic and policy literature.

2.2 Effects on the Secondary Market

In this section I analyze the effect of guarantee fee increases on the secondary market. Recall from Section 1.2.2 that investors trade mortgages as MBS on the secondary market well after they are originated, and that the prices these MBS command in part determines the interest rates originators are willing to offer. In order to analyze how guarantee fees affect the prices or MBS on the secondary market, I use a dataset consisting of all 30-year fixed-rate MBS issued via FNMA or FHLMC in a two-week window around each of the 2012 guarantee fee increase dates, as well another consisting of all FNMA 30-year fixed-rate MBS issued in a two-year period bracketing the two 2012 guarantee fee increase dates, as described in Section 1.5.3. In my analysis, I first consider how originators respond in the short-term to the fee increase via their securitization schedules and via the excess servicing spread they receive. I subsequently analyze the effect of the fee increase on MBS prices, the largest component of originator profits, in order to measure how the fee increase affects the incentives to originate new mortgages.

2.2.1 MBS Issuance

Several patterns are immediately apparent from the MBS issuance data. Table 2.1 shows total MBS issuance volume in the weeks just prior to and just after the two guarantee fee increases implemented in 2012. These fee increases essentially function like taxes on mortgage securitization. Hence if originators are aware of the fee increase, as we should expect in this case as both were announced months in advance, they should seek to avoid the tax as much as possible by shifting the timing of their securitization. Table 2.1 provides evidence for such shifting, as the total volume of MBS issuance falls considerably just after the fee increase. In both month-long windows, originators issued half of their securities, or twice what we would expect if originations were uniformly distributed across weeks, in the week just prior to the fee increase. Figure 2.1 shows similar shifting for FNMA securities. In the months directly after the fee increase, denoted by vertical lines, issuance fell sharply, with a corresponding peak in the month just prior. Taken together these patterns suggest that originators accelerated their securitization schedules forward in order to dodge the increase in fees, providing suggestive evidence that these fees do have an effect of lending incentives. Figure 2.1 also shows that over a longer horizon, fee increases appear to have no effect on issuance volume. Both fee increase dates result in temporary disruptions to an otherwise smooth trend as some securitization is shifted into the prior month, but with little evidence of long-term market disruption. The absence of obvious effects at horizons longer than a month may suggest the presence of a compensating market response on the part of originators. Table 2.1 and Figure 2.1 thus suggest a short-term shock to the market followed by reequilibration. In this section, I examine the nature of the short-term shock on secondary markets in detail, while Sections 2.3 and 2.4 considers the longer-term response by, respectively, lenders and borrowers.

2.2.2 MBS Characteristics

In this section I assess the short-run impact on the secondary mortgage market of the two guarantee fee increases implemented in 2012. My hypothesis is that guarantee fees affect MBS prices and hence originator profitability through the cash flows for MBS deals outlined in Figure 1.1 and Section 1.2.3. When guarantee fees increase, we should expect that, holding mortgage interest rates constant, some combination of excess

Table 2.1: Short-Term MBS Issuance Bunching

	MBS Issuance ($ Billion)		MBS Issuance (% of Total)	
	April 1st, 2012	December 1st, 2012	April 1st, 2012	December 1st, 2012
2 Weeks Prior	9.4	10.7	13.7	14.1
1 Week Prior	36.6	37.9	53.4	49.9
1 Week After	2.3	16.2	3.4	21.3
2 Weeks After	20.3	11.2	29.6	14.8
Prior	46.0	48.6	67.0	63.9
After	22.6	27.4	33.0	36.1

Notes: Total FNMA and FHLMC 30-year FRM MBS issuance in two-week window on either side of guarantee fee increase implementation date. "% of Total" denotes measured relative to 4-week total.

Figure 2.1: Long-Term MBS Issuance Bunching

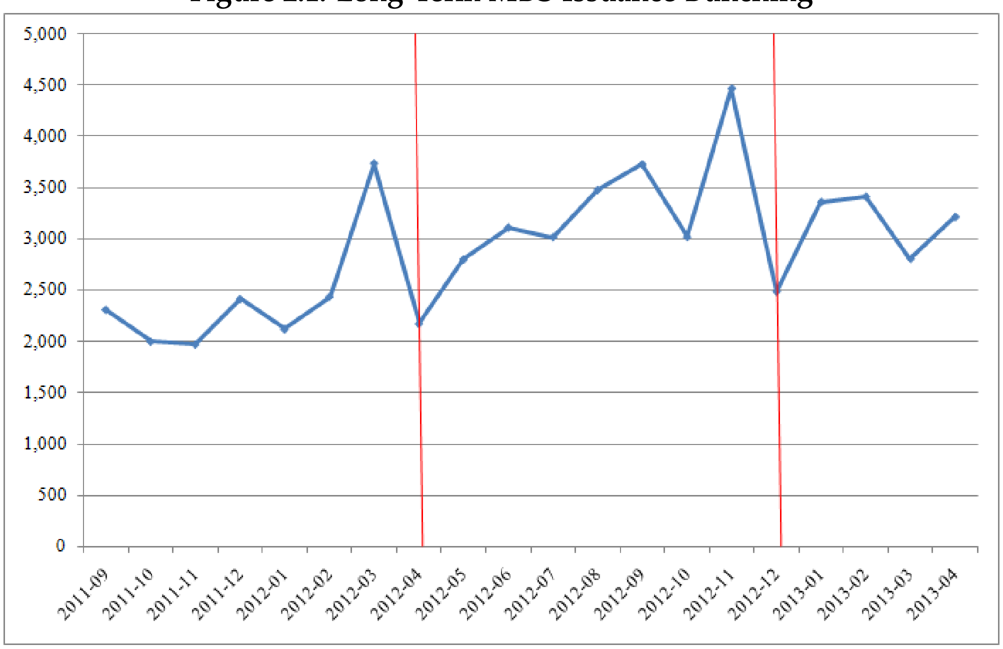

Notes: Total FNMA 30-year FRM MBS issued by month.

servicing spreads and coupons must decline in exact proportion, in either case leading to lower overall origination revenues for mortgage lenders[1]. My framework for analyzing these effects is an event study using the two fee increase implementation dates, April 1st, 2012, and December 1st, 2012[2]. While lenders may have begun responding to the fee increases as early as the announcement dates[3], we should expect the policy change to have visible effects on the secondary market only after the implementation date, as the guarantee fee increase did not apply to securities issued prior to implementation[4]. Assuming, as discussed in Section 1.2.4, that the guarantee fee increase was due to an exogenous policy change, in what follows I treat the implementation date as an instrument for guarantee fees. To test the validity of using this date as an instrument, I estimate a first-stage regression of guarantee fees on whether the the security was issued after the implementation date:

$$FEE_i \quad = \quad ISSUEPOST_i\beta_1 + RISK_i\beta_2 + IR_i\beta_3 + e_i$$

where i indexes securities, $ISSUEPOST_i$ is an indicator for whether the security was issued post-implementation, $RISK_i$ are risk characteristics and IR_i is the coupon rate. The guarantee fee, FEE_i, is imputed as the difference between the WAC, the average interest rate paid by borrowers whose mortgages compose the pool, and the coupon paid on the security; hence it measures the sum of the guarantee fee and any excess servicing spreads. The parameter of interest is β_1, the effect of the security being issued post-implementation on this combination of guarantee fees and excess servicing spreads.

Table 2.2 shows the estimated coefficients and standard errors of β_1 for five separate

[1] My data does not allow me to observe the costs of mortgage origination, hence I can analyze only the effect of guarantee fees on originator revenues and not on profits.

[2] Guarantee fees increased by the same amount on November 1st, 2012 rather than December 1st for mortgages sold for cash rather than delivered into MBS.[80] Since my secondary market analysis uses only data on MBS, the relevant date for this study should be December 1st, 2012.

[3] Respectively, December 30th, 2011 and August 31st, 2012

[4] In Sections 2.3 and 2.4 I discuss the timing of supply-side responses to the fee increase, and for most of those specifications I assume lenders respond beginning on the announcement date.

Table 2.2: Guarantee Fee Increase Event Study

	Model 1	Model 2	Model 3	Model 4	Model 5
Issued Post-April 2012	0.0532***	0.0838***	-	-	0.0748***
	(0.0042)	(0.0138)	-	-	(0.0020)
Issued Post-December 2012	-	-	0.0379***	0.0450***	0.0430***
	-	-	(0.0047)	(0.0059)	(0.0019)
Implementation Date	April 2012	April 2012	December 2012	December 2012	Both
Sample Window	2 Weeks	1 Week	2 Weeks	1 Week	6 Months
Fixed Effects	State	State	State	State	State, Orig. Year
N	2,696	1,554	3,646	2,338	58,106
R^2	0.1975	0.2045	0.1488	0.1615	0.1459

Notes: Dependent variable is difference between MBS coupon and WAC. Standard errors in parentheses. Additional controls include quartiles of LTV, FICO and interest rate, weighted-average maturity, and % refinance and owner-occupied. ***/**/* indicates significance at 99%/95%/90%-level respectively.

versions of the event study above[5]. Models 1 and 2 use all FNMA and FHLMC MBS issued in respectively a two-week or one-week window around April 1st, 2012, while Models 3 and 4 use a similar sample of MBS issued around December 1st, 2012. Model 5 uses a sample of FNMA MBS from the longer-term MBS dataset issued beween six months prior to the first implementation date (October 2011) and six months after the second date (June 2013) and measures the effects of both increases jointly in the same model[6]. In both April 2012 and December 2012 guarantee fees increased by 10 bp, so if there were no adjustment on excess servicing spreads the estimated coefficient would be .1. In each model, the estimated coefficient is somewhat smaller than .1, suggesting that in the short-term some fraction of the guarantee fee increase is offset by originators receiving lower servicing spreads. The extent of this offset is greater in December 2012, so that more of the impact of the guarantee fee increase is absorbed by lower servicing spreads, but my estimates suggest that on average lenders reduced their excess servicing by between 20% and 60% of the guarantee fee increase. Because servicing spreads directly determine originator revenues, the effect of the guarantee fee increase is to reduce the incentive to issue securities regardless of whether it results in lower coupons

[5]The other estimated coefficients and standard errors are shown in Table A8 in Appendix D; all five models in Table A8 are the same as those shown here.

[6]Since the implementation dates correspond to the beginning of months, I can measure with certainty which regimes individual securities were issued under even if, as in the case of the longer-term MBS dataset, I can only observe the month of issue.

or lower servicing spreads.

While the results from Model 5 indicate that there was a sustained increase in average guarantee fees following the GSE's policy change, two key questions remain regarding the timing of this effect. First, whether the partially offsetting reduction in servicing spreads was a permanent or temporary response on the part of originators, and second, whether the estimate effect was part of a pre-existing trend or due entirely to the policy change. To address these questions, I estimate an event study similar to Model 5 using month-level fixed effects in order to better understand the timing of the effect of the policy change. I estimate:

$$FEE_i \quad = \quad \Sigma_{s=1}^{S} \mathbf{1}[ISSUEDATE_i = s]\delta_s + RISK_i\beta_2 + IR_i\beta_3 + e_i$$

with all variables as before. The parameters of interest are the δ_s terms, the month-level fixed effects. Figure 2.2 plots these fixed effects by month. There does not appear to be a trend in guarantee fees prior to the first fee increase date, suggesting that the results from Table 2.2 are not driven by pre-existing trends. Table A9 in Appendix D shows the estimated fixed effects underlying Figure 2.2, and the estimate for the month prior to the first policy change date is only 1 bp above the level from five months prior. On implementation, the estimated fixed effect jumps upward, but by considerably less than 10 bp. The estimated fixed effects continue rising after implementation for some time, reaching 8.5 bp above the pre-period level after 3 months and the full 10 bp above the per-period level after 7 months. For the second implementation date, the difference between the last estimated fixed effect, March 2013, and the final pre-period month, November 2012, is roughly as large (6.25 bp) as the difference between comparable months for the first implementation date (7.5 bp difference between July 2012 and March 2012), indicating that the timing of adjustment is similar in both cases. While there is a pre-trend in the months leading up to the second guarantee fee increase in December

Figure 2.2: Guarantee Fee Fixed-Effects Estimates

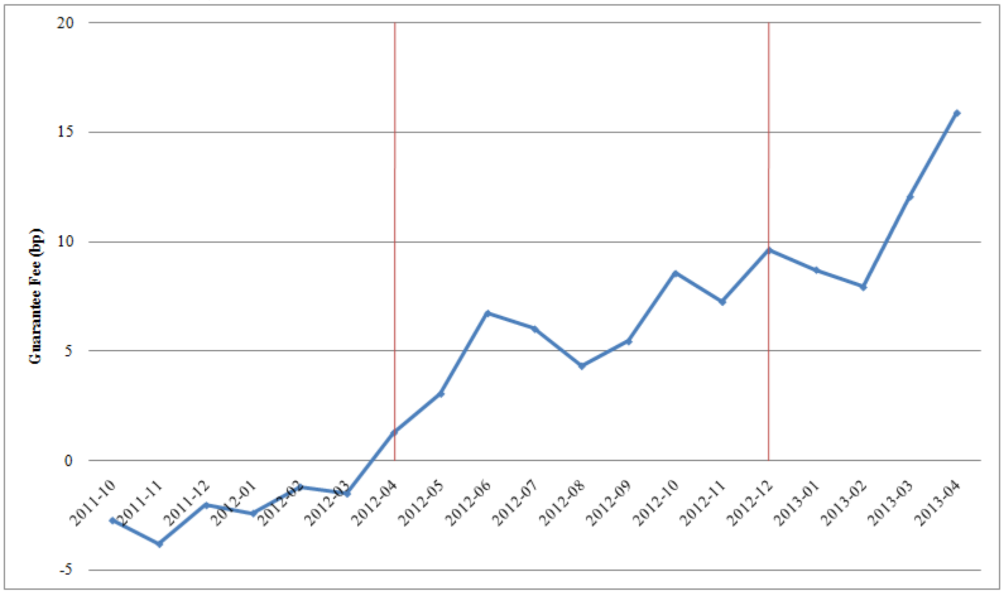

Notes: Estimated month fixed effects for guarantee fee increase event study.

2012, I interpret this trend as the gradual continuing effect of the first fee increase in April. Figure 2.2 thus suggests that in the short-term originators reduce their excess servicing spreads in response to the guarantee fee increase, but gradually adjust by absorbing less and less of the fee increase via lower spreads, with full adjustment after seven months. Taken together, the models presented in Table 2.2 and Figure 2.2 provide reasonable evidence that an event study using the two guarantee fee increase implementation dates is a valid empirical design for estimating the effects of GSE policy changes on originator revenues, as the estimates from Table 2.2 indicate that the implementation date is a relevant instrument and the lack of a clear pre-trend visible in Figure 2.2 suggests that the change is exogenous.

With the validity of the event-study approach established, I next assess the short-run impact of guarantee fee increases on originator revenues. Recall from Figure 1.1 and the discussion from Section 1.2.3 that at a given level of interest rates, guarantee fees can affect lender revenues and hence profits through either the servicing spread or the coupon on the security, which in part determines the security's price. The results from Table 2.2 and Figure 2.2 suggest the two policy changes had some effect on servicing

Table 2.3: MBS Coupon Event Study

	Model 1	Model 2	Model 3	Model 4	Model 5
Issued Post-April 2012	-0.0614***	-0.0637***	-	-	-0.0759***
	(0.0042)	(0.0108)	-	-	(0.0022)
Issued Post-December 2012	-	-	-0.0276***	-0.0443***	-0.0432***
	-	-	(0.0042)	(0.0060)	(0.0020)
Implementation Date	April 2012	April 2012	December 2012	December 2012	Both
Sample Window	2 Weeks	1 Week	2 Weeks	1 Week	6 Months
Fixed Effects	State	State	State	State	State/Orig. Year
N	2,696	1,554	3,646	2,338	58,106
R^2	0.9447	0.9518	0.9258	0.9326	0.9528

Notes: Dependent variable is MBS coupon. Standard errors in parentheses. Additional controls include quartiles of LTV, FICO and interest rate, weighted-average maturity, and % refinance and owner-occupied. ***/**/* indicates significance at 99%/95%/90%-level respectively.

spreads, so the natural next question is what effect they had on coupons. To test this response, I estimate models of the form:

$$COUPON_i = ISSUEPOST_i\beta_1 + RISK_i\beta_2 + IR_i\beta_3 + e_i$$

where all variables are as described above. As before, the key parameter of interest is β_1, the estimated effect on MBS coupons of being issued after the fee increase date.

Table 2.3 shows the estimated coefficients and standard errors from this event study[7]. Because the regression controls for the WAC of the pool, this is largely an accounting exercise. We know from tracing the flow of money in Figure 1.1 that holding fixed the payment borrowers make into the pool, the sum of the coupon and servicing spread must decline by the amount of the guarantee fee increase. The estimated coefficients from Models 1 through 5 in Table 2.3, which correspond to the samples used in Models 1 through 5 in Table 2.2, bear this out, as they are largely the same. Coupons tend to fall in the short-term by roughly half of the 10 bp increase in guarantee fees, corroborating the results from Table 2.2 that the remaining half was realized as lower servicing spreads. As before, coupons respond somewhat more to the April 2012 fee increase, and the overall effect is visible at longer horizons in both cases. I also estimated a similar version of this

[7]The other estimated coefficients and standard errors are shown in Table A11 in Appendix D, with the five models corresponding as above.

model using 2-stage least squares by first predicting the change in guarantee fees at the implementation date using the estimates from Table 2.2 and then including these predicted values as a regressor. I estimate the model:

$$\begin{aligned} FEE_i &= ISSUEPOST_i\alpha_1 + RISK_i\alpha_2 + IR_i\alpha_3 + e_i \\ COUPON_i &= \widehat{FEE}_i\beta_1 + RISK_i\beta_2 + IR_i\beta_3 + e_i \end{aligned}$$

where $\widehat{FEE}_i$ is the predicted fee. Those estimates are shown in Table A10 of Appendix D, and as would be expected, the estimated coefficients are roughly -1, implying a direct proportional response of coupons to guarantee fees.

2.2.3 Origination Profitability

The significance of the findings from Section 2.2.2 that higher guarantee fees reduce MBS coupons is that MBS coupons are the primary determinant of MBS prices, the largest component of originator revenues, and hence directly determine lender origination incentives. Thus in this section I analyze the short-term effect of guarantee fee increases on MBS prices. Using panel data on secondary market prices, I estimate security-day level models of the form:

$$PRICE_{i,t} = ISSUEPOST_{i,t}\beta_1 + RISK_i\beta_2 + IR_i\beta_3 + DATE_{i,t}\beta_4 + AGE_{i,t}\beta_5 + e_{i,t}$$

where $PRICE_{i,t}$ is the observed price of security i on trading day t, $RISK_i$ and IR_i its risk characteristics and coupon rate at origination, $AGE_{i,t}$ its time since issuance, included to correct for any effects of the staggered timing of coupon payments, and $DATE_{i,t}$ are trading-day fixed effects, included to capture any market-wide factors.

Table 2.4 presents the results from this specification, with Models 1 through 5 using the security-day-level panel dataset versions of the datasets used in Models 1 through 5

Table 2.4: MBS Price Event Study

	Model 1	Model 2	Model 3	Model 4	Model 5
Issued Post-April 2012	-0.2251***	-0.3272***	-	-	-0.2058***
	(0.0152)	(0.0286)	-	-	(0.0045)
Issued Post-December 2012	-	-	-0.4447***	-0.3676***	-0.3515***
	-	-	(0.0221)	(0.0469)	(0.0057)
Implementation Date	April 2012	April 2012	December 2012	December 2012	Both
Sample Window	2 Weeks	1 Week	2 Weeks	1 Week	6 Months
Fixed Effects	State/Day	State/Day	State/Day	State/Day	State/Day/Orig. Year
N	20,512	6,582	23,777	7,922	499,321
$\mathbf{R^2}$	0.8718	0.8958	0.7453	0.7529	0.8036

Notes: Dependent variable is MBS price. Standard errors in parentheses. Additional controls include quartiles of LTV, FICO and interest rate, weighted-average maturity, days from issuance, and % refinance, third-party and owner-occupied. ***/**/* indicates significance at 99%/95%/90%-level respectively.

of Table 2.3[8]. In each case, the decrease in prices for securities issued just after the implementation date is a plausible multiple of the fee increase or, in turn, the decline in coupons. Comparing Models 2 and 4 in Tables 2.3 and 2.4, the implied capitalization rate of prices over coupons is in the range of 5 to 8 times. This difference in capitalization rates between April and December may reflect market factors, but the range overall is quite reasonable. As discussed in Section 1.2.3, FNMA and FHLMC both offer originators the opportunity to reduce their ongoing guarantee fees with up-front payments at a conversion rate known as a buy-down multiple. This multiple, which effectively measures the same capitalization ratio, typically varies between 4 and 7 depending on market factors, suggesting that the estimates from Table 2.4 accord well with FNMA's and FHLMC's own documentation. As in Section 2.2.2, I also estimate a 2-stage least squares version of the model above:

$$\begin{aligned}
FEE_i &= ISSUEPOST_i\alpha_1 + RISK_i\alpha_2 + IR_i\alpha_3 + e_i \\
PRICE_{i,t} &= \widehat{FEE}\beta_1 + RISK_i\beta_2 + IR_i\beta_3 + DATE_{i,t}\beta_4 + AGE_{i,t}\beta_5 + e_{i,t}
\end{aligned}$$

where the coefficient β_1 in this case directly measures the present-value multiple. These estimates are shown in Table A12 of Appendix D, and in most cases they also fall within

[8]The other estimated coefficients and standard errors are shown in Table A13 in Appendix D, with the five models corresponding as above.

the range offered by FNMA and FHLMC. As with the estimates from Table 2.3, by controlling for the WAC of the MBS these regression specifications essentially remove any compensating change in interest rates on the part of originators, making this largely an accounting exercise. Nevertheless, these results do provide direct evidence for the transmission of guarantee fees to originator revenues through the MBS price channel.

As with MBS coupons in Section 2.2.2, I also estimate an event study similar to Model 5 from Table 2.4 using week-level fixed effects in order to assess the timing of transmission from guarantee fees to MBS prices. The empirical model is:

$$\begin{aligned} PRICE_{i,t} &= \Sigma_{s=1}^{S} \mathbf{1}[ISSUEDATE_i = s]\delta_s + RISK_i\beta_2 \\ &+ IR_i\beta_3 + DATE_{i,t}\beta_4 + AGE_{i,t}\beta_5 + e_{i,t} \end{aligned}$$

where δ_s coefficients, the estimated effects on prices of issuance at any week before or after the fee increase date, are the parameters of interest. Figure 2.3 plots these estimated fixed effects by week, with Panel A plotting the effects in a two-month window around April 1st, 2012, and Panel B respectively around December 1st, 2012. The estimates suggest a negligible pre-trend in prices leading up to the implementation date, providing further validity for the event study design. Following implementation (Week 0), prices for newly issued securities are not immediately affected, although the negative effects in both cases become visible in Week 1, with subsequent negative estimates in Weeks 3, 5 and 7 after the fee increase. These estimates primarily reflect the structure of the secondary mortgage market. As discussed in Section 1.2.2, much of the trading in agency MBS takes place through the TBA market, a standardized futures market with a pre-set schedule of dates on which settlement occurs. As a result, the highest volume of issuance and trading for 30-year FRM-backed agency MBS (Class A MBS) is on the settlement date[9], which is typically the second Thursday of each month. These

[9]Vickrey and Wright (2013)[186] provide additional detail on the mechanics of TBA market settlement and trading.

settlement dates correspond with the Week 1 and Week 5 fixed effects, while the Week 3 and Week 7 fixed effects correspond to settlement dates for another class of FHLMC and FNMA mortgages[10]. Since there are no settlement dates for any agency products in the first or third week of the month, corresponding to Weeks 0, 2, 4 and 6 in Figure 2.3, the absence of strongly visible negative effects on those dates could reflect a lack of issuance volume or a sample selection problem where only certain unusual MBS are issued off the settlement date schedule. Recall also from Table 2.1 that issuance volume fell sharply in the weeks just after the fee increase, suggesting that the sample corresponding to Week 0 is highly selected[11]. Hence if we treat Week 5, the next Class A settlement date, as a good indicator for the effect of the fee increase, the estimated effect of the guarantee fee increase on MBS prices is between -30 bp and -40 bp, roughly in line with the estimates from Table 2.4. In conjunction with the evidence presented in Table 2.4, these estimates suggest a clear negative relation between guarantee fee increases and MBS prices at fixed levels of mortgage interest rates.

The results from this section provide some evidence for the transmission of guarantee fees through the secondary mortgage market. At the policy implementation date, originators shift their securitization schedule forward to avoid the higher fees, although this shifting appears to be confined to the month directly before and after the fee increase. Over short horizons, originators absorb roughly half of the guarantee fee increase through lower servicing spreads, although over longer horizons essentially all of the fee increase manifests as lower MBS coupons. The effect of the fee increase on MBS prices is roughly what would be predicted from applying FNMA and FHLMC's documented present-value multiples to the change in guarantee fees. In all cases, I control for the effect of the WAC on MBS coupons and prices, effectively removing the

[10]This class, Class D, includes many other FNMA and FHLMC products, but no Class D instruments are included in this sample. Plausibly, there is additional trading volume for all Class A products, which compose my data, on settlement dates for other classes.

[11]At longer horizons and lower frequencies, there does not appear to be a sample selection issue. Table A7 in Appendix D shows the average characteristics of newly-issued MBS by month, and there is no obvious change in the months just following the fee increase implementation dates.

Figure 2.3: Price Impact Fixed-Effects Estimates

Panel A: April 2012 Increase

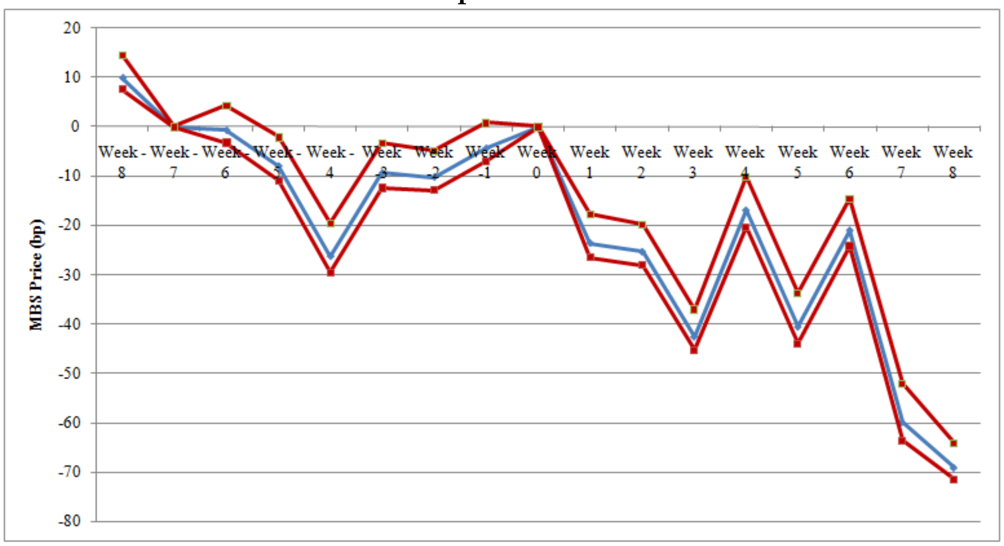

Notes: Estimated week fixed effects for MBS price increase event study. April 2012 increase. Mean estimated effect shown in blue, 95% confidence intervals shown in red.

Panel B: December 2012 Increase

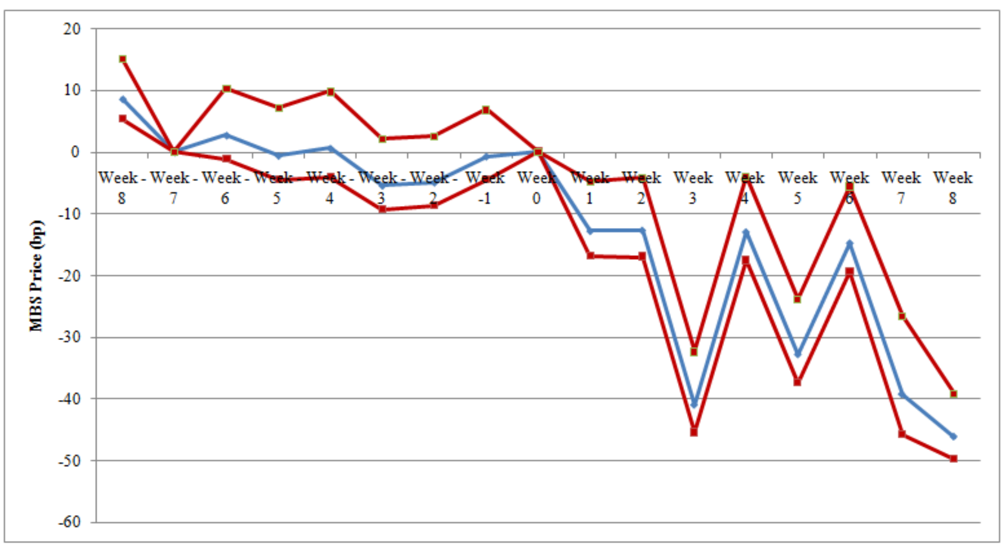

Notes: Estimated week fixed effects for MBS price increase event study. December 2012 increase. Mean estimated effect shown in blue, 95% confidence intervals shown in red.

effects of any pass-through from guarantee fees to interest rates. In the next section, I study the equilibrium response of mortgage originators on this margin and what effects this response has on credit constraints.

2.3 Effects on Primary Market Interest Rates

In this section I assess the equilibrium effect on primary market interest rates the two guarantee fee increases implemented in 2012. The results from the previous section indicate that, holding mortgage interest rates fixed, originator revenues should decline when guarantee fees increase. However, we should not expect originators to hold mortgage interest rates fixed in response to the policy change, and some fraction of their increased costs should be passed-through to prices. Hence in this section I analyze the extent to which mortgage originators respond to the fee increase by changing their offered mortgage interest rates, or effectively the extent of pass-through. Following Fuster et al (2013),[84] I assume that originators respond immediately after the announcement of the policy change. In the case of the two fee increases I analyze, these dates correspond to adjustments beginning in January 2012 and September 2012[12]. This assumption accords well with the expected timeline for an agency securitization, as the whole process of originating and securitizing a new mortgage typically takes several months. Hence, if originators learned of the fee increase three months prior to implementation (as they did in this case), some subsequent new originations would be securitized under the new fee schedule, providing an incentive to adjust prices immediately. To measure the timing and extent of pass-through to mortgage rates, I analyze the dataset on refinance mortgages originated in a two-year period around the time of the two 2012 guarantee fee increases described in Section 1.5.1, as well as state-day level data on average jumbo and conforming mortgage rates described in the

[12]The actual announcement dates themselves were December 30th, 2011 and August 31st, 2012, so I assume originators respond on the next business day.

same section. In my analysis I first consider the effect of the guarantee fee increase on the primary-secondary spread, and subsequently on the jumbo-conforming spread, both measures that should isolate the effect of guarantee fees.

2.3.1 Primary-Secondary Spread

The next step in my analysis of the effects of guarantee fee increases is to test the extent of pass-through to mortgage rates. However, guarantee fees are only one component of originator costs; by far the larger component is the originator's funding cost. Many originators, especially smaller nonbank lenders, actually borrow the cash required to originate a mortgage from short-term warehouse lenders before securitizing the mortgage and repaying the loan, while larger banks with internal funding sources have their own costs of capital. This is turn raises the question of how to properly measure originator funding costs. In principal, any point on the yield curve with a maturity similar to mortgages[13] would be a reasonable candidate, but for these purposes I assume that originators finance themselves at the secondary market rate. There are two reasons for this choice. First, as discussed in Vickrey and Wright (2013)[186] originators can substantially hedge their risk of interest rate movements by selling their mortgages forward in the TBA market, allowing them to lock in a funding rate based on secondary market yields. Second, while the spread of the primary market rate over the secondary market rate will directly reflect guarantee fees[14], other interest rate spreads will tend to reflect other macroeconomic factors. In particular, because mortgage yields include a premium for bearing prepayment risk, the spread of secondary mortgage yields over a suitable risk-free alternative interest rate will tend to vary with the extent of prepayment

[13] Typically between five and ten years.

[14] Fuster et al (2013)[84] discuss reasons why this is the case, and in their analysis they assume immediate 100% pass-through to the secondary-primary spread.

risk[15]. As evidence for this relationship, I estimate the following model:

$$SECONDARYRATE_t - SWAPRATE_t \quad = \quad MOVE_t\beta_1 + \epsilon_t$$

where the secondary yield ($SECONDARYRATE_t$) and 10-year swap rate ($SWAPRATE_t$), the instrument typically used to construct the yield curve for durations similar to 30-year MBS, are measured on a weekly basis as 3-week centered moving averages and the Merrill Lynch Options Volatility Index ($MOVE_t$), a measure of interest rate volatility, is measured similarly. Interest rate volatility functions in this case as a proxy for prepayment risk, as an increase in volatility increases the likelihood that interest rates will fall below the cutoff required for borrowers to profitably refinance. The results of this regression, shown in Figure A6 of Appendix E, show a clear correlation between the secondary-swap spread and interest rate volatility, as variation in the MOVE index explains 43% of the variation in the spread[16]. Similar problems would affect the choice of any other zero curve rate, especially in light of the fact that, as discussed in Section 1.2.4, Federal Reserve policy over this time period exerted a strong influence on long-term interest rates. Thus in what follows I measure the effect of guarantee fees on the secondary-primary spread, assuming as in Fuster et al (2013) [84] that this spread reflects only guarantee fees and the effects of any other supply frictions at the primary-market level.

Motivated by the assumptions stated above, I estimate empirical loan-level models of

[15] Recall that owing to their guarantee agency-guaranteed mortgages effectively have no credit risk, or all of the credit risk has been converted to prepayment risk. Hence the spread of the secondary rate over a risk-free rate should not reflect a default risk premium.

[16] Other market-specific demand and liquidity factors could potentially explain the rest of this spread.

mortgage interest rates of the following form:

$$
\begin{aligned}
SPREAD_i &= POST_i\beta_1 + RISK_i\beta_2 + MACRO_i\beta_3 + e_i \\
SPREAD_i &= POST_i\beta_1 + [POST_i \times CASHOUTREFI_i]\beta_2 + RISK_i\beta_3 + MACRO_i\beta_4 + e_i \\
SPREAD_i &= POST_i\beta_1 + [POST_i \times HHI_i]\beta_2 + RISK_i\beta_3 + MACRO_i\beta_4 + e_i \\
SPREAD_i &= POST_i\beta_1 + [POST_i \times NONBANK_i]\beta_2 + RISK_i\beta_3 + MACRO_i\beta_4 + e_i
\end{aligned}
$$

where i indexes mortgages, $SPREAD_i$ measures the spread of the interest rate on mortgage i over the average current-coupon secondary market yield, $POST_i$ is an indicator for whether the mortgage was originated after the fee increase announcement, $RISK_i$ are mortgage-specific risk factors and $MACRO_i$ are market-specific risk factors. The sample of mortgages used in each case is the set of all agency fixed-rate refinances originated in a period from six months prior to the first announcement to six months after the second announcement (July 2011 to March 2013). In subsequent models, I interact the post-announcement indicator with either an indicator for whether the loan is a cash-out refinance ($CASHOUTREFI_i$), with the Herfindahl index for agency mortgage refinance originations for the borrower's state (HHI_i), or with an indicator for whether the lender originating the loan is a nonbank mortgage lender ($NONBANK_i$)[17]. These interaction terms are intended to shed light on the mechanisms governing pass-through from guarantee fees to mortgage interest rates, and some discussion of each is warranted.

First, as discussed in Fuster et al (2013) [84] and briefly in Section 1.2.4, a potential determinant of the primary-secondary spread is putback risk, or the risk that originators may have to repurchase defaulted mortgages from the GSEs under a stress scenario. As compensation for this risk, lenders may charge additional interest rate premiums to

[17] I identify nonbank mortgage lenders using the lender's name. I only observe the name for lenders originating a sufficient quantity of loans; lenders making less than roughly 1,500 loans over the full sample, or roughly .02% of the total, are coded as "Other". In the main specification reported here, I treat these lenders as banks, although the results are robust to excluding them from the study.

risky borrowers, even if they intend to securitize the mortgage, in case the loan is later put back to them. I proxy for the default risk factor using an indicator for cash-out refinance, as these tend to be the refinance product with the highest default risk, and interact this term with the policy change indicator. Second, prior research such as Scharfstein and Sunderam (2013)[178] argues that pass-through to mortgage interest rates is attenuated in more concentrated mortgage markets due to imperfect competition. As a measure of market concentration and hence the degree to which imperfect competition impedes pass-through, I interact the policy change with the Herfindahl index, a measure of market concentration, for mortgage lending in that state and year. While one might argue that the Herfindahl index interaction presents an endogeneity problem, in that areas with a few low-cost firms will tend to have higher concentration ratios as the more efficient firms dominate the market, the fact that the guarantee fee increases applied equally to all lenders and affected their cost structure in the same way suggests that this is not a major concern. Finally, as discussed in Black et al (2010)[39] and Loutskina and Strahan (2009),[149] lenders who are more reliant on securitization should respond more to changes in secondary market conditions. Those authors focus on the sources and costs of funds lenders use to originate new mortgages, and show that credit supply tends to respond more for lenders with low levels or high costs of capital. Among lenders I observe in the data, banks, credit unions and large insurance companies should be better-equipped due to their size and capital base to hold mortgage loans on balance-sheet. Non-bank mortgage lenders, by contrast, do not have access to deposits or insurance premia as a stable funding source, borrow from more expensive warehouse lenders, and only originate loans with the intention of securitizing them; thus, they should be more exposed to increases in guarantee fees. As such, origination costs should rise by relatively more for nonbank mortgage lenders than for other institutions, so I interact the policy change with an indicator for whether the lender is a nonbank mortgage specialist.

Table 2.5: Refinance Interest Rate Event Study

Post-Jan. 2012	.1043***	.0992***	.1064***	.1020***
	(1.58E-3)	(1.59E-3)	(2.20E-3)	(1.59E-3)
Post-Sept. 2012	.0847***	.0628***	.1720***	.0809***
	(5.98E-4)	(6.21E-4)	(1.50E-3)	(6.22E-4)
FICO	-.0017***	-.0017***	-.0017***	-.0017***
	(6.13E-6)	(6.09E-6)	(6.12E-6)	(6.12E-6)
LTV	.0029***	.0029***	.0029***	.0029***
	(1.33E-5)	(1.33E-5)	(1.33E-5)	(1.33E-5)
DTI	.0011***	.0011***	.0011***	.0011***
	(2.10E-5)	(2.09E-5)	(2.10E-5)	(2.10E-5)
Cash-Out Refi.	.0775***	-.0084***	.0771***	.0775***
	(5.35E-4)	(1.25E-3)	(5.35E-4)	(5.35E-4)
Post-January 2012 x Cash-Out	-	.0720***	-	-
	-	(1.35E-3)	-	-
Post-Jan. 2012 x Cash-Out	-	.1564***	-	-
	-	(1.20E-3)	-	-
Post-Jan. 2012 x HHI	-	-	-.0268***	-
	-	-	(6.54E-3)	-
Post-Sept. 2012 x HHI	-	-	-.3449***	-
	-	-	(5.65E-3)	-
Post-Jan. 2012 x Non-Bank	-	-	-	.0114***
	-	-	-	(9.10E-4)
Post-Sept. 2012 x Non-Bank	-	-	-	.0224***
	-	-	-	(1.25E-3)
N	2,077,099	2,077,099	2,077,099	2,077,099
R^2	0.2717	0.2792	0.2732	0.2724

Notes: Dependent variable is spread of primary mortgage rate over secondary mortgage rate, sample is all loans originated between September 2011 and March 2013 in GSE dataset. Standard errors in parentheses. Additional controls include origination balance, interest rate volatility, state house price indices and unemployment rates, and indicators for investor-owned, third-party origination, and condominium properties. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table 2.5 presents the estimated coefficients and standard errors from these regressions. Model 1 includes only indicators for whether the mortgage was originated after the announcement date, while Models 2, 3 and 4 interacts these terms respectively with the cash-out refinance indicator, state-level HHI, and nonbank indicator. The estimates from Model 1 suggest that in the long-run, guarantee fees pass through substantially to mortgage interest rates. If pass-through were complete, the estimated coefficients would be .1 (10 bp); the mean estimates in Model 1 imply 100% pass-through following the first guarantee fee increase and roughly 85% pass-through following the second. It is possible given the evidence from Figure 2.2 that since guarantee fees take some time to become fully incorporated, the smaller estimate for the September 2012 announcement reflects the fact that the post-period sample is not long enough to fully measure the extent of pass-through. The estimates from Model 2 suggest that pass-through was more complete for cash-out refinances, and taken literally, suggest that mortgage rates for cash-out refinances increased by more than the amount of the guarantee fee increase. This somewhat curious result may reflect originators charging an additional markup over the amount of the guarantee fee increase to cash-out refinance borrowers to compensate for putback risk. Model 3 indicates that the extent of pass-through was lower in more concentrated markets, as we would expect if mortgage interest rates were set via oligopolistic competition. This result accords with the evidence from Scharfstein and Sunderam (2013) [178] that pass-through from MBS yields to mortgage rates and refinancing activity is attenuated in highly-concentrated markets. The effect of market concentration on the extent of pass-through is considerably stronger for the second guarantee fee increase than the first. This provides another potential explanation for why measured pass-through from the second guarantee fee increase is somewhat smaller in Model 1: because of the greater effect of concentration[18], the combined effect pooled across all markets is reduced. Finally, the estimates for Model 4

[18]The estimates, literally interpreted, suggest zero pass through for markets with an HHI above .5 after the September 2012 increase.

suggest that nonbank mortgage lenders increased their interest rates by more than did conventional mortgage lenders. These results corroborate those from Black et al (2010) [39] and Loutskina and Strahan (2009), [149] and are as one would expect given nonbank lenders' greater exposure to increases in the cost of securitization. Across all models, the effects of other covariates are much as one would expect, with higher-FICO borrowers receiving lower rates and higher-LTV and -DTI and cash-out refinance borrowers receiving higher rates. Overall, these results are broadly supportive of the assumption of essentially complete pass-through in the long-term from guarantee fees to mortgage rates[19] and the findings from other papers regarding the effects of increased market concentration and secondary market reliance on mortgage pricing.

In some ways, the high estimated degree of pass-through from guarantee fees to mortgage interest rates is difficult to square with existing stylized facts regarding the degree of imperfect competition in mortgage markets. As discussed in Scharfstein and Sunderam (2013), [178] mortgage markets in the U.S. are typically rather concentrated, with the average household living in a county with a four-firm concentration of roughly 30%. In a standard model of imperfect competition, such a degree of concentration would result in a lower rate of pass-through from costs to prices, and indeed Scharfstein and Sunderam (2013) [178] find evidence for such reduced pass-through in the mortgage market. However, while my results from Table 2.5 suggest that pass-through is lower in more concentrated markets, they also indicate a higher degree of pass-through in the aggregate than we should expect given overall levels of market concentration. Nevertheless, my results in some ways accord well with prior estimates. Scharfstein and Sunderam (2013) [178] in fact estimate[20] complete pass-through from secondary market yields to primary market interest rates for markets with a four-firm concentration ratio

[19] For example, Fuster et al (2013) [84] and DeFusco and Paciorek (2014) [64] both assume 100% pass-through.

[20] The authors estimate in Table 7, Panel A a pass-through rate of $1.349 - 2.956 \times CR4$ and in Panel B of $1.349 - 2.956 \times CR4$, implying 100% pass-through for markets with four-firm concentration ratios below 11.8% or17.2%, respectively.

below, depending on the specification, roughly 12 to 17%. Moreover both Fuster et al (2013)[84] and DeFusco and Paciorek (2014)[64] assume 100% pass-through from guarantee fees to mortgage interest rates in their analysis. These findings could potentially be explained by a low degree of pass-through from idiosyncratic costs to prices but a high degree of pass-through of aggregate cost increases. As discussed in Miller et al (2015),[164] the degree of industry-pass through depends variously on both idiosyncratic and aggregate shocks, and in their empirical setting the authors find that aggregate cost shocks are typically passed-through completely to prices. Because the increase in guarantee fees I analyze applied nearly equally to all firms, with only minor heterogeneity across firms during the prevailing market environment in the ability to avoid GSE securitization by holding originated loans in portfolio, it is not unreasonable to expect that it should be entirely passed-through to consumers.

While there could in theory be other explanations for the measured increase in the primary-secondary spread following the announcement of guarantee fee increases, some simple diagnostics support the validity of this regression design. Importantly, there is little evidence of a compositional shift post-announcement in the credit risk characteristics of originated mortgages that could be driving these results. Table A14 in Appendix E summarizes average risk characteristics by month and shows that the characteristics of new refinance originations were essentially identical over this entire period. Moreover, in addition to observable characteristics, there is little evidence for a compositional shift on unobservables: the correlation of the regression residuals from Model 1 in Table 2.5 with month of origination is just -.0131, suggesting that unobserved credit risk factors did not vary systematically over this period. Finally, the results presented in Table 2.5 are robust to a number of alternative specifications, including measuring the effect of the guarantee fee increase on interest rates paid on conforming purchase loans as well as refinances.[21]

[21] The estimates from those robustness checks are shown in Table A15 in Appendix E.

2.3.2 Jumbo-Conforming Spread

Though the results from Table 2.5 suggest that over a medium-term horizon guarantee fees are substantially passed through to mortgage rates, they offer scant evidence on the timing of this pass-through. Because the FNMA and FHLMC data are observed at a monthly frequency, it is impossible to measure short-term pass-through using this microdata. As an alternative, I consider intertemporal variation in the spread of average jumbo mortgage rates over average conforming mortgage rates at the state level. Prices for jumbo mortgages, which are too large to qualify for the agency guarantee, should not tend to reflect changes in originator costs that affect only conforming mortgages; hence any difference between the two should reflect agency-market-specific factors only, such as guarantee fees. To test what effect guarantee fees had on conforming mortgage rates, I estimate the following empirical model:

$$JUMBO_{i,t} = \alpha_i + CONFORMING_{i,t}\beta + \Sigma_{s=1}^{S} \mathbf{1}[DATE_t = s]\delta_s + e_{i,t}$$

where the unit of observation (i, t) is a state-day pair, α_i is an unobserved state-specific factor, β measures the tendency of conforming and jumbo mortgages to co-move, and the coefficients δ_s measure variation in the jumbo-conforming spread across days. Figure 2.4 plots the estimated values of δ_s by trading day in a 100-day window around the two announcement dates.[22] In each case, the pre-trend in the two weeks leading up to the announcement is relatively flat and hovers around 0, following which the jumbo mortgage rate falls relative to the conforming mortgage rate over the next several weeks by about 10 bp[23]. Thereafter the estimated fixed effects hover around -10 bp. This decline is plausibly driven by market factors that increase the spread of conforming

[22] I use 5-day centered moving averages for jumbo and conforming mortgage rates to smooth out noise in the series.

[23] There is a clear spike in jumbo mortgage rates visible in the January 2012 series roughly 3-4 weeks after the announcement date. This spike coincides exactly with the month of February, and I suspect it is mis-coded data.

mortgage rates over some base interest rate while having no effect on jumbo mortgage rates. The increase in guarantee fees on these dates is a reasonable candidate for such a factor, especially in light of the results from Table 2.5 showing an increase in the spread of conforming mortgage interest rates over the secondary rate, and FNMA's own analysis attributes much the decline in the jumbo-conforming spread to changes in guarantee fees.[82] The estimated extent of pass-through, as in Table 2.5, is effectively 100% pass-through after three weeks in both cases. The estimated decline in jumbo rates relative to conforming rates appears to persist over longer time horizons as well. Table A16 in Appendix E estimates a similar regression on indicators for post-announcement using a longer 6-month window before and after the fee increase, and the results again confirm that jumbo mortgage rates fell relative to conforming rates by roughly the full amount of the guarantee fee increase (10 bp)[24]. As discussed in Section 2.3.1, this high degree of pass-through even in a relatively concentrated industry is sensible if we assume the guarantee fee increase applied equally to all market participants. Combined with the results from Table 2.5, these estimates provide suggestive evidence for essentially complete pass-through from guarantee fees to conforming mortgage interest rates.

2.4 Effects on Credit Constraints

After considering the effect of the policy change on interest rates, I next study the subsequent effect on credit constraints, and in particular the volume of new mortgage originations. For these purposes I treat propensity to refinance as an indicator of credit supply, as among other reasons refinance mortgages provide a much clearer "base rate"than do home purchase loans: the set of potential customers for refinance loans is all individuals with mortgages, which I observe for a certain market segment, but the set

[24]The point estimates imply pass-through of between 100% and 115%, or between 10 and 11.5 bp. In both cases the difference between the point estimate and 10 bp is roughly one standard error.

Figure 2.4: Jumbo-Conforming Loan Spread Fixed-Effects Estimates

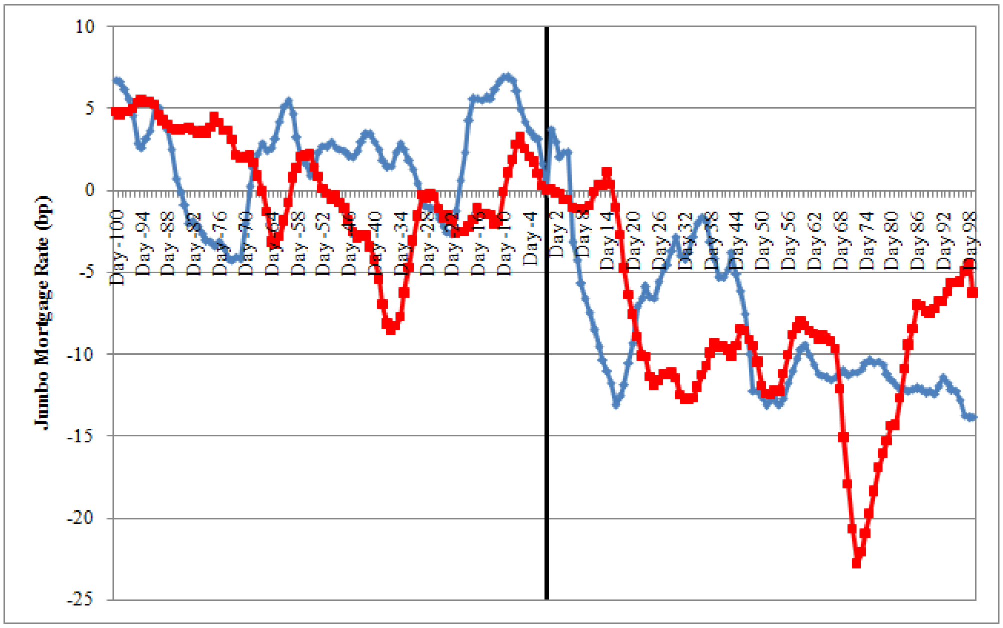

Notes: Estimated daily fixed effects for jumbo-conforming spread event study. January 2012/September 2012 estimated fixed effects shown in Blue/Red. Black bar denotes event date.

of potential customers for new purchase loans is unobserved.[25] Therefore in what follows most of the discussion will center on refinance mortgages rather than home purchase mortgages. Finally, building on the findings from several previous studies[26] documenting a connection between refinancing credit constraints and default, I analyze the effect of increased guarantee fees on default behavior. To measure the effect of guarantee fee increases on refinancing and default, I employ the full dynamic dataset of loans that were ineligible for HARP originated between June 2009 and June 2013, as described in Section 1.5.1. In my analysis I consider in turn the effect of the GSE's policy changes on refinancing rates and default behavior, and conclude with a discussion of the magnitude of these effects relative to prior estimates.

[25]I will discuss the assumptions required to interpret changes in refinancing activity as reflecting "credit supply".

[26]See Keys et al (2014),[136] Laderman (2012),[143] Krainer and Laderman (2011),[139] or Tracy and Wright (2012)[185] for examples.

2.4.1 Refinancing Behavior

Given the findings from Section 2.3 that guarantee fee increases result in higher mortgage rates, the natural next question is what effect this has on origination volume. If demand for mortgage credit is not perfectly inelastic, the increase in interest rates (prices) should result in lower volumes (quantity). One challenge, however, is that it is difficult to determine what demand for mortgage credit would have been in the absence of the guarantee fee increase; in particular, the set of potential home-buyers is unobserved, so changes in the volume of home purchase loans may reflect either changes in the number of home buyers or changes credit supply. To avoid this problem, I use panel data on refinancing decisions, which allows me to define a "base rate" of potential refinance borrowers. I then treat changes in the propensity to refinance as an indicator of credit supply.

This approach presents several crucial endogeneity concerns. First, on the supply side, their may be other factors coincident with the guarantee fee increase that affect access to refinancing credit. Indeed, an expansion to the HARP program implemented in January 2012 (HARP 2.0) made it significantly easier for certain borrowers with little or no home equity to refinance. The start of this program coincides exactly with the first announced guarantee fee increase, complicating the use of this date for an event study. As such, I restrict attention to a sample of borrowers who did not qualify for HARP refinances, or those with loans originated after June 2009[27]. There were no other policy changes that would affect refinancing credit supply around this time, and the hope is that in the short-term, other credit supply factors were essentially fixed[28]. Second, on the demand side, there may be factors correlated with the borrower's value of refinancing

[27]The HARP program contains several other cutoffs for eligibility even among GSE-backed mortgages, but the other cutoffs, which are based on LTV and delinquency status, are manipulable and likely to raise additional endogeneity concerns.

[28]I also control for month-to-month changes in interest rates, which do vary over this short time period. In particular, the third round of Federal Reserve asset purchases (QE3) was announced in September 2012, although that should be controlled for via interest rates.

that vary systematically between the two guarantee fee regimes. In my empirical model, I control for most of the key observed variables determining the borrower's value of refinancing: changes in interest rates, home equity, and local unemployment rates. However, there may be unobserved variables that affect the value of refinancing; in particular, Mian and Sufi (2011) [159] highlight unobserved local-area income growth expectations as an unobserved factor correlated with household refinancing decisions. I include state-level fixed effects to control for fixed differences across geographic areas in the hope that income growth expectations do not vary considerably in over the time horizons used for this analysis. More troublingly, however, unobserved household expectations may be systematically correlated with unemployment rates and house prices, as they would be in a typical macroeconomic model. This potential endogeneity issue makes the inclusion of house prices and unemployment rates as control variables problematic. To address this concern, I instrument for house prices using housing-supply-elasticity instruments and for unemployment rates using Bartik instruments that should be exogenous to income growth expectations[29]. Finally, in my data I observe only the rate of prepayment, which includes both refinance and sales. Therefore, in order to interpret changes in prepayment rates as indicative of changes in refinancing rates I must assume that the propensity to sell for this group is held fixed over the sample period. Since the borrowers in this sample have relatively new loans originated just two years prior, the rate of home sales should be low for the group as a whole. While I control for loan age to address the issue of sales, any increase in the rate of sales over the sample period due solely to the passing of time would bias my estimates of the effect of guarantee fees upward. As such, if there is an upward trend over time in home sales that is not proxied via house prices and loan age, my estimates of the effect of guarantee fee increases in refinancing will represent a lower bound of the true effect, and a negative estimated effect on prepayment would still reflect a negative

[29]The full details of the process of constructing these instrumental variables is explained in Appendix B.

shift in credit supply.

My empirical model of prepayment takes the following form:

$$REFI_{i,t} = POST_{i,t}\beta_1 + EQ_{i,t}\beta_2 + AGE_{i,t}\beta_3 + RISK_i\beta_4 + \Delta RATE_{i,t}\beta_5 + e_{i,t}$$

where the unit of observation (i, t) is a borrower-month, $REFI_{i,t}$ is an indicator for borrower i prepaying her mortgage in month t, $POST_{i,t}$ is an indicator for whether month t is after the guarantee fee increase announcement date, $EQ_{i,t}$ is the borrower's home equity and $AGE_{i,t}$ the age of their loan, $RISK_i$ are fixed borrower risk characteristics observed at origination, and $\Delta RATE_{i,t}$ is the change in market interest rates since borrower i took out her last mortgage. As noted above, I instrument for home equity and include the first-stage predicted value as a regressor[30].

Table 2.6 shows the estimated coefficients and standard errors from this regression. Models 1 and 2 respectively use a sample of observations from a three-month or six-month window around the January 2012 fee increase announcement date, while Models 3 and 4 use similar samples around the September 2012 announcement date. The estimated coefficients are largely as one would expect. Less risky borrowers, measured as those with either higher credit scores or more home equity, are more likely to prepay, perhaps reflecting their greater access to refinancing credit, as are borrowers with older loans. Borrowers whose initial mortgage carried an interest rate with a high spread over the average are more likely to prepay, as they can derive greater value from refinancing, and borrowers are also more likely to prepay if the prevailing national average mortgage rate has declined since they received their initial loan. This specification omits certain variables, including instrumented local unemployment rates, indicators for initial mortgage purpose, house type and origination channel. Table A17 in Appendix E shows the estimated coefficients when these variables are included, and the results are

[30]The main results are all robust to estimating the model with ordinary least squares using the observed variables rather than instrumental variables, as shown in Table A19 of Appendix E.

Table 2.6: Refinancing Probability Event Study

	Model 1	**Model 2**	**Model 3**	**Model 4**
Post-Announcement	-0.0016***	-0.0028***	-0.0013***	-0.0032***
	(0.0000)	(0.0001)	(0.0001)	(0.0001)
FICO	8.03E-05***	6.58E-05***	8.52E-05***	8.25E-05***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\widehat{\textbf{Equity}}$	2.44E-07***	2.14E-07***	2.89E-07***	2.75E-07***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Loan Age	0.0003***	0.0004***	0.0004***	0.0004***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Initial LTV	0.0014***	0.0013***	0.0018***	0.0017***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\widehat{\textbf{Unemp. Rate}}$	-1.86E-4	4.74E-3***	-3.17E-3***	3.78E-3***
	(2.21E-4)	(1.24E-4)	(2.27E-4)	(1.30E-4)
Initial Int. Rt. Spread	0.0142***	0.0121***	0.0111***	0.0115***
	(0.0002)	(0.0001)	(0.0002)	(0.0002)
Δ Mtg. Rt.	-0.0281***	-0.0191***	-0.0208***	-0.0191***
	(0.0002)	(0.0001)	(0.0001)	(0.0001)
State FE?	Yes	Yes	Yes	Yes
Lender FE?	Yes	Yes	Yes	Yes
Announcement Date	Jan. 2012	Jan. 2012	Sept. 2012	Sept. 2012
Sample Window	3 Months	6 Months	3 Months	6 Months
N	29,872,035	52,340,967	36,057,155	63,184,479
$\mathbf{R^2}$	0.0755	0.0660	0.0893	0.0850

Notes: Dependent variable is indicator for prepayment, sample is all loans originated after June 2009 in GSE dataset. Standard errors clustered at individual-level in parentheses. Hats denote instrumented variables. ***/**/* indicates significance at 99%/95%/90%-level respectively.

substantially the same[31]. Moreover my results are also very similar quantitatively when I estimate the model with ordinary least squares rather than instrumental variables, as shown in Table A19 of Appendix E.

Turning to the estimated effect of the fee increase, each of the four models suggests that prepayment probabilities declined following the guarantee fee increase. The estimates suggest that the short-term effect at 3-month horizons was somewhat smaller than the medium-term effect at 6-month horizons and that the effects of the two guarantee fee increases were similar in magnitude. The estimates with additional controls in Table A17 of Appendix E are also similar in magnitude. Under my identifying assumptions and treating the estimates from Table A17 as a baseline, I estimate that the effect of a 10 bp increase in guarantee fees is to reduce the probability of refinance by between .17% and .24%, or between 2% and 2.6% on an annualized basis (CPR)[32]. While a decline of just 2% CPR in refinancing may not seem economically significant, on a base level of 27.5% CPR, which is roughly the sample average, this corresponds to a decline of between 7.3% and 9.5% in refinancing rates. Moreover, the quantity of outstanding agency FRMs is so large that even a small change in the probability of refinancing entails a rather large change in origination volumes. Applying the estimates from Models 1 and 3 from Table A17 in Appendix E to the total quantity of agency FRMs outstanding shown in Figure A7 of Appendix J generates an implied decline in refinancing volume of $7.19 billion per month from the first guarantee fee increase and $10.01 billion per month from the second increase, for a total decline in mortgage refinancing volume of $206.4 billion per year. As such, my results suggest that

[31] Appendix E also contains estimates using the implementation date as a regressor, shown in Table A18. The estimated effect of the April 2012 guarantee fee increase are similar at three-month horizons, but at six-month horizons they are substantially larger, possibly because the post-period in this case includes months from after the September announcement, and the estimate has conflated the two. Moreover the estimated effect at three-month horizons for the December implementation date is zero while the estimate at a six-month horizon is negative, which is exactly what would result if the effect had actually begun three months earlier in September. As such, I interpret these results as suggesting that the announcement date is when the effects on credit supply begin.

[32] Mortgage prepayment rates are typically quoted as Conditional Prepayment Rates (CPR) rather than Single Monthly Mortality (SMM), where CPR = $1-(1-\text{SMM})^{12}$.

increases in guarantee fees have economically significant effects on refinancing activity.

My estimates of the effect of guarantee fees on refinancing also accord well with prior estimates. In a simulation exercise, DeFusco and Paciorek (2014) [64] apply their estimated mortgage demand elasticity to a hypothetical increase in guarantee fees of 11 bp, which they assume is fully passed-through to mortgage rates. They estimate a change in mortgage credit demanded of between 0.17% and 0.22% as a result of this policy change, with an implied figure for $\frac{\partial \text{Demand}}{\partial IR}$, with interest rates measured in percentage points, of 2.2-3.1%[33]. I estimate a pass-through rate from guarantee fees to mortgage interest rates of between 85% and 100%, or a .085% to .1% increase in mortgage rates, and a change in mortgage credit demand of between .17% and .24%, corresponding to a $\frac{\partial \text{Demand}}{\partial IR}$ elasticity of 1.7-2.8%, indicating my estimates suggest a change in demand for refinancing credit that is very much in line with their estimates. While DeFusco and Paciorek (2014) [64] consider demand for home-purchase mortgages rather than refinance loans, Tracy and Wright (2012) [185] use a proportional-hazard model to estimate a more comparable figure for $\frac{\partial \text{Demand}}{\partial IR}$. Their estimated elasticity of between 2.3% and 2.9%[34] almost exactly matches my own. The close correspondence between my own estimates of the effect of an effective increase in mortgage rates on mortgage credit demand and prior estimates suggests that my own results are a reasonable estimate of the effect of the GSE's policy change.

As with my results from Table 2.5 in Section 2.3, I also estimate an event-study type model to demonstrate the timing of transmission from guarantee fee increases to

[33] The authors write that "Under these assumptions, our estimates imply that the proposed increase in the g-fee would reduce the total dollar volume of fixed- rate conforming mortgage originations by roughly 0.17 to 0.22 percent relative to what it otherwise would have been," and "As an example, our preferred estimates imply that an increase in the mortgage rate from 5 percent to 6 percent - 100 basis points - would lead to a decline in first mortgage demand of 2 to 3 percent, which strikes us as a reasonably small but plausible estimate."

[34] The authors estimate prepayment hazard coefficients for interest rates (in percentage points) of 1.000229 for borrowers above 80% LTV and 1.000288 for borrowers below 80% LTV, suggesting that a 100 bp increase in rates should reduce prepayment by 2.3% to 2.9%.

refinancing. This model takes the form:

$$
\begin{aligned}
REFI_{i,t} &= \Sigma_{s=1}^{S} \mathbf{1}[DATE_t = s]\delta_s + EQ_{i,t}\beta_2 \\
&+ AGE_{i,t}\beta_3 + RISK_i\beta_4 + \Delta RATE_{i,t}\beta_5 + e_{i,t}
\end{aligned}
$$

where all variables are defined as before, and the coefficient of interest are the δ_s terms measuring month-by-month changes in refinance rates. These fixed-effects are plotted in Figure 2.5. While there is some month-to-month variability in refinance rates preceding both the January and September 2012 guarantee fee increases, neither series shows a clear pre-trend, and to the extent that there is a downward trend prior to the September 2012 announcement, it appears to be driven by the effects of the January 2012 announcement, as the sample windows over which these fixed-effects are estimated overlap. Moreover directly following the fee increase, in both cases we see a substantial reduction in refinance probabilities over the following six months of roughly 5.8% CPR or roughly .5% per month. The timing of the observed decline in refinancing rates, combined with the lack of a clear pre-trend, suggests that it was indeed the increase in guarantee fees which was responsible for the decline.

2.4.2 Default Rates

Building on the results from Sections 2.3 and 2.4.1 that increases in guarantee fees tighten household credit constraints, the final step in my analysis is consider how these credit constraints affect default behavior. Several previous studies, including Keys et al (2014),[136] Laderman (2012),[143] Krainer and Laderman (2011),[139] Agarwal et al (2015),[12] and Tracy and Wright (2012)[185] have documented a connection between credit supply and default. Each of these studies shows that increasing the supply of refinancing credit can result in lower default rates (or vice-versa), as for instance would be the case in an option-theoretic forward-looking model of refinancing and default such

Figure 2.5: Refinancing Probability Fixed-Effects Estimates

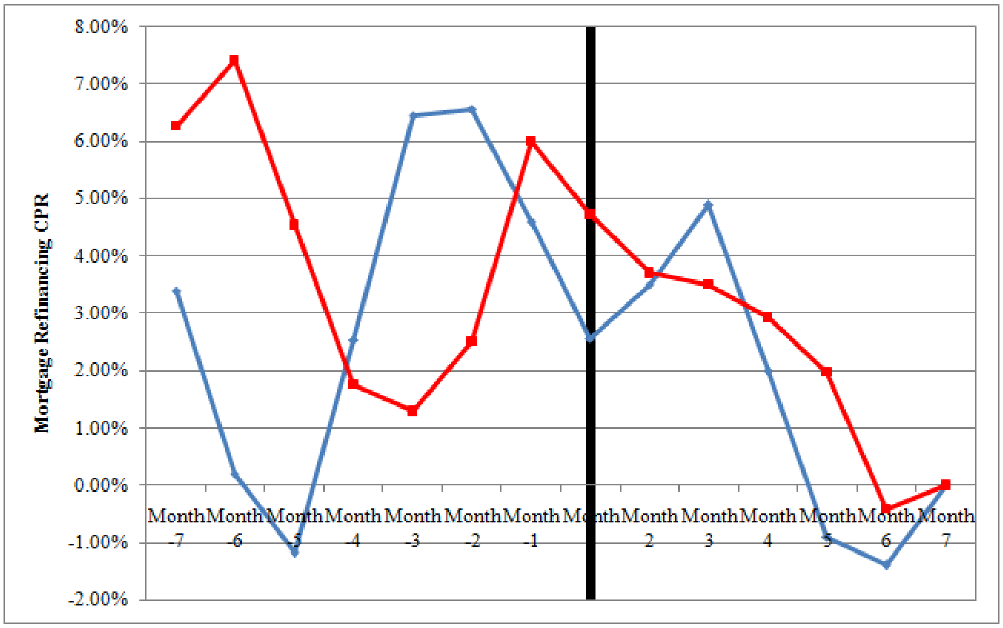

Notes: Estimated daily fixed effects for refinancing probability event study. January 2012/September 2012 estimated fixed effects shown in Blue/Red. Black bar denotes event date.

as Bajari et al (2013) [22] when the refinancing option becomes less valuable. To test what effect the change in credit supply induced by higher guarantee fees has on default, I estimate the following model:

$$DEFAULT_{i,t} \quad = \quad POST_{i,t}\beta_1 + APP_{i,t}\beta_2 + AGE_{i,t}\beta_3 + RISK_i\beta_4 + MACRO_{i,t}\beta_5 + e_{i,t}$$

where $DEFAULT_{i,t}$ is an indicator for whether borrower i defaulted in month t[35], $APP_{i,t}$ is the percent home price appreciation since the borrower initially took out ther mortgage, and $MACRO_{i,t}$ are state-specific or other aggregate factors affecting default such as unemployment rates or mortgage interest rates.

The estimated coefficients and standard errors from this regression are presented in Table 2.7. My estimates suggest that default probabilities increased marginally following the guarantee fee increase announcement dates. The estimated effects over a 6-month

[35] I define the month of default as the month directly following the final month in which a borrower made a payment for the final time before defaulting. For example, if a borrower makes a final payment in January, then makes no further payments before being declared in default in April, I date the default month to February.

Table 2.7: Default Probability Event Study

	Model 1	Model 2	Model 3	Model 4
Post-Announcement	4.04E-5***	1.17E-4***	4.65E-5***	1.11E-4***
	(8.02E-6)	(5.84E-6)	(6.36E-6)	(9.29E-6)
FICO	-2.70E-6***	-2.48E-6***	-2.80E-6***	-2.58E-6***
	(8.67E-8)	(6.61E-8)	(6.02E-8)	(7.71E-8)
$\widehat{\text{Home Price}}$ App.	-5.83E-4**	4.52E-5	-8.24E-4***	-4.26E-4***
	(2.96E-4)	(1.66E-4)	(7.81E-5)	(1.10E-4)
Loan Age	8.17E-6***	9.18E-6***	5.60E-6***	4.75E-6***
	(3.95E-7)	(2.83E-7)	(2.55E-7)	(3.32E-7)
Initial LTV	4.80E-6***	4.33E-6***	4.57E-6***	4.10E-6***
	(1.73E-7)	(1.35E-7)	(1.21E-7)	(1.57E-7)
$\widehat{\text{Unemp.}}$ Rate	-1.66E-4***	1.26E-5	-5.28E-5***	8.94E-6
	(2.96E-5)	(1.52E-5)	(9.67E-6)	(2.27E-5)
Initial Int. Rt.	4.89E-5***	5.37E-5***	6.29E-5***	7.90E-5***
	(8.96E-6)	(6.85E-6)	(6.15E-6)	(8.07E-6)
Δ Mtg. Rt.	3.30E-5	1.97E-4***	2.80E-4***	5.25E-4***
	(6.01E-5)	(1.40E-5)	(1.51E-5)	(3.56E-5)
State FE?	Yes	Yes	Yes	Yes
Lender FE?	Yes	Yes	Yes	Yes
Announcement Date	Jan. 2012	Jan. 2012	Sept. 2012	Sept. 2012
Sample Window	3 Months	6 Months	3 Months	6 Months
N	29,872,035	52,340,967	36,057,155	63,184,479
R^2	0.0006	0.0007	0.0005	0.0006

Notes: Dependent variable is indicator for default, sample is all loans originated after June 2009 in GSE dataset. Standard errors clustered at individual-level in parentheses. Hats denote instrumented variables. Additional controls include initial DTI ratio, state-level instrument home value indices, indicators for owner-occupied, third-party origination, initial cash-out refinance, initial purchase loan, condominiums, manufactured housing, and planned-unit developments. ***/**/* indicates significance at 99%/95%/90%-level respectively.

window are considerably larger than those for the 3-month window, suggesting that the full effect on default takes some time to manifest. This is sensible, since any impact of guarantee fees on default would have to come through removing the refinancing option, and households who are initially denied credit may continue paying their mortgage while attempting to refinance on multiple occasions. The other coefficients are as one would expect, with higher default rates for borrowers with low home price appreciation since origination, high-LTV loans, high initial interest rates or high prevailing interest rates, although the estimated effect of unemployment on default is inconsistent across specifications[36]. The results from Table 2.7 are also robust to estimating the model with ordinary least squares rather than instrumental variables, as shown in Table A20 of Appendix E. Taken together with the results from Sections 2.3 and 2.4.1, these results suggest that the tightening of credit constraints associated with increases in guarantee fees lead to higher rates of mortgage default, as we would expect if credit-constrained borrowers with high and burdensome monthly mortgage payments turned to default as a second-best alternative.

My estimates for the effect of tightened credit constraints on default are also both quantitatively reasonable and economically significant. Applying the estimated changes in the probability to the average monthly default probability in the data for this sample from 2012[37] implies an increase in the probability of default of between 20.8% and 60.4%. While this figure may seem somewhat high, it accords well with prior estimates. Agarwal et al (2015) [12] estimate that as a result of the HARP program, the probability of refinancing for eligible borrowers increased by .24% per month, while the probability of default declined by .0038% to .0095% per month[38]. These estimates correspond to an

[36] Recall from Section 1.3.1 that the effect of unemployment on default is a contentious issue in the literature.

[37] Default is very rare for this sample in 2012, partly because these loans were originated from 2009 to 2011 during a period of tight credit and hence represent a less risky sample, and partly because the loans are relatively young. Hence the average default rate for the sample in 2012 is .02% per month, or .24% CPR.

[38] Agarwal et al (2015) [12] state that "We find that a one percentage point absolute increase in the ex-ante share of eligible loans for HARP is associated with an increase of about 0.24 percentage points in the fraction of loans that refinance under the program... [and] a reduction of about 0.38 basis points in the average zip

estimated range of $\frac{\partial P(\text{Default})}{\partial P(\text{Prepay})}$ of between 1.58% and 3.95%, while my own estimates from the main post-announcement coefficients of Models 1 and 3 in Tables 2.7 and A17 suggest a figure for $\frac{\partial P(\text{Default})}{\partial P(\text{Prepay})}$ of 1.94% to 2.38%, indicating that our results are very much in line. This correspondence between estimated effects of guarantee fee increases and the HARP program provides further suggestive evidence that guarantee fees affect default primarily by closing off the refinancing channel for certain borrowers.

The results from Sections 2.3 and 2.4 demonstrate that guarantee fee increases result in tighter mortgage credit supply conditions. Following the announcement of guarantee fee increases, originators increase their offered mortgage rates almost completely, passing though between 85% and 100% of their increase in costs depending on the specification. The extent of pass-through appears to be greater for riskier borrowers and for less-well-capitalized lenders, but lower in less competitive markets. As a result of this equilibrium supply-side response, credit supply tightens, and the propensity of borrowers to refinance falls as a result. My estimates of the magnitude of this effect, which largely align with previous work, suggests that total refinancing volume declined by roughly $205 billion as a result of the two 2012 guarantee fee increases. As suggested by prior studies, I also estimate that guarantee fee increases and the attendant tightening of credit constraints is associated with an economically-meaningful increase in the probability of default.

2.5 Conclusion

The analysis presented in this chapter traces the effect of changes in guarantee fees on the agency mortgage market and on credit supply in particular. The results from Section 2.2 suggest that in the short-term, increases in guarantee fees lead originators to

code mortgage interest rate," while their estimates imply that a 1 bp change in average area interest rates translates to a .01% to .025% change in the probability of default. Hence a 1% change in the share of eligible borrowers results in a change in the probability of default of between .0038% and .0095%

accelerate their securitization schedules and reduce excess servicing spreads to try to avoid the effects of the fee increase. However, I find that holding interest rates fixed, the prices of newly-originated MBS fall, resulting in lower revenues for originators. The results from Section 2.2 indicate that originators respond to this potential decline in revenues in predictable ways. I find that nearly all of the fee increase is passed through to mortgage interest rates, although in line with previous research, I find that this affect is stronger in more competitive lending markets, for higher-risk borrowers, and for nonbank lenders. I shown in Section 2.4 that as a a result of this increase in interest rates, the probability of refinance decreases and the probability of default increases for affected borrowers. The magnitudes of these effects are of potential macroeconomic significance, with an estimated annual reduction in refinancing volume of $205 billion. While my documentation of these effects in response to guarantee fee increases is novel, my results largely confirm prior studies of mortgage demand elasticity and the effects of credit constraints on household default behavior. From discussions with industry participants as well as from the assumptions made in several prior studies,[64,84] it would appear that, novel though they may be, all of these results align with common sense as well.

A key limitation of my analysis of changes in guarantee fees is that the effects of such a policy change are dependent on the GSE's market share and hence not regime-invariant. At the time of the policy changes analyzed in this study, agency securitization accounted for 60% of all new first-lien mortgages, with an even greater market share of 30-year fixed-rate mortgages, and effectively zero private-label securitization.[97] As indicated both by prior research[39,149] and by my own results from Section 2.3, the effect of changes in secondary-market conditions will depend on the extent to which lenders are reliant on the secondary market as a funding source. Given that a high GSE market share reflects widespread reliance on GSE securitization, the strong measured effects of guarantee fee increases during this period are not surprising. However, it is unlikely my findings regarding the extent of pass-through or the effect of guarantee fees on credit constraints would continue to hold quantitatively during periods in which private-label securitization or other origination options are on an equal footing with GSE securitization. Since large, discrete, and exogenous changes in guarantee fees are relatively rare events, I cannot evaluate this claim directly. However, I would hesitate to generalize these finding across market conditions. As an alternative, in Chapter 4 I attempt to estimate the effect of changes in guarantee fees on mortgage credit supply in a regime-invariant fashion in order to make quantitative claims about the out-of-sample effects of counterfactual GSE pricing policies.

My results in this chapter are nonetheless of potential significance for macroeconomic, monetary and housing finance policy. First, the size of the estimated change in refinancing volume, $205 billion per year, is large enough to warrant consideration of the macroeconomic effects of this policy change. Second, an analysis of the effects of guarantee fees should be taken into account in predicting the effects of changes in short-term interest rates or of other GSE policies on refinancing volume. Finally, we might suspect in light of the estimated effects of guarantee fees on credit supply that even given a congressional mandate to raise additional revenue, the optimal policy for a housing finance regulator would seek to preferentially recoup losses with higher fees when credit constraints are least likely to bind. In Chapter 4 I design a model in order to simulate the effects of such a policy directly and tie those results back to those from this chapter.

Chapter 3

Unemployment and the Value of Refinancing Credit

3.1 Introduction

In this chapter I analyze how the value of access to refinancing credit depends on employment status. Many macroeconomic models assume [13] and several empirical studies suggest [103] that as a result of differences across borrowers in the marginal value of cash-on-hand, the value of access to credit also differs across borrowers. An important dimension of heterogeneity in this regard is income or employment status. However, because access to credit depends on employment status, measuring the direct effects of unemployment rates on refinancing activity will not entirely capture the effects of borrower preferences. Indeed, levels of almost every type of household borrowing declined during the 2008 recession, and while it is difficult to disentangle the demand effect from variations in access to credit, some evidence suggests that the supply effect of tighter credit constraints was the driving factor[1]. The goal of this chapter is to estimate in a reduced-form manner how much the value of credit access varies across borrowers, with a particular focus on how it varies with unemployment. In combination with the results from Chapter 2 regarding the effects of guarantee fees on credit constraints, the estimates from this chapter will inform the structural model I develop in Chapter 4 by providing motivation for how credit supply policies can have different impacts on employed and unemployed borrowers.

Building on the insight from Section 1.4 regarding the interaction between liquidity preference, or demand for refinancing credit, and credit supply in determining refinancing outcomes, in this chapter I employ several distinct approaches to address the question of who benefits from expansions in mortgage credit supply. First I analyze the related but slightly different question of whether the value of refinancing credit is higher for the unemployed than for the employed. As noted in prior research, [51,123] consumption-smoothing motives could increase the value of cash-on-hand, and hence

[1]See Brown et al (2013) [45] for a discussion of recent trends in consumer indebtedness and Gropp et al (2014) [102] for evidence of credit supply effects.

refinancing, for borrowers with temporarily depressed incomes, such as the unemployed. I develop a simple model to demonstrate the intuition that the relationship between the value of refinancing and unemployment can be directly estimated from the data, at least up to the sign of the effect, among populations for whom credit access is not negatively affected by employment status. I then focus on two populations, borrowers with considerable home equity wealth and borrowers who are eligible for HARP refinances, and estimate the sensitivity of their refinancing behavior to local unemployment rates. I interpret the results, under the framework of my simple model, as indicative of the covariance between unemployment and the value of refinancing credit. Second I use a difference-in-differences design to analyze how takeup under the HARP program differed across high- and low-unemployment areas. I exploit sharp discontinuities in eligibility for the HARP program and a discrete one-time change in the efficacy of the program, as discussed in Section 1.2.4, to generate exogenous variation in access to credit for some but not all borrowers. I interpret variation in the marginal propensity to refinance under HARP across borrowers as evidence that credit supply policies have disproportionate impact on certain borrowers, namely those for whom takeup is highest.

My results indicate both that preferences for refinancing credit are higher among the unemployed, and, potentially as a consequence, that credit supply policies such as HARP disproportionately benefit the unemployed. Using data on refinancing applications, I find that the probability of being credit constrained varies positively with unemployment rates for the population overall. However, my results indicate that this effect differs greatly depending on other credit characteristics, and in particular, that among borrowers with large home equity positions unemployment has no effect on credit constraints. Similarly I estimate that for the population as a whole, the probability of refinancing is declining in local unemployment rates, although among borrowers with substantial home equity the two vary positively. I further develop this insight by exploiting the design of the HARP program, which provides a unique opportunity to examine how borrowers refinance when credit constraints are removed or severely attenuated. I find that among this unconstrained population, a 1% increase in local unemployment rates leads to a 3.4% to 4.4% increase in refinance CPR, with borrowers in the top decile of the unemployment-rate distribution refinancing at speeds of 34.3% CPR higher than those in the bottom decile. While in the context of the model from Section 1.4 these results suggest that the value of refinancing credit is higher for the unemployed, they offer no indication of how the benefits of credit supply policies are distributed. It could be the case, for example, that although the unemployed have stronger preferences for refinancing credit, policies such as HARP do little to alleviate credit constraints among this population. However, my baseline estimates of takeup under the HARP program suggest that this is not the case, and that that the program both substantially stimulated refinancing and that takeup increased substantially with unemployment rates. I estimate that following a change to the HARP program in 2012 that dramatically increased its effectiveness, the probability of refinancing increased 26.9% to 31.7% more for eligible borrowers than for similar ineligible borrowers. I also find that a 1% increase in local unemployment rates leads to a 3% to 7% increase in program takeup, and while I find the largest differences between borrowers in low-unemployment areas and borrowers in moderate-unemployment areas, differences in program takeup are monotonic across the unemployment distribution, with borrowers in the highest quartile of unemployment rates 40% to 50% more likely to refinance under HARP than those in the lowest quartile. These estimates suggest that the HARP program disproportionately benefited the unemployed, although I cannot address whether this is due to greater demand for refinancing credit among the unemployed or a greater reduction in the credit constraints facing the unemployed.

This chapter makes several distinct contributions both to the academic literature regarding the interaction between unemployment, liquidity preference and credit constraints, and to the policy literature concerning the effectiveness of the HARP program. First, this paper is among the first to provide direct empirical evidence on the covariance between unemployment and demand for credit. While some previous studies, such as Sullivan (2008) [183] and Bethune (2014), [35] do make use of micro data to analyze the effect of unemployment on credit utilization, my analysis extends their results in two key ways. One, I study the effect of unemployment on demand for a type of credit, home equity borrowing, that is both much a much larger market than the types of credit those authors analyze[2] and that prior research has identified as the most important mechanism in the U.S. for consumption-smoothing-motivated precautionary saving. [50] Two, my analysis controls directly for credit constraints in a way that prior studies have not by studying populations of borrowers for whom credit constraints should be minimal. While my two-stage approach to estimating the effects of unemployment on both access to and utilization of credit bears some similarities with the work of Antoniades (2015) [19] and Jiménez et al (2012, 2014), [125, 126] my methodology is new to the literature, and to my knowledge no prior research has used the HARP program as an instrument for credit availability. Second, on the policy side, to my knowledge this is the first study to analyze how takeup under the HARP program varied across borrowers, in itself an interesting policy question in its own right. These estimates may be useful in the design and evaluation of other housing finance or credit supply policies, and as noted in previous research [1] my finding that the value of easing credit constraints varies positively with unemployment (or negatively with income) has important implications for the design of stabilization policy. On the theoretical side, while a rather large class of macroeconomic models with endogenous credit constraints assumes that the value of credit access varies with income[3], relatively few papers are

[2]Both studies focus on unsecured consumer credit, largely consisting of credit cards. As noted in Brown et al (2013), [45] over my sample period U.S. total mortgage debt outstanding varied between 7 and 12 times the volume of total credit card debt outstanding.

[3]See Aiyagari (1994), [13] Krusell and Smith (1998), [141] Gourinchas and Parker (2002), [99] and Lustig and Van Nieuwerburgh (2005) [152] for examples.

able to provide empirical evidence either in favor of this assumption or of the magnitude of the effect. The analysis in this chapter is intended to provide empirical support for the microfoundations underlying these models, as well as to inform the policy literature concerning the effects of credit supply policies and the assumptions implicit in the structural model I develop in Chapter 4.

The analysis in this chapter proceeds as follows. In Section 3.2 I outline a simple model of the refinancing process based on that presented in Section 1.4 and use the model to generate several hypotheses relating unemployment to the value of refinancing credit and subsequently to objects that I observe in the data. Section 3.2.1 uses data on refinance applications to identify groups of borrowers for whom credit supply should be insensitive to unemployment, and Section 3.2.2 builds on these results by analyzing how such borrowers' refinancing speeds covary with unemployment. Likewise Section 3.3 analyzes the refinancing behavior of another group of borrowers, those eligible for the HARP program, who as a result of the program were even less impeded by unemployment-based credit constraints. Finally in Section 3.4 I analyze using a difference-in-differences design how takeup under the HARP program varied across borrowers in high- and low-unemployment-rate areas. Section 3.5 concludes with a discussion of how my results inform housing finance policy analysis, the academic macroeconomic literature, and the structural model I develop in Chapter 4.

3.2 Supply and Demand Interactions in the Effect of Unemployment on Refinancing

In this section I provide descriptive evidence on how supply and demand effects interact in the agency refinancing market and discuss what implications these have for the relationship between unemployment and the value of refinancing credit. There are two goals for these descriptive regressions. The first is to provide evidence in favor of a model of borrower behavior similar to that in Aiyagari (1994).[13] In the context of such a model, the shadow value of easing credit constraints should be greatest for households facing negative labor income shocks, and indeed, that the value of refinancing credit itself is greater for the unemployed. However, the results from Section 1.4 suggest that identifying such an effect will be complicated by credit constraints. The second goal is describe the patterns in the data that will identify the structural model I develop in Chapter 4, as will be discussed in more detail in Section 4.3.

To help frame these results, as well as those to be presented in Sections 3.3 and 3.4, consider a more simplified version of the model presented in Section 1.4. Borrowers with an employment status u, where $u = 1$ denotes unemployment, and other characteristics X derive value from refinancing $V_R(X,u)$, and prefer refinancing to continuation only if that $V_R(X,u) + \epsilon \geq 0$, where ϵ is a preference shock with symmetric distribution Λ. Hence the probability that refinancing is optimal is given by $P_{\text{Refi. Optimal}}(X,u) = \Lambda(V_R(X,u))$. Borrowers are observed to refinance only if refinancing is optimal and they can access credit, where the probability that they can access credit is given by $1 - P_{\text{Const}}(X,u)$. Thus the probability that a borrower with characteristics (X,u) is observed to refinance is given by:

$$P_{\text{Refi.}} = P_{\text{Refi. Optimal}}(X,u) \times (1 - P_{\text{Const}}(X,u))$$

with first derivative:

$$\begin{aligned}
\frac{\partial P_{\text{Refi.}}}{\partial u} &= \frac{\partial P_{\text{Refi. Optimal}}(X,u)}{\partial u} \times (1 - P_{\text{Const}}(X,u)) - P_{\text{Refi. Optimal}}(X,u) \times \frac{\partial P_{\text{Const}}(X,u)}{\partial u} \\
&= \underbrace{\Lambda'(V_R(X,u))}_{\geq 0} \frac{\partial V_R(X,u)}{\partial u} \times \underbrace{(1 - P_{\text{Const}}(X,u))}_{\geq 0} - \underbrace{P_{\text{Refi. Optimal}}(X,u)}_{\geq 0} \times \frac{\partial P_{\text{Const}}(X,u)}{\partial u}
\end{aligned}$$

The equation above highlights the key econometric problem for identifying $\frac{\partial V_R(X,u)}{\partial u}$: if $\frac{\partial P_{\text{Const}}(X,u)}{\partial u}$ is sufficiently large, it may be that $\frac{\partial P_{\text{Refi.}}}{\partial u}$ and $\frac{\partial P_{\text{Const}}(X,u)}{\partial u}$ have the opposite sign, or that unemployed borrowers are not more likely to refinance than employed borrowers despite valuing the refinancing option more. In this section, I provide evidence for how $\frac{\partial P_{\text{Refi.}}}{\partial u}$ and $\frac{\partial P_{\text{Const}}(X,u)}{\partial u}$ vary with other credit-risk characteristics X and attempt to interpret these results as evidence in favor of an Aiyagari [13]-type model, while Sections 3.3 and 3.4 approach the same question from alternative angles.

My empirical approach is to estimate the quantities $\frac{\partial P_{\text{Refi.}}}{\partial u}$ and $\frac{\partial P_{\text{Const}}(X,u)}{\partial u}$ directly from the data. I estimate $\frac{\partial P_{\text{Refi.}}}{\partial u}$ using the matched dataset described in Section 1.5.4. This dataset consists of 6,911,395 monthly loan-level panel records for 200,000 FRM borrowers from California from 2000 to 2012, and provides detailed credit-risk and location information on the borrower as well as their choices of when to refinance and into what product. To estimate $\frac{\partial P_{\text{Const}}(X,u)}{\partial u}$ I employ the HMDA dataset described in Section 1.5.2, which consists of 17,109,796 applications for conforming refinance mortgages from California between 2000 and 2012. This dataset also provides detailed credit-risk and location information, as well as information on whether the loan application was denied or whether the borrower rejected an extended loan offer, both of which I treat as indicative of credit constraints.

3.2.1 Credit Supply Models

The first step in my analysis is to measure how credit constraints vary with employment status. My goal is to identify groups of individuals for whom $\frac{\partial P_{\text{Const}}(X,u)}{\partial u}$ is zero or close to zero. The intuition is that among these groups, we should be able to infer how the value of refinancing credit varies with unemployment based solely on how actual refinancing varies with unemployment, at least up to the sign of the effect. Motivated by this intuition, I estimate probit models of the probability that a borrower is credit-constrained of the following form:

$$P(\text{Const}_{i,t} = 1) \quad = \quad \Phi(\text{Risk}_{i,t}\gamma_1 + \text{Unemp}_{i,t} \times \text{Risk}_{i,t}\gamma_2 + Z_{i,t}\gamma_3)$$

where $\text{Const}_{i,t}$ is an indicator for whether the borrower was credit-constrained, or applied for a refinance loan unsuccessfully, $\text{Risk}_{i,t}$ are credit-risk characteristics such as home equity, LTV, or FICO, $\text{Unemp}_{i,t}$ is the local-area unemployment rate, and $Z_{i,t}$ are other controls such as interest rate spreads. I estimate the model using the HMDA dataset, and in my baseline specifications treat both denied applications and rejected offers as indicative of credit constraints[4]. I treat rejected offers as indicative of credit constraints because, as discussed in Section 1.4, they represents cases in which the borrower cannot obtain affordable credit, or is effectively constrained by prices rather than outright denial. The key parameters of interest are γ_2, which measure the interaction between unemployment and other credit risk factors, as these parameters will identify groups for whom the effect of unemployment on credit constraints is negligible.

A key question is why we should expect the effect of unemployment on credit constraints to ever be negligible. Lenders may be less concerned by the employment status of borrowers who, by virtue of their equity position or FICO, are otherwise good credit risks, in part because they discount the borrower's ability to repay, and hence employment status, if they have sufficient home equity. Two market trends suggest that this is indeed the case. First, as discussed in Finberg (2003),[81] a series of financial fraud cases in the mid-2000s suggests that a popular predatory lending strategy is to target borrowers with high home equity but low income for large cash-out refinances in hope that the borrower defaults and the lender can foreclose on the house. Second, there is a product called a reverse mortgage, described extensively in Shan (2011),[179] that is especially designed for borrowers with no income but abundant home equity, and until very recently these products were not underwritten at all and lenders did not even verify their applicant's income or credit history. These two examples suggest that in certain cases, particularly in cases where the applicant's home is worth substantially

[4]My results are robust to only using application denials, as shown in Table A23 in Appendix E.

more than their mortgage debts, lenders simply do not care whether the borrower has an income or not, as the worst-case scenario for the lender entails foreclosing on the home and selling it at a profit.

For the purposes of identifying the true effect of employment on both credit constraints and unemployment, the probit model written above faces two key endogeneity problems. First, there are potentially unobserved variables correlated with both credit supply and borrower demand for refinancing credit that may themselves be correlated unemployment rates or credit risk. Mian and Sufi (2011) [159] note that in a permanent-income framework, households expecting higher future income should borrow against it and spend, with the associated consumption boom pushing up home prices and employment. As such, refinancing activity and home prices, and thus home equity, should be positively correlated with unobserved income growth expectations while unemployment should be negatively correlated with the same. To address this concern, I instrument for house prices and unemployment using measures that are unrelated with local-area income growth expectations[5]. In my baseline specifications, I instrument for home prices using the housing-supply-elasticity instrument from Gyourko et al (2008) [106] and for unemployment using the instruments from Bartik (1991), [25] which measure the change in unemployment relative to a base year that would be predicted by national-level trends[6]. Second, as discussed in Section 1.5.4 the evidence from Gyourko and Tracy (2014) [107] suggests that using a market-wide unemployment rate as a proxy for individual-specific employment status results in attenuation bias in the estimated effect of employment status, or that the true effect of unemployment is much larger in magnitude than would be suggested from using market-level aggregate probabilities. While I do not have access to data on individual-level employment status with which to address this concern, I do not think it

[5]My results are robust to using observed values for unemployment rates and home price appreciation, as shown in Table A22 in Appendix E.

[6]Details on how these instruments are constructed are contained in Appendix A.

invalidates my results, as my estimates thus represent a lower-bound on the true effect of employment status on credit constraints and refinancing, as the bias should be towards zero. Moreover, my estimates in this section and in Section 3.2.2 use place-level unemployment rates rather than MSA-level as in Gyourko and Tracy (2014).[107] Places are much smaller than MSAs and provide a better signal of of employment status, allowing me to distinguish between, for example, Beverly Hills and Compton, both part of the Los Angeles MSA, where the average unemployment rates over my sample period are 5.4% and 13.4% respectively.

My estimates of the effect of unemployment rates on credit constraints across various credit-risk bins are shown in Figure 3.1. I divide the sample into five equally-sized bins based on three credit-risk factors: FICO, DTI, and LTV. I then estimate the effect of unemployment on credit constraints within each bin via a probit regression. Each model also include controls such as LTV, DTI and FICO for credit risk and mortgage interest rate spreads and guarantee fees to control for secondary market conditions[7]. Figure 3.1 displays the estimated coefficients on unemployment with standard error bars, where lower quintiles (Q1) of risk factors correspond to higher credit risk, as measured by a low FICO score, high DTI ratio, or high LTV ratio, and higher quintiles (Q5) correspond to lower credit risk. The mean effect of unemployment on the probability of being credit-constrained, as we might expect, is significantly positive and marked with a horizontal line. However, this mean effect masks considerable heterogeneity, as the positive effect of unemployment on credit constraints is concentrated among borrowers in the lowest three quintiles of creditworthiness. For borrowers with LTV ratios below 60% (Q3-Q5), DTI ratios below 34% (Q5), or FICO scores above 800 (Q5), the estimated effect of unemployment on credit constraints is negative, suggesting that credit constraints do not tighten with unemployment for these borrowers. While the estimated negative effect may seem counterintuitive, because these estimates combine the effect on

[7]The estimated coefficients for these controls are not shown, but the full set of parameter estimates including these covariates is shown in Table A21 in Appendix E.

Figure 3.1: Credit Constraint Probit Unemployment Effects by Credit Risk

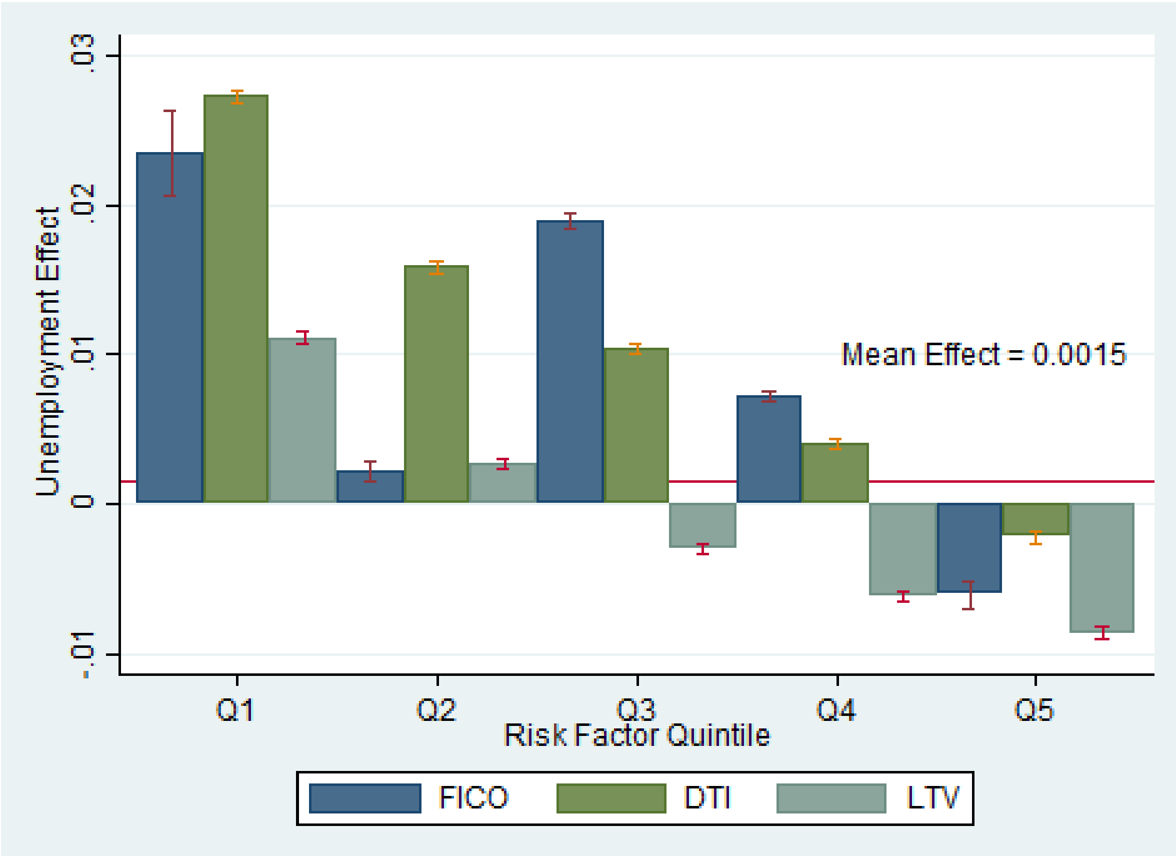

Notes: Estimated effect of unemployment on probability of credit constraints by quintiles of credit-risk factors. Blue/Green/Grey bars measure effect of unemployment within quintiles of FICO/DTI/LTV respectively. Higher quintiles denote lower credit risk (increase FICO, decreasing DTI, LTV). Red line shows mean effect across all applicants. Bars indicate standard error bands.

both application denial and offer acceptance this result may indicate that unemployed borrowers have high liquidity preference and are thus more willing to accept high-interest-rate offers[8]. Moreover the effect of unemployment on credit constraints is essentially monotonic in credit risk, suggesting that lenders preferentially deny their unemployed high-credit-risk applicants but approve their unemployed low-credit-risk applicants.

The estimates from Figure 3.1 are robust to several alternative specifications. First, Table A21 in Appendix E provides additional parameter estimates from the same model

[8]As evidence for this interpretation, the estimates using only application denial shown in Table A23 of Appendix E are only negative for borrowers with very low LTV ratios, although the overall pattern of attenuation in the effect of unemployment as credit risk descreases is similar to that from Figure 3.1.

as Figure 3.1 that shed further light on this finding. The full set of parameters largely conform to expectations, as borrowers are more likely to be credit constrained if they have low income, a low FICO score, or a high DTI ratio, and borrowers are also less willing to lend when lending is less profitable due to high guarantee fees or low mortgage interest spreads over swap rates. When I estimate the model with ordinary least squares rather than instrumental variables, as shown in Table A22 in Appendix E, the estimates are largely as before, with the effect of unemployment attenuating sharply as borrowers become more credit-worthy, and negative for borrowers with FICO scores above 740. When I estimate the same model using only application denial as an indicator of credit constraints, shown in Table A23 of Appendix E, I observe the same pattern as in Figure 3.1, as the effect of unemployment on denial declines nearly monotonically in creditworthiness. A key difference with these results relative to my baseline specification is that the effect of unemployment is on average considerably higher and is only negative for borrowers with LTV ratios below 53%. These findings suggest that unemployed borrowers may be less willing t reject higher-interest rate offers. However, as discussed previously, I believe that the best indicator of credit constraints is any outcome which results in the borrower not obtaining credit, and since this includes situations in which the borrower is priced out of the market and thus rejects credit offers, I view my results with respect to a combination of denial and rejection as a baseline. Moreover the fact that the effect of unemployment on application denial is positive except for applicants with very low LTV ratios confirms the intuition discussed previously with respect to lender incentives. Applicants with low LTV ratios have homes worth substantially more than the balance on their mortgage, and it is for these borrowers that we should expect lenders to be unconcerned with the applicant's ability to repay the loan. In combination, the results from Tables A21 and A23 suggest that there are certain borrowers, specifically borrowers with strong home-equity positions, for whom the effect of their employment status on credit constraints is negligible.

3.2.2 Refinancing Models

The next step in my analysis is to examine how the sensitivity of actual refinancing rates to unemployment varies across credit risk groups. The results from Section 3.2.1 suggest that among borrowers with abundant home equity, or potentially high credit scores, $\frac{\partial P_{\text{Const}}(X,u)}{\partial u}$ should be negative or close to zero. Among such borrowers, the signs of $\frac{\partial P_{\text{Refi.}}}{\partial u}$, which I can estimate from the data, and $\frac{\partial V_R(X,u)}{\partial u}$, which is the object of interest, should be the same. Building on this insight, I estimate probit models of the probability of refinance with the following form:

$$P(\text{Refi}_{i,t} = 1) \quad = \quad \Phi(\text{Risk}_{i,t}\beta_1 + \text{Unemp}_{i,t} \times \text{Risk}_{i,t}\beta_2 + X_{i,t}\beta_3)$$

where $\text{Refi}_{i,t}$ is an indicator for refinance, $X_{i,t}$ are controls such as loan age, and all other variables are defined as above. I estimate the model using the matched dataset described in Section 1.5.4. While this dataset contains fewer observations than the full GSE dataset, all of which are from California rather than nationwide, the key advantage it offers relative to the full GSE dataset for probit regressions of this type is that I can observe refinances separately from sales. As such, unlike the results presented in Section 2.4.1, I directly observe refinance as my outcome rather than prepayment, and hence need no additional assumptions on the effect of unemployment rates on home sales in order to identify the parameters of interest. Those parameters are β_2, which measure how the effect of unemployment rates vary with other credit risk characteristics. As in Section 3.2.1, I use place-level variation in unemployment rates and house price appreciation and instrument for these variables using respectively the Bartik (1991)[25] instruments and WRLURI[106] instruments discussed in Appendix A.

My estimates of how the effect of unemployment rates on refinancing differs by credit risk are shown in Figure 3.2. As with Figure 3.1, each bar indicates the mean estimated effect of unemployment within five equally-spaced credit risk bins, ordered by decreasing credit risk, with standard error bars shown as well. I divide the sample into bins based on FICO scores, home price appreciation since the borrower took out their initial mortgage, and home equity. As we should expect, the mean effect of unemployment on refinancing across all borrowers is significantly negative and is denoted by a horizontal line. Consistent with my stylized model of the refinancing process and with the results from Figure 3.1, I estimate that the negative effect of unemployment on refinancing attenuates sharply and nearly monotonically in borrower credit risk. The effect of unemployment is just under half as large for borrowers in the top FICO quintile (Q5, above 800) as for borrowers in the bottom quintile (Q1, below 620) and just over half as large for borrowers in the highest house price appreciation quintile (Q5, greater than 32% growth) as for borrowers in the lowest quintile (Q1, greater than 13% decline). Moreover, the effect of unemployment on refinancing is only negative for borrowers in the bottom two quintiles of home equity (Q1 and Q2, below $120,000), indicating that for borrowers with sufficient home equity $\frac{\partial P_{\text{Refi.}}}{\partial u}$ is positive. In combination with those from Section **??**, these results provide some evidence that the value of refinancing does increase with unemployment.

Other estimates based on the same refinancing probit specification also largely conform to expectations. Table A24 in Appendix E shows the full set of covariates from the same model as in Figure 3.2, and the estimates suggest that borrowers behave as we would expect if they refinanced to maximize their refinancing option value subject to a credit constraint. Borrowers are more likely to refinance if they are more creditworthy, measured as either higher FICO scores or lower DTI ratios, or if the refinancing option is more in the money, measured as either high home equity or large declines in mortgage interest rates. The results from Table A24 are also largely robust to estimating the model with ordinary least squares rather than with instruments for unemployment rates, house price appreciation and home equity, as shown in Table A25 in Appendix E. The results from estimating the model without instruments largely confirm the baseline estimates, as

Figure 3.2: Refinance Probit Unemployment Effects by Credit Risk

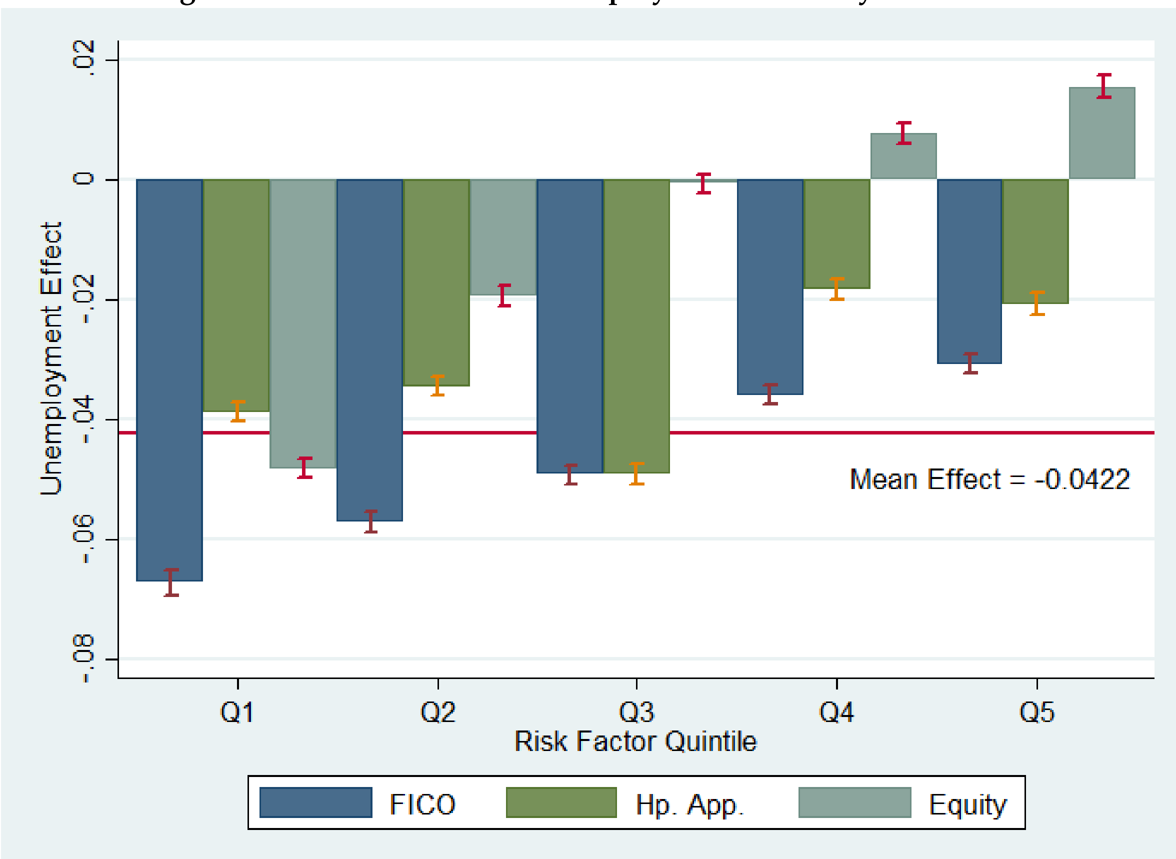

Notes: Estimated effect of unemployment on probability of refinance by quintiles of credit-risk factors. Blue/Green/Grey bars measure effect of unemployment within quintiles of FICO/House-price Appreciation/Equity respectively. Higher quintiles denote lower credit risk (increasing FICO, HP. App., Equity). Red line shows mean effect across all applicants. Bars indicate standard error bands.

I find the same pattern of a negative overall effect of unemployment on refinancing that strongly attenuates as borrowers become more creditworthy, although I estimate that the relationship is not significantly negative only for borrowers in the highest home-equity bin (Bin 5). The general pattern of a negative overall effect of unemployment on refinancing that attenuates in home price appreciation, equity, or FICO is also more clearly visible if we use ten rather than five credit risk categories, as shown in Table A26 in Appendix E. Indeed from Table A26 we can see that borrowers with high levels of either home price appreciation (Bins 8 and 9, between 18% and 32% growth) or home equity (Bins 5 and above, or greater than $120,000) are more likely to refinance, as would be the case if lenders effectively discounted the employment status of high-collateral borrowers who apply for refinancing credit. As with the results from Table A24, I also estimate that the effect of unemployment on refinancing increases essentially monotonically with creditworthiness. These results suggest that for borrowers where the effect of unemployment on credit constraints is negligible, particularly for those with high home equity, the probability of refinancing actually increases in unemployment.

3.2.3 Interpretation of Results

Combined with the stylized model presented previously, the results from Sections 3.2.1 and 3.2.2 provide suggestive evidence regarding the effect of unemployment rates on the value of refinancing credit. The sensitivity of actual observed refinances to unemployment depends on both the value of refinancing and on credit supply, and with how those objects vary with unemployment. As suggested by certain stylized facts regarding mortgage underwriting practices, I estimate that the sensitivity of credit constraints to unemployment, $\frac{\partial P_{\text{Const}}(X,u)}{\partial u}$, is very low for borrowers with low LTV ratios or, analogously, considerable home equity wealth. Although on average I estimate that the effect of unemployment on refinancing activity is negative, for borrowers with such large home equity positions I estimate that the sensitivity of refinancing to unemployment, $\frac{\partial P_{\text{Refi.}}}{\partial u}$, is negative. These results are not driven by either unobserved local-area income growth expectations or by home sales, both of which I can control for. In the context of my model, these results suggest that $\frac{\partial V_R(X,u)}{\partial u} > 0$, or that the value of refinancing credit is increasing in unemployment, as would be the case if liquidity preference was higher for borrowers facing unemployment spells due to consumption-smoothing motivations. However, my supply-demand framework enables me to identify only the sign of $\frac{\partial V_R(X,u)}{\partial u}$, as I do not observe any sub-group for whom credit constraints overall have been removed. Thus in the next section, I analyze the refinancing behavior of another set of borrowers who, by virtue of government policy, are even less likely to be affected by unemployment-based credit constraints.

3.3 Refinancing Behavior of Unconstrained Borrowers

In this section I study how refinancing activity depends on unemployment for HARP-eligible borrowers following the implementation of HARP 2.0. Recall from Section 1.2.4 that the HARP program essentially allowed all eligible borrowers to refinance, particularly following the removal of LTV ratio caps in January 2012. In particular, the program guidelines removed almost all scope for lenders to deny applicants on the basis of their employment status, breaking the dependence of credit constraints on employment status. In the context of the model presented in Section 3.2, this program design implies at a minimum that for HARP-eligible borrowers $\frac{\partial P_{\text{Const}}(X,u)}{\partial u}$ should be close to zero, and a stronger interpretation would be that $P_{\text{Const}}(X,u)$ is close to zero. This cohort provides an ideal laboratory with which to study how unemployment affects the value of refinancing credit, as for eligible borrowers the sensitivity of observed refinancing rates to unemployment should reflect only variation in preferences and not in credit constraints. Building on this insight I analyze the sensitivity of refinancing rates to unemployment using a regression on a subset of the GSE dataset described in Section 1.5.1 consisting only of loans eligible for HARP in the period following the implementation of HARP 2.0.

A key issue for this analysis is how to measure whether a given loan is eligible for HARP, as I do not observe this directly. Prior to 2013, the cutoff for determining eligibility was determined based on the date at which the loan was delivered to FNMA or FHLMC[9]. However, I observe only the origination date, not the date at which the loan was delivered to the GSEs, and hence must make some assumptions about the timeline for delivery. Stanton et al (2014)[181] suggest that the typical GSE securitization takes no longer than 45 days, as originators often borrow the cash to originate new loans for that term and must complete the securitization in order to repay their own funding line. In the context of the GSE dataset, this timeline translates to a one- to two-month cutoff for eligibility; however, for three reasons, in my baseline specification I use a cutoff of three months prior to June 2009 (i.e. pre-March 2009) to measure eligibility[10]. First, I do not actually need to measure eligibility, I simply need a valid instrument for credit access. The chosen three-month cutoff is more conservative than a one-month cutoff in the sense that I should observe a certain number of false negatives as borrowers who are in fact eligible will be labeled as ineligible. However, I can be fairly certain that there are very few false positives: Amromim and Kearns (2014)[17] suggest that 95% of loans that are ever delivered to the GSEs are delivered within four months of origination, so nearly all of the loans I label as eligible will in fact be eligible. If instead I used a one-month cutoff (i.e. pre-June 2009), the instrument would be weaker as my eligibility criteria would be less predictive of actual credit access due to false positives. Second, Fannie Mae actually began originating HARP loans in March 2009, although only a very small number of refinances were completed prior to June. Although the eligibility criteria remained the

[9] In 2013 this rule was modified to base eligibility on the actual origination date, in part to grant more borrowers eligibility.

[10] In certain robustness tests, I show that all of my results are both qualitatively and quantitatively similar using a one-month cutoff.

Table 3.1: Count of Mortgage Originations by Month

Pre-HARP		Post-HARP	
Month	# Originations	Month	# Originations
2008-06	136,928	2009-06	296,716
2008-07	105,094	2009-07	207,276
2008-08	114,085	2009-08	174,923
2008-09	119,994	2009-09	168,540
2008-10	113,360	2009-10	167,701
2008-11	135,936	2009-11	165,077
2008-12	214,947	2009-12	151,407
2009-01	329,514	2010-01	120,959
2009-02	362,719	2010-02	110,024
2009-03	342,493	2010-03	137,321
2009-04	358,264	2010-04	132,821
2009-05	358,234	2010-05	139,886
Total	2,691,568	Total	1,972,651

Notes: Number of unique mortgages originated by month in dataset.

same, it is possible that as a result the better predictor of credit access is whether a loan was delivered to Fannie Mae by March 2009. Finally, as I will show later the clearest differences in prepayment behavior can be observed between loans originated before and after March 2009. In this sense, my definition of eligibility is an empirical one, inferring the presence of eligible borrowers from subsequent refinancing.

Table 3.1 shows the number of unique mortgages contained in the GSE dataset within one year of the deadline for HARP eligibility; in what follows, I refer to this as the "HARP dataset". There is a clear decline in new issuance visible beginning in June 2009. This corresponds to the start of the HARP program, and the decline is likely caused by the exclusion of HARP refinances from the data[11][12]. As a result, the dataset contains

[11] FNMA and FHLMC remove loans that began as HARP refinances from their loan-level data, but loans that terminated via a HARP refinance are included.

[12] Alternatively the difference could be caused by borrowers refinancing prior to the eligibility cutoff. However, because refinances take some time to complete and the program was announced only a few months in advance, and because refinancing credit was extremely tight during this period, it is highly unlikely that a great number of borrowers refinanced just prior to the cutoff in order to remain eligible.

about 36% more originations from the year prior to the implementation of the HARP program than from the year after. However, using the preferred eligibility cutoff date of 3 months prior to the delivery cutoff of June 2009, or originated before April 2009, the sample is roughly balanced. The difference in loan counts between the pre- and post-period within a three-month window around this cutoff date is just 2%, or 4% within a six-month window[13].

To evaluate how refinancing behavior varies with unemployment among HARP-eligible borrowers, I estimate linear regression specifications of the following form:

$$\text{Refi}_{i,t} = \text{Unemp}_{i,t}\beta_1 + X_{i,t}\beta_2 + e_{i,t}$$

where $\text{Refi}_{i,t}$ is an indicator for refinance, $\text{Unemp}_{i,t}$ is the area unemployment rate, $X_{i,t}$ are other controls that proxy for the value of refinancing such as changes in mortgage interest rates and home equity. I estimate the model on a sample of all loans from the HARP dataset originated prior to the April 2009 cutoff during the period from January 2012 to December 2013, when HARP 2.0 was active. As such, I can be relatively certain that borrowers in this sample can refinance under the HARP program with no constraints, allowing me to interpret refinancing activity as indicative of borrower preferences for refinancing credit.

Table 3.2 shows the coefficient estimates from this model. I estimate four separate models: a baseline model (Model 1) including only a limited set of control variables, a model with separate coefficients for each unemployment decile (Model 2) in order to better understand non-linearities in the effect of unemployment, a model using instruments for unemployment and home equity (Model 3), and a model with these same instruments and a number of additional control variables (Model 4). Models 3 and 4 use the same Bartik (1991)[25] and WRLURI[106] instruments for unemployment and

[13] In most specifications I treat a three-month window around the cutoff as the baseline.

house price appreciation respectively as in Section 3.2, but with unemployment rates and house price appreciation measured at the MSA- or state-level as in Section 2.4, as I do not observe finer geographic detail for loans in the HARP dataset[14]. As in Section 2.4, I also do not separately observe refinancing and sales, but I treat trends in prepayment as largely indicative of trends in refinancing. Across all specifications, my results indicate that refinancing activity increases with unemployment rates. The estimates from Model 1, which I treat as a baseline, indicate that a 1% increase in unemployment rates leads to a 4.4% increase in refinance CPR, while the estimates from Model 4, which I treat as a final robustness check, implying a somewhat lower 3.4% increase in CPR. The estimates from Model 2 suggest that this relationship is monotonic and in fact convex in unemployment rates; the difference in prepayment between the sixth decile, with an average unemployment rate of 8.1%, and the first decile, with an average of 5.4% is 7.2% CPR, while the difference between the tenth decile, with an average rate of 10.6%, and the sixth decile is over 30% CPR. In addition, across all specifications the effects of the other control variables are similar and largely in line with intuition, as I estimate that borrowers are more likely to refinance if they have high home equity, if mortgage rates have declined, or if the spread on their initial mortgage over average rates was high.

I perform two additional robustness tests in order to confirm that the positive measured effect of unemployment on refinancing truly comes from demand rather than supply effect. First I estimate the same specification but restricting the sample only to borrowers who were eligible for HARP based on both their mortgage origination date and current LTV ratio. HARP eligibility was restricted to borrowers with LTV ratios above 80%, but I proxy for eligibility with an indicator for whether the borrower's current LTV ratio is above 75%, as current LTV must be estimated based on imputed house prices and hence has some amount of noise, so as with the origination-date-based eligibility criterion I treat this measure solely as an instrument for credit availability[15].

[14] The construction of these instruments is discussed in more detail in Appendix A.
[15] 75% is also almost exactly the sample median, and hence seemed a suitable threshold.

The estimates from this model, shown in Table A27 in Appendix E, are considerably larger than those from Table 3.2, and imply that a 1% increase in unemployment rates leads to a 12.5% to 13.8% increase in refinance CPR. As before, the estimated effect of unemployment rates on refinancing appears to be monotonic, and the larger estimates from Table A27 relative to Table 3.2 imply that the HARP-eligible cohort is largely unconstrained by credit supply. Second I estimate the same models as in table 3.2 but using only a sample of HARP-elgible borrowers observed post-2013. The concern is that if borrowers in higher-unemployment-rate areas were more likely to be constrained prior to the implementation of HARP 2.0, there may be pent-up demand for refinancing credit and hence we may observe an uptick in refinancing even if demand were insensitive to unemployment. The results from Table A28 in Appendix E show the results from this model, and these estimates indicate that pent-up demand does not drive my baseline results, as the positive effect of unemployment on refinancing rates is still roughly as large even a year after the implementation of HARP 2.0. On the whole, my estimates suggest that the value of refinancing credit is increasing in unemployment, with the strongest effects concentrated in the areas with highest unemployment.

The results from this section and from Section 3.2 provide some support in favor of the proposition that unemployed borrowers derive greater value from refinancing credit than do unemployed borrowers. Across all borrowers, I find that the average effect of unemployment on refinancing is negative, but my estimates suggest that this likely reflects credit constraints that prevent the unemployed from refinancing. Among two distinct groups of borrowers, those with high home equity and those eligible for the HARP program, access to refinancing credit is essentially decoupled from employment status. My results indicate that among these groups, refinancing activity actually increases in local unemployment rates, suggesting that the unemployed place higher value on refinancing credit, potentially due to their consumption-smoothing motives. These findings will help inform the model I develop in Chapter 4, in which I allow the

Table 3.2: Refinance Behavior of HARP-Eligible Borrowers

	Model 1	Model 2	Model 3	Model 4
FICO	9.41E-5***	9.41E-5***	9.27E-5***	9.20E-5***
	(2.09E-6)	(2.09E-6)	(2.20E-6)	(2.39E-6)
DTI	-2.40E-4***	-2.38E-4***	-2.57E-4***	-2.46E-4***
	(1.17E-5)	(1.15E-5)	(1.29E-5)	(1.17E-5)
Initial LTV	2.67E-3***	2.67E-3***	2.74E-3***	3.76E-3***
	(2.56E-4)	(2.57E-4)	(2.63E-4)	(3.97E-4)
Int. Rt. Spread	5.54E-3***	5.71E-3***	6.54E-3***	8.00E-3***
	(8.27E-4)	(8.45E-4)	(9.32E-4)	(9.65E-4)
Loan Age	4.75E-7***	4.76E-7***	4.89E-7***	4.93E-7***
	(4.64E-8)	(4.66E-8)	(4.80E-8)	(4.88E-8)
Δ Mtg. Rt.	-6.77E-3***	-6.65E-3***	-6.54E-3***	-
	(1.46E-4)	(1.47E-4)	(1.52E-4)	-
Equity	6.11E-4***	5.88E-4***	5.10E-4***	1.49E-3***
	(2.06E-5)	(2.35E-5)	(2.91E-5)	(9.08E-5)
Unemp. Rate	3.73E-3***	-	1.55E-3***	2.91E-3***
	(5.62E-4)	-	(4.03E-4)	(4.24E-4)
Dec. 2 Unemp.	-	7.37E-4**	-	-
	-	(3.48E-4)	-	-
Dec. 3 Unemp.	-	1.95E-3***	-	-
	-	(6.56E-4)	-	-
Dec. 4 Unemp.	-	3.80E-3***	-	-
	-	(8.95E-4)	-	-
Dec. 5 Unemp.	-	6.34E-3***	-	-
	-	(1.12E-3)	-	-
Dec. 6 Unemp.	-	6.20E-3***	-	-
	-	(1.25E-3)	-	-
Dec. 7 Unemp.	-	9.63E-3***	-	-
	-	(1.66E-3)	-	-
Dec. 8 Unemp.	-	1.15E-2***	-	-
	-	(1.99E-3)	-	-
Dec. 9 Unemp.	-	2.08E-2***	-	-
	-	(3.05E-3)	-	-
Dec. 10 Unemp.	-	3.55E-2***	-	-
	-	(5.29E-3)	-	-
Lender FE?	Yes	Yes	Yes	Yes
State FE?	Yes	Yes	Yes	Yes
Instruments?	None	None	Unemp./Equity	Unemp./Equity/App.
N	24,063,716	24,063,716	24,063,716	24,063,716
R^2	0.1173	0.1175	0.1202	0.1209

Notes: Dependent variable is indicator for prepayment, sample is all loans in GSE dataset originated in 12 months prior to April 2009 observed from January 2012 to December 2013. Standard errors clustered at individual-level in parentheses. Additional controls include include initial loan interest rate and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. Models 2 and 4 use instrumented values for unemployment rates and home equity, and Model 4 also includes controls for combined LTV, age squared, instrumented home price appreciation, and quartiles of mortgage-rate changes as indicators and interacted with changes in mortgage rates. ***/**/* indicates significance at 99%/95%/90%-level respectively.

value of refinancing to vary with employment status. However, a key question remains: who benefits from policies that increase credit supply? It may be the case that while the unemployed derive greater value from credit access, policies that ease credit constraints disproportionately benefit employed borrowers while providing little incentive for lenders to extend credit to the unemployed. In the next section, I address this question by examining how takeup under the HARP program varies with local unemployment rates.

3.4 Takeup Under the Home Affordable Refinance Program

In this section I analyze how takeup under the HARP program varies with local unemployment rates[16]. The goal of this exercise is to understand who benefits from credit supply policies. While the results from Sections 3.2 and 2.4 suggest that unemployed borrowers derive greater value from refinancing than their employed counterparts, they offer no insight into how the benefits of credit supply policies are distributed. Ex ante it is not unreasonable to suspect that programs such as HARP do very little to ease credit constraints facing the unemployed, and hence disproportionately benefit better credit risks. However, if takeup under the program is greater for borrowers in high-unemployment areas, a reasonable conclusion is that policies which increase mortgage credit supply disproportionately benefit the unemployed. If true, this result could be the product of a last-in, first-out style of lending on the part of mortgage originators in which high-credit-risk borrowers are both the first to be locked out of the market as lending conditions tighten and the last to re-enter as

[16] The analysis in this section represents a substantial part of a separate ongoing project with my co-author Sandesh Dhungana, currently entitled, “Unemployment and the Value of Refinancing Credit: Evidence from the Home Affordable Refinance Program.” I thank him for his contributions towards this project and the results from this section.

credit conditions improve. In the context of the model from Section 3.2, an estimate of how takeup under the HARP program varies with unemployment rates will be a product both of how preferences for credit vary with unemployment and how the effect of the program on credit constraints vary with unemployment. To a first approximation, if the HARP program essentially removed all credit constraints, then the latter effect should be based solely on how the tightness of credit constraints varied with unemployment rates prior to the start of the program. Nevertheless, we can interpret heterogeneous changes in the marginal propensity to refinance as indicative of heterogeneous benefits from program implementation, which, in combination with the results from Sections 3.2 and 3.3 will provide evidence in favor of the policies I analyze in Chapter 4

I estimate a difference-in-differences specification using the full HARP dataset in order to measure the extent of takeup under the HARP program. My research design employs two hard cutoffs that determine credit access, as discussed in Section 1.2.4. The first is a cutoff for HARP eligibility based on the loan origination date, while the second the implementation date of HARP 2.0, after which many more borrowers were able to refinance under the program. The control group is a set of borrowers who, by virtue of their later origination dates, were ineligible for the HARP program, and I compare how the sensitivity of refinancing to unemployment differed between eligible and ineligible borrowers before and after the implementation of HARP 2.0. Although prior studies such as Agarwal et al (2015) [12] and Zhu (2012) [192] study the effects of both the original HARP program and HARP 2.0, for several reasons I focus exclusively on HARP 2.0. First, the original HARP program was widely viewed as unsuccessful: relatively few borrowers were able to refinance under the program, and lenders were given substantially more latitude to reject applicants based on credit risk[17]. Second, unlike

[17] Conversations with industry participants suggest that lenders never actually "rejected" candidates based on credit risk under the original HARP program, but by failing to advertise the program and making the process complicated and difficult they achieved largely the same outcome.

Agarwal et al (2015),[12] who use as a control group a sample of loans that were ineligible for HARP by virtue of not being guaranteed by the GSEs, the time-based eligibility criteria I use is almost entirely co-linear with the implementation of the original HARP program[18]. As such, I would essentially have no pre-period for ineligible borrowers with which to compare refinancing rates following the implementation of the original HARP program.

The pre- and post-HARP samples are also roughly balanced on most credit risk characteristics. Table A29 in Appendix G shows the average characteristics of eligible and ineligible loans originated within a 6-month window[19] of the eligibility cutoff with t-statistics for the differences. Note that in what follows, I use the term "eligible" to indicate "originated prior to April 2009". This cutoff date was chosen in part based on the results from Table A30 in Appendix G, which shows, for each month within a six-month window of the HARP eligibility cutoff, the probability that a loan originated in that month has refinanced within a certain length of time. The probability that a given loan has refinanced within 2, 3, or 4 years clearly declines leading into the HARP cutoff date, but the nadir is actually reached in May 2009, which is part of the reason I label loans originated in May as ineligible[20]. As shown in Table A29, on most credit characteristics, there are statistically significant but economically negligible differences between the eligible and ineligible cohorts, with two key exceptions. First, ineligible borrowers tended to receive lower interest rates than eligible borrowers. This difference is a function of the eligibility criteria, which is based only on origination vintage: because mortgage rates declined over the sample period, loans that were originated later

[18]There are advantages and disadvantages to using a time-based versus channel-based eligibility definition. On most credit risk metrics, GSE loans are likely more similar to other GSE loans than to non-agency loans. However, prepayment rates are sensitive to loan age and other cohort effects that I must control for when comparing different origination vintages.

[19]Throughout this analysis, I treat loans originated within 3- and 6-month windows of the eligibility cutoff as the "baseline" sample.

[20]The other eligibility criteria either apply to all loans in the HARP dataset (e.g. guaranteed by the GSEs) or are dynamic variables (e.g. not delinquent, over 80% current LTV) and hence not meaningful at the origination date.

(and were thus ineligible) tended to receive lower rates. Second, ineligible borrowers are far more likely to be first-time home buyers or to take out home purchase loans. These related discrepancies are due to the exclusion of HARP refinances from the GSE dataset. As a result, new mortgages from borrowers who refinanced under HARP, which would all be originated after the start of the HARP program and would thus be ineligible, are absent from the sample of ineligible loans, leading to a lower percentage of refinance loans and proportionately more purchase loans. Barring this selection issue, all other differences in credit characteristics are essentially negligible. To the extent that there are differences the eligible cohort tends to be of marginally worse credit quality, which would tend to bias estimates of the effect of the HARP program downward without proper controls. I control both for these observable characteristics and for cohort effects in all specifications , and do not believe sample selection should be a serious concern.

3.4.1 Difference-in-Difference Estimates of Program Takeup

In this section I employ a differences-in-differences model to estimate takeup under the HARP program. Table 3.3 shows a non-parametric difference-in-differences estimate of how takeup under the HARP program varies with local unemployment rates. For each quartile of unemployment rates, I compare mean CPR for eligible and ineligible borrowers in the year before and after the introduction of HARP 2.0. Although these estimates do not control for other important determinants of refinancing, especially loan age, the main results are all clearly visible. Refinancing rates tend to be highest for borrowers in the lowest unemployment rate quartile (Q1), while within each quartile refinancing rates are higher both for eligible borrowers and after the introduction of HARP 2.0. The difference-in-differences for mean refinancing rates, which I treat as an estimate of program takeup, increases monotonically in unemployment rates, with CPR roughly twice as high for borrowers in the top unemployment category (above 10% unemployment) relative to the bottom category (below 7.5% unemployment).

Table 3.3: Unconditional Difference-in-Differences

Q1 Unemp.	Pre-HARP	Post-HARP	Difference	Diff.-in-Diff.
Ineligible	14.79%	28.28%	13.49%	-
Eligible	21.39%	37.22%	15.83%	2.34%
Q2 Unemp.	**Pre-HARP**	**Post-HARP**	**Difference**	**Diff.-in-Diff.**
Ineligible	11.84%	26.96%	15.12%	-
Eligible	18.13%	36.32%	18.19%	3.07%
Q3 Unemp.	**Pre-HARP**	**Post-HARP**	**Difference**	**Diff.-in-Diff.**
Ineligible	12.08%	24.84%	12.76%	-
Eligible	17.45%	34.35%	16.91%	4.15%
Q4 Unemp.	**Pre-HARP**	**Post-HARP**	**Difference**	**Diff.-in-Diff.**
Ineligible	10.54%	24.79%	14.25%	-
Eligible	15.05%	34.01%	18.96%	4.71%

Notes: Average CPR for subsets of mortgage loans in HARP dataset. "Eligible" and "Ineligible" refer to loans originated in 6-month window before and after April 2009, "Pre-HARP" and "Post-HARP" refer to monthly observations before and after January 2012. "Difference" denotes difference in CPR between loans observed during pre- and post-HARP period, "Diff.-in-Diff." denotes difference in differences between eligible and ineligible loans. "Q1 Unemp." through "Q2 Unemp." refer to quartiles of the unemployment rate distribution, in order of increasing unemployment rates.

The next step in my analysis is to add additional controls to this baseline non-parametric estimate. Before turning to how takeup varies with unemployment rates, I first estimate the effect of the program itself using the following specification:

$$\begin{aligned}\text{Refi}_{i,t}^{Obs} &= \alpha + X_{i,t}\delta + \beta_1\mathbf{1}[t \geq T_{Imp}] + \beta_2\mathbf{1}[\text{Elig}_{i,t}] + \beta_3\text{Unemp}_{i,t} \\ &+ \beta_4\mathbf{1}[\text{Elig}_{i,t}] \times \mathbf{1}[t \geq T_{Imp}] + e_{i,t}\end{aligned}$$

where $\text{Refi}_{i,t}$ is an indicator for prepayment, $X_{i,t}$ are controls, T_{Imp} is the time of program implementation, and $\text{Unemp}_{i,t}$ is the local unemployment rate. In this model, β_1 captures any group-invariant intertemporal variation in propensity to refinance and β_2 measures cohort-specific differences in refinancing, while β_4, the coefficient of interest, measures the effect of program implementation on refinancing rates for the eligible cohort.

Before presenting the results from this specification, a discussion of several endogeneity concerns with this empirical approach is warranted. First and most importantly, there may actually be some residual dependence of credit access on unemployment following the implementation of the HARP program, biasing any estimates of program efficacy. All evidence suggests that the policies implemented with HARP 2.0 removed the incentive for lenders to deny applicants on the basis of income or employment status. However, I control for the possibility that there are certain lenders who did not participate in or fully comply with the HARP program, leading to reduced credit access for borrowers whose original mortgage was from these lenders[21], by including lender fixed-effects. Second, as discussed in Section 2.4, there may also be endogeneity between the refinance decision, house prices, and unemployment rates that would bias estimates of the latter two on the former. In certain specifications, I instrument for house prices and unemployment rates using the same WRLURI [106] and Bartik (1991) [25] and instruments, respectively, as in Section 2.4[22]. A final endogeneity concern is with the assumptions required for identification of a difference-in-differences model, in particular that there is no manipulation of the criteria for program eligibility and that trends for eligible and ineligible borrowers were similar prior to the implementation of HARP 2.0. While there is somewhat of a decline in loan origination just after the June 2009 cutoff, as discussed in Section 3.3 it is highly unlikely that borrowers took out loans in advance of the cutoff in order to remain eligible, and in any event there are no such declines at the chosen cutoff of three months prior to implementation. Moreover in my baseline specification I ignore other eligibility criteria that may be more manipulable, such as delinquency status or current LTV[23], under the assumption that origination date is plausibly exogenous and still a sufficiently strong instrument for credit access. While I cannot test the parallel trends assumption directly,

[21] As discussed in Amromin and Kearns (2014), [17] most borrowers refinanced under HARP via their original lender.

[22] The construction of these instruments is discussed in more detail in Appendix A.

[23] In certain robustness tests I also include an eligibility criteria based on current LTV, and the results are very similar.

in Section 3.4.2 I present evidence regarding the timing of the effect I measure.

Table 3.4 sets out the regression estimates from the difference-in-differences specification. I estimate the specification above on three different samples of loans originated within 3, 6, and 12 months of the eligibility cutoff. In what follows I treat the estimates using the 3-month window as the baseline specification, since borrowers in the eligible and ineligible cohorts should be most-similar on dimensions such as loan age, house price appreciation and changes in mortgage rates that are important determinants of the refinance decision. Across all specifications, the estimates indicate that the effect of the eligibility and post-HARP 2.0 implementation interaction term (corresponding to β_4) is positive and highly significant. The baseline estimates suggest that an eligible mortgage has a 0.76% greater probability of prepaying in any month following the implementation of HARP 2.0, translating to an economically-significant 8.7% increase in CPR. By way of comparison, the difference-in-differences estimates from Agarwal et al (2015) [12] suggest that HARP-eligible borrowers refinanced under the program at a rate of roughly 6.2% CPR[24], so the estimates accord well with theirs[25]. The estimates from the 3- and 6-month window specifications correspond to increased refinancing speeds due to HARP takeup of between 26.9% and 31.7% on a base of roughly 27.5% CPR.

Regarding the other control variables, I estimate relatively small effects for program eligibility (β_1) and essentially no effect for HARP 2.0 implementation (β_2). The latter result indicates on that credit did not loosen significantly for all borrowers in 2012 and beyond, as the HARP 2.0 dummy effectively represents a time fixed-effect on credit supply in this specification. The former result, and the finding that the effect of the original HARP program on prepayment was positive, may represent persistent

[24] The headline figure quoted in their paper is 1.6% per quarter.

[25] There are two potential sources of discrepancy. First, I estimate the effect of HARP 2.0 alone, while the results from Agarwal et al (2015) [12] include the effects of both the original HARP program and HARP 2.0 jointly; to the extent that HARP 2.0 was more effective, my estimates should be correspondingly larger. Second, I use an imperfect instrument for eligibility based on origination date, while their administrative data allows them to infer HARP eligibility directly. This may lead to some attenuation in the estimates, although these results indicate that the first effect clearly outweighs the second.

Table 3.4: Difference-in-Differences Estimate of HARP Program Effect

	3-Month Window	6-Month Window	1-Year Window
Eligible	0.0012***	0.0015***	0.0019***
	(4.47E-5)	(4.34E-5)	(5.42E-5)
Post-HARP	0.0050***	0.0092***	0.0111***
	(3.65E-5)	(3.99E-5)	(3.63E-5)
Post-HARP 2.0	-0.0002*	0.0000	0.0016***
	(8.03E-5)	(9.74E-5)	(1.32E-4)
Post-HARP 2.0 x Eligible	0.0076***	0.0064***	0.0038***
	(5.76E-5)	(6.77E-5)	(7.93E-5)
FICO	5.31E-5***	5.41E-5***	5.34E-5***
	(6.75E-7)	(6.08E-7)	(5.31E-7)
DTI	-1.66E-4***	-1.70E-4***	-1.64E-4***
	(4.57E-6)	(5.52E-6)	(6.89E-6)
LTV	1.16E-3***	1.25E-3***	1.31E-3***
	(6.67E-5)	(7.94E-5)	(9.87E-5)
Int. Rt. Spread	3.41E-3***	5.57E-3***	6.82E-3***
	(9.77E-5)	(1.06E-4)	(1.31E-4)
Unemp.	-1.16E-4***	-1.19E-4***	-1.38E-4***
	(1.18E-5)	(1.20E-5)	(1.34E-5)
WALA	3.29E-4***	3.60E-4***	2.83E-4***
	(5.74E-6)	(7.81E-6)	(7.96E-6)
Δ Mtg. Rt.	-1.23E-2***	-1.04E-2***	-1.07E-2***
	(3.57E-5)	(4.31E-5)	(6.48E-5)
Equity	1.69E-7***	1.88E-7***	2.08E-7***
	(1.00E-8)	(1.13E-8)	(1.36E-8)
Lender FE?	Yes	Yes	Yes
State FE?	No	No	No
N	83,824,948	121,448,185	177,172,126
R^2	0.0623	0.0590	0.0556

Notes: Dependent variable is indicator for prepayment, sample is all loans in HARP dataset. Standard errors clustered at individual-level in parentheses. Additional controls include loan interest rate and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. ***/**/* indicates significance at 99%/95%/90%-level respectively.

unobserved time- or cohort-level fixed effects, but more likely is a function of the fact that the difference-in-differences specification uses HARP 2.0 only, hence any estimate of the effect of eligibility is picking up some effect of the original HARP program. My estimates for other control variables are very similar to those presented in Sections 2.4 and 3.2, as I find that borrowers are more likely to prepay if they have higher FICO scores, lower DTI ratios, higher initial interest rate spreads, more home equity, or if mortgage rates have declined by more since they took out their initial mortgage loan. I estimate two alternative specifications as a robustness test. First, I estimate the same specification using a cutoff for eligibility one month prior to implementation rather than three months prior, or before June 2009 rather than before April. Because not all mortgages are delivered to the GSEs within a month of origination[26], some loans classified as "eligible" by this criterion will actually be ineligible. Second, I estimate an analogous specification using instruments for unemployment rates and home price appreciation. The results of these two specifications, shown respectively in Tables A32 and A33 in Appendix G accord well with the baseline results, as the estimates both of takeup under the HARP program and of the other control variables are comparable quantitatively and qualitatively to the results in Table 3.4.

3.4.2 Heterogeneity in Program Takeup

While the results from Table 3.4 indicate that the extent of takeup under the HARP program is economically significant and of a similar magnitude to prior estimates, the ultimate question is how takeup under HARP varies with unemployment. To address this question, I estimate a difference-in-differences model that allows for heterogeneity in the treatment effect by local unemployment rates. Following Chetty et al (2009),[54] I

[26]Stanton et al (2014)[181] suggest that the process typically takes 30-45 days.

estimate the following model with heterogeneous treatment effects:

$$\begin{aligned} Refi^{Obs}_{i,t} &= \alpha + X_{i,t}\delta + \beta_1 \mathbf{1}[t \geq T_{Imp}] + \beta_2 \mathbf{1}[Elig_{i,t}] + \beta_3 Unemp_{i,t} \\ &+ \beta_4 \mathbf{1}[t \geq T_{Imp}] \times \mathbf{1}[Elig_{i,t}] + \beta_5 \mathbf{1}[t \geq T_{Imp}] \times Unemp_{i,t} \\ &+ \beta_6 \mathbf{1}[Elig_{i,t}] Unemp_{i,t} + \beta_7 \mathbf{1}[t \geq T_{Imp}] \times \mathbf{1}[Elig_{i,t}] \times Unemp_{i,t} + e_{i,t} \end{aligned}$$

where all variables are defined as above. In addition to estimates of the effect of local unemployment rates on program takeup (β_7), which is the main coefficient of interest, I also allow the effect of unemployment rates to vary based on time (β_5) and by cohort (β_6) in order to control for differential trends in the dependence of credit supply on unemployment for various origination vintages and across time.

Table 3.5 shows the estimates of the difference-in-differences specification with heterogeneous treatment effects, as with the results from Table 3.4 estimated on samples originated within 3, 6, and 12 months of June 2009. As shown in the table, the key coefficient (β_7) on the interaction between eligibility, unemployment rate and the HARP 2.0 period is positive and highly statistically significant across all specifications. The baseline results indicate that a 1% increase in local unemployment rates[27] increases refinancing by .35% CPR. Normalizing by the baseline estimate for program takeup (β_4) suggests that for every 1% increase in local unemployment rates, HARP program takeup increases by roughly 7%. While the estimates using 6- and 12-month sample windows are broadly similar, they do suggest less heterogeneity in program takeup by unemployment rate. Nevertheless, interpreting the results using a longer sample window as a lower-bound, these results still indicate that program takeup varied by an economically meaningful 3% with each 1% change in local unemployment rates. To frame my estimate of the effect of unemployment on takeup quantitatively, the baseline results suggest that a 1% increase in unemployment rates for HARP-elgible borrowers

[27] In all results from this section, unemployment is measured on a 0-100 scale so that the regression coefficient can be interpreted directly as a the effect of a 1% change in the unemployment rate.

Table 3.5: Heterogeneity in Program Effect by Unemployment Rate

	3-Month Window	6-Month Window	1-Year Window
Eligible	0.0011***	-0.0001	-0.0005
	(3.49E-4)	(3.17E-4)	(2.81E-4)
Post-HARP	0.0051***	0.0092***	0.0110***
	(3.46E-5)	(3.75E-5)	(3.54E-5)
Post-HARP 2.0	-0.0025***	-0.0028***	-0.0011***
	(3.93E-4)	(4.63E-4)	(5.73E-4)
Post-HARP 2.0 x Eligible	0.0049***	0.0054***	0.0032***
	(2.73E-4)	(3.06E-4)	(3.51E-4)
Eligible x Unemp.	9.722E-6	1.708E-4***	2.588E-4***
	(3.47E-5)	(3.20E-5)	(2.98E-5)
Post-HARP 2.0 x Unemp.	2.745E-4***	3.317E-4***	3.074E-4***
	(3.80E-5)	(4.53E-5)	(5.54E-5)
Post-HARP 2.0 x Eligible x Unemp.	3.521E-4***	1.636E-4***	1.202E-4***
	(3.03E-5)	(3.53E-5)	(4.16E-5)
FICO	5.31E-5***	5.41E-5***	5.34E-5***
	(6.73E-7)	(6.05E-7)	(5.31E-7)
DTI	-1.66E-4***	-1.69E-4***	-1.64E-4***
	(4.55E-6)	(5.51E-6)	(6.89E-6)
LTV	1.16E-3***	1.25E-3***	1.31E-3***
	(6.67E-5)	(7.94E-5)	(9.88E-5)
Int. Rt. Spread	3.41E-3***	5.56E-3***	6.81E-3***
	(9.61E-5)	(1.06E-4)	(1.31E-4)
Unemp.	-2.24E-4***	-2.99E-4***	-3.42E-4***
	(1.54E-5)	(1.74E-5)	(1.83E-5)
WALA	3.33E-4***	3.63E-4***	2.85E-4***
	(5.92E-6)	(8.02E-6)	(8.29E-6)
Δ Mtg. Rt.	-1.22E-2***	-1.04E-2***	-1.06E-2***
	(3.38E-5)	(4.57E-5)	(5.98E-5)
Equity	1.69E-7***	1.88E-7***	2.08E-7***
	(1.00E-8)	(1.13E-8)	(1.36E-8)
Lender FE?	Yes	Yes	Yes
State FE?	No	No	No
N	83,824,948	121,448,185	177,172,126
R^2	0.0623	0.0590	0.0556

Notes: Dependent variable is indicator for prepayment, sample is all loans in HARP dataset. Standard errors clustered at individual-level in parentheses. Additional controls include loan interest rate and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. ***/**/* indicates significance at 99%/95%/90%-level respectively.

during the period following HARP 2.0 implementation is roughly equivalent to an additional $1,600 in home equity or a 2.3 bp decline in mortgage interest rates. More likely however these results under-estimate the actual effect of employment status on the marginal propensity to refinance under HARP. As discussed in Section 1.5.4, in this specification I treat local unemployment rates as a proxy for borrower-level unemployment, as this is the best available estimate, but the results from Gyourko and Tracy (2014) [107] suggest that using sub-aggregate rates as a proxy will lead to downward attenuation bias (towards zero) in the estimated effect of individual unemployment. As such, while I believe that these results are a strong indication that unemployed borrowers are more likely to takeup under HARP than similar employed borrowers, quantitatively these results may provide an under-estimate of the true effect. All estimates of the other control variables and interaction terms are consistent both across specifications and also with the results from Table 3.4. The estimates for what I refer to as “second-order" interaction terms (β_5 and β_6) suggest that the negative effect of unemployment was attenuated for borrowers after 2012 (the HARP 2.0 interaction) and for borrowers from older origination vintages (the eligibility interaction)[28]. These results, which may be indicative of generally loosening credit standards post-2012, highlight the importance of controlling for these second-order interactions in my main specifications. Nevertheless, the results from Table 3.5 provide reasonable evidence that the marginal propensity to refinance under HARP did vary positively with local unemployment rates.

The results from Table 3.5 are the primary contribution of this analysis, and so in the remainder of this section I focus on the robustness of these results to various alternative specifications. I focus on three different sets of robustness tests. First, I consider the same standard set of diagnostics employed for the difference-in-differences specification in Table 3.4 by using a one-month cutoff for eligibility or by instrumenting for house price appreciation and unemployment rates. Second, I test for non-linearity in the relationship

[28]Note also that in the baseline specification, the combined effect of all unemployment terms is still negative.

between unemployment rates and takeup under the HARP program by analyzing response to the program within quartiles of unemployment rates. Third, I add an additional eligibility criteria based on the same LTV-derived measure. Finally I test for parallel pre-trends in prepayment rates across eligible and ineligible cohorts by estimating a unique eligibility-unemployment interaction term for each month around the implementation of HARP 2.0. I plot these difference-in-difference fixed effects estimates in order to show graphically that differences in takeup by unemployment rate arose at around the time that HARP 2.0 was implemented and were not a continuation of long-run trends.

Tables A34 and A35 in Appendix G show the estimates from the first set of robustness tests. As with the difference-in-differences specification, I verify that the response-heterogeneity results from Table 3.5 are robust to using a one-month cutoff for eligibility or instrumenting for house prices and unemployment rates. These estimates largely confirm the baseline results, with parameter estimates for both the key variable of interest, the eligibility-post-implementation-unemployment interaction, and the other control variables that are quantitatively comparable. Indeed the estimates of heterogeneity in the marginal propensity to refinance under HARP are similar but somewhat stronger than the baseline estimates from Table 3.5 when I use a less strict measure of program eligibility in Table A34 or instruments in Table A35, with the results using a three-month sample window indicating that a 1% increase in unemployment rates leads to respectively a 13.6% or 7.3% increase in takeup. The second set of robustness tests checks for non-linearity in the relationship between unemployment rates and program takeup. I estimate a model similar to the baseline model from Table 3.5 but allowing flexibly for non-linearity by interacting the difference-in-differences term with indicators for each unemployment rate quartile rather than with a linear rate term. The results from three separate specifications, a baseline specification and two alternative specifications using a one-month HARP eligibility cutoff and instruments for unemployment rates and house prices, are shown in A36 in Appendix G[29]. Across all three specifications, the estimates using the primary three month window sample indicate that HARP program takeup increases essentially monotonically in local unemployment rates. In each case, the greatest difference is between the lowest unemployment-rate quartile[30] and the other moderate-to-high unemployment-rate quartiles. Nonetheless, with the exception of the specification using instruments shown in Panel C, where program takeup is highest among borrowers in the penultimate quartile of the unemployment rate distribution, I do find that program takeup increases with each quartile. The baseline estimates shown in Panel A indicate that takeup for borrowers in the highest quartile (Q4, above 10% unemployment) are 40-50% more likely to refinance under HARP than borrowers in the lowest quartile (omitted group, below 6.5% unemployment). Appliying the baseline estimates from Table 3.5 of a 7% increase in takeup for each 1% increase in unemployment rates to the difference in means for these two quartiles (6% for Q1 and 11% for Q4) gives a back-of-the-envelope estimate of a 35% difference in program takeup. This estimate is relatively similar to observed difference, with the remaining 5-15% discrepancy likely due to the concavity of the relationship between marginal propensities to refinance and unemployment rates. Intuitively, because in the quartile-based specification I estimate a non-linear, concave relationship between unemployment rates and program takeup, drawing a hypothetical best-fit line through the four quartile point estimates would tend to over-estimate takeup for the lowest quartile and under-estimate both takeup for the highest quartile and the effect of unemployment rates overall. Nevertheless I feel that the baseline estimates using quartiles as well as the two alternative specifications provide reasonable support for that conclusion that the marginal propensity to refinance under HARP is increasing in unemployment rates.

[29] For brevity I display only the estimated program effect; estimates for the control variables are very comparable to those from the baseline model in Table 3.5.

[30] Q1, the omitted group, has an implied estimate of zero

For the third set of robustness tests I re-estimate the baseline difference-in-differences specification with heterogeneous response rates using a definition of eligibility based on a combination of current LTV and origination date. Specifically, I re-estimate the results from Tables 3.5, A34, and A35 but interacting the origination-date-based eligibility indicator with an indicator for current LTV ratios above 75%. As mentioned previously, in addition to origination date, one of the criteria for HARP eligibility is that the borrower's current LTV ratio be above 80%, and I proxy for this criterion with a 75% LTV cutoff. Table 3.6 shows the results of the models with LTV-based eligibility controls corresponding to the estimates shown in Table 3.5. As can be seen in the table, I continue to estimate that HARP takeup rates increase with unemployment, and the baseline sample estimates suggests that a 1% increase in local unemployment rates increases refinancing by .4% CPR, nearly the same as the estimates from Table 3.5. However, with high-LTV eligibility controls I find that the effect of eligibility itself (β_2) is much stronger compared to comparable results with the baseline eligibility definition, while the effect of the difference-in-differences interaction term between HARP 2.0 implementation and the eligibility (β_4) is insignificant with the baseline sample window and slightly negative for longer sample windows. Taken literally, this result would imply that the HARP program had a negative or zero impact on refinancing activity. This counterintuitive result likely stems from the effect of the original HARP program; prior to the implementation of HARP 2.0, high current LTV ratios[31] could disqualify a borrower from the HARP program. When I control for a difference-in-differences interaction between implementation of the original HARP program and LTV-based eligibility in Table A37 in Appendix G, I estimate that the effect of the interaction is negative for the original HARP program and positive for HARP 2.0. This result suggests that high-LTV borrowers gained increased access to credit through the HARP 2.0 program but, on balance, not from the original HARP program, which tended to exclude those with the

[31] Above 105% from program inception, increased to above 125% post-December 2009.

worst equity positions. Absent this discrepancy regarding the effect of the LTV-based eligibility criterion, the other estimates accord well with my prior results from Table 3.5. I also estimate the model with LTV-based eligibility controls from Table 3.6 using either a one-month cutoff for program eligibility or instruments for house prices and unemployment rates. The results from these specifications, shown in Tables A38 and A39 of Appendix G, largely confirm the analogous specifications shown in A34 and A35 of Appendix G, although again with the same discrepancy regarding the estimated HARP 2.0 difference-in-differences term (β_4). Nevertheless I interpret these results as largely confirming the main result from Table 3.5 that takeup under HARP 2.0 increases in unemployment rates.

As a final robustness test, I verify that the effect of unemployment on refinancing activity for the HARP-eligible cohort increased only following the introduction of HARP 2.0. This test roughly corresponds with a test of the parallel trends assumption implicit in my difference-in-differences with heterogeneous treatment effects design. In order to identify the effect of unemployment on program eligibility, I rely on the assumption that the effect of unemployment on refinancing for the eligible cohort was constant prior to the implementation of HARP 2.0. To test this assumption, I estimate the following event-study variation on the main specification:

$$\begin{aligned} Refi^{Obs}_{i,t} &= \alpha + X_{i,t}\delta + \beta_1 \mathbf{1}[t \geq T_{Imp}] + \beta_2 \mathbf{1}[Elig_{i,t}] + \beta_3 Unemp_{i,t} \\ &+ \beta_4 \mathbf{1}[t \geq T_{Imp}] \times \mathbf{1}[Elig_{i,t}] + \beta_5 \mathbf{1}[t \geq T_{Imp}] \times Unemp_{i,t} \\ &+ \beta_6 \mathbf{1}[Elig_{i,t}] Unemp_{i,t} + \Sigma_{\tau=T_{Imp}-S}^{\tau=T_{Imp}+S} \beta_{7,\tau} \mathbf{1}[t=\tau] \times \mathbf{1}[Elig_{i,t}] \times Unemp_{i,t} + e_{i,t} \end{aligned}$$

where the coefficients $\beta_{7,-S}$ to $\beta_{7,S}$ are fixed-effects estimates of the effect of the unemployment-eligibility interaction in a window of length $2S$ around the implementation date T_{Imp}. Figure 3.3 plots these estimated coefficients for a model corresponding to the baseline specification shown in Table 3.5. My estimates suggest

Table 3.6: Heterogeneity in Program Effect by Unemployment with LTV Controls

	3-Month Window	6-Month Window	1-Year Window
Eligible x High LTV	0.0059***	0.0063***	0.0061***
	(7.88E-5)	(1.00E-4)	(1.32E-4)
Post-HARP	0.0059***	0.0100***	0.0117***
	(4.63E-5)	(4.92E-5)	(4.48E-5)
Post-HARP 2.0	-0.0006	-0.0003	0.0013***
	(2.44E-4)	(3.28E-4)	(4.45E-4)
Post-HARP 2.0 x Eligible x High LTV	0.0005	-0.0010*	-0.0046***
	(4.86E-4)	(5.13E-4)	(5.84E-4)
Eligible x High LTV x Unemp.	-7.216E-4***	-7.911E-4***	-7.404E-4***
	(1.79E-5)	(2.13E-5)	(2.45E-5)
Post-HARP 2.0 x Unemp.	3.538E-4***	2.667E-4***	1.397E-4***
	(2.34E-5)	(3.29E-5)	(4.44E-5)
Post-HARP 2.0 x Eligible x High LTV x Unemp.	3.391E-4***	4.349E-4***	6.503E-4***
	(5.05E-5)	(5.31E-5)	(6.05E-5)
FICO	5.34E-5***	5.46E-5***	5.44E-5***
	(6.51E-7)	(5.93E-7)	(5.28E-7)
DTI	-1.64E-4***	-1.67E-4***	-1.61E-4***
	(4.50E-6)	(5.44E-6)	(6.82E-6)
LTV	1.22E-3***	1.32E-3***	1.37E-3***
	(6.86E-5)	(8.20E-5)	(1.02E-4)
Int. Rt. Spread	3.36E-3***	5.18E-3***	6.45E-3***
	(8.92E-5)	(9.89E-5)	(1.32E-4)
Unemp.	1.01E-5	6.31E-5***	5.25E-5***
	(1.18E-5)	(1.16E-5)	(1.16E-5)
WALA	3.41E-4***	3.73E-4***	2.96E-4***
	(6.04E-6)	(8.11E-6)	(8.35E-6)
Δ Mtg. Rt.	-1.23E-2***	-1.06E-2***	-1.09E-2***
	(3.69E-5)	(4.32E-5)	(6.31E-5)
Equity	1.69E-7***	1.89E-7***	2.08E-7***
	(1.01E-8)	(1.14E-8)	(1.37E-8)
Lender FE?	Yes	Yes	Yes
State FE?	No	No	No
N	83,824,948	121,448,185	177,172,126
R^2	0.0625	0.0591	0.0556

Notes: Dependent variable is indicator for prepayment, sample is all loans in HARP dataset. Standard errors clustered at individual-level in parentheses. Additional controls include loan interest rate and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Figure 3.3: Estimated Eligible-Unemployment Interaction by Month

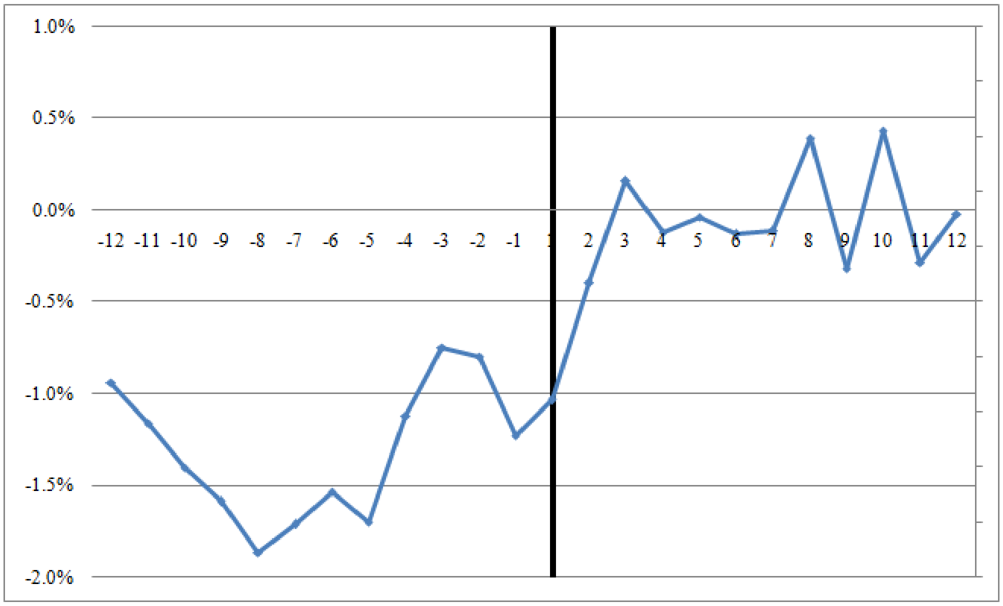

Notes: Estimated fixed-effects for eligibility-unemployment interaction. Dependent variable is indicator for prepayment, sample is all loans in HARP dataset. Additional controls include HARP eligibility and post-HARP 2.0 implementation indicators, interactions between eligibility and HARP 2.0 and post-HARP 2.0 and unemployment rates, FICO, DTI, LTV, unemployment rate, WALA, change in mortgage rates, equity, loan interest rate and interest rate spread, and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. Estimation sample is set of 83,824,948 monthly observations from loans originated within 3 months of April 2009.

that the effect of unemployment on refinancing activity is roughly constant prior to the introduction of HARP 2.0, with little evidence of pre-existing trends. Although there is a jump in prepayment between four and two months prior to the implementation of HARP 2.0, corresponding to September through November of 2011, this is likely the result of a sharp drop in mortgage interest rates by nearly 60 bp between July 2011 and December 2011, and in any event the refinancing wave that resulted appears to burn out prior to the introduction of HARP 2.0. On implementation, the estimated interaction effect rises by roughly 1% CPR, where it remains in subsequent months, leading me to estimate a positive overall impact for the eligibility-post-implementation-unemployment term. I interpret these patterns as suggestive evidence in favor of the parallel trends assumption, as they seem to confirm that differences in takeup by unemployment rates only emerged following the introduction of HARP 2.0.

The results from this section provide some evidence that the HARP program disproportionately benefited unemployed borrowers. In this section I first measured the effects of HARP 2.0 itself by estimating the differential increase in refinancing for eligible borrowers following program implementation. This estimate, which I treat as a measure of program takeup, suggested that HARP 2.0 increased prepayment rates by between 34 and 39% relative to baseline speeds. Next, by interacting this difference-in-differences term with local unemployment rates, I showed that takeup rates increased with area unemployment rates, with a 1% increase in unemployment rates translating to a 7% increase in takeup relative to the baseline. Through a variety of robustness tests, I showed that these results were monotonic across the unemployment rate distribution, although with the largest differences between areas with elevated unemployment rates and areas with full employment, were not sensitive to changes to the specification, and were driven entirely by the advent of HARP 2.0 and not be pre-existing trends. I interpret these results as evidence that the HARP program allowed unemployed borrowers who were previously unable to refinance due to credit constraints to obtain credit. The greater marginal propensity to refinance under HARP for borrowers who are likely to be unemployed suggests that the program disproportionately benefited such borrowers. Importantly however, I cannot say whether these differences in takeup are due to unemployed borrowers deriving greater value from refinancing, as suggested by the results from Sections 3.2 and 3.3, or due to the HARP program loosening credit constraints by more for unemployed borrowers. In Chapter 4, I address this question directly via a structural model. Nevertheless, the results from this section do complement those from Sections 3.2 and 3.3 by suggesting that policies aimed at easing credit supply provide disproportionate benefits to the unemployed.

3.5 Conclusion

This chapter provides new evidence on how the benefits of credit supply policies vary with employment status. The results from Section 3.2 suggest that the credit constraints facing borrowers with high levels of home equity wealth are relatively insensitive to unemployment rates, and that among such borrowers, the probability of refinancing is increasing in unemployment rates. These findings suggest that demand for refinancing credit is greater for unemployed borrowers, but that for all except the most creditworthy borrowers the supply effect of reduced access to credit overwhelms the demand effect. In this sense my findings are in line with prior researce [190] documenting that borrowing constraints interfere with consumption-smoothing behavior operating through the refinancing channel. [51] The results from Section 3.3 confirm this intuition, as borrowers who are eligible for HARP refinances and hence essentially unconstrained are more likely to refinance when unemployment rates are high. I estimate that a 1% increase in local unemployment rates leads to a 3.4% to 4.4% increase in refinance CPR and find that the effect of unemployment on the vaue of refinancing is monotonic across the unemployment rate distribution. Finally in Section 3.4 I show that takeup under the HARP program was greatest in high-unemployment areas, indicating that the benefits of policies that ease credit constraints accrue disproportionately to the unemployed. My estimate indicate that a 1% increase in local unemployment rates leads to a 3% to 7% increase in HARP program takeup, with takeup once again monotonically increasing in unemployment.

While these results are useful for policy analysis in particular, several key limitations inherent in my research design make it difficult to extrapolate these results to widely divergent settings. First, the empirical design employed in Section 3.2 can identify, at best, the sign of the relationship between unemployment and refinancing rates, and not the quantitative magnitude. While the results from Section 3.3 are somewhat more applicable externally, even under the most charitable assumptions that credit constraints are totally negligible for HARP-eligible borrowers the model measures only the sensitivity of refinancing to unemployment rates and cannot be used to analyze welfare. Moreover these results look only at a small set of eligible borrowers during a relatively unique period just after the implementation of HARP 2.0, and as such these results may not be generalizable. Second, the estimates from Section 3.4 provide no indication on why takeup is higher among the unemployed. While the results from Sections 3.2 and 3.3 suggest that part of the estimated effect is due to greater demand for refinancing credit among unemployed borrowers, as much if not more of the effect could just as well come from the fact that the HARP program reduced credit constraints more for unemployed borrower than it did for employed borrowers. In Chapter 4 I attempt to address these concerns by estimating the sensitivity of borrower valuations for refinancing credit directly via a structural model in order to make quantitative claims regarding borrower welfare. I then use the estimated model, which explicitly models credit constraints, to simulate the effects of a HARP-type policy by holding borrower valuations fixed and varying credit supply conditions.

The results in this chapter nonetheless contribute to the existing literature in two key ways. First, to my knowledge this chapter provides the first evidence on how takeup rates under the HARP program varied with borrower characteristics. My estimates can shed light on policy questions regarding ways to improve the efficicacy of HARP and other similar programs, as well as our understanding of the benefits and distributional effects of credit supply programs. Second, the estimates in this chapter inform the existing macroeconomic literature on models of idiosyncratic labor income risk by providing evidence that unemployment spells do in fact increase the value of refinancing credit. While several prior studies have shown similar suggestive evidence using other types of consumer credit, I employ two novel empirical techniques that allows me to analyze the effect of unemployment on demand for mortgage credit, the largest and perhaps most macroeconomically-important consumer credit market, absent the confounding effects of credit constraints. These estimates also inform the design of the model I present in Chapter 4, which both allows preferences for refinancing credit to vary with employment status in a microfounded way and allows observed refinances to be a product both of borrower demand and access to refinancing credit.

Chapter 4

Guarantee Fees as a Countercyclical Policy Tool

4.1 Introduction

In this chapter I analyze the effects of an alternative pricing policy for the GSEs guarantee fees. As discussed in Section 1.2.3 and shown in Figures 1.2 and 1.3, over the most recent business cycle the GSEs have pursued a highly procyclical policy with respect to their guarantee fees, with low fees charged during credit market booms and high fees charged during periods of tight credit. These periods also correspond, respectively, to periods of low and high unemployment, and to the extent that guarantee fees pass through to credit constraints, GSE policy resulted in looser credit constraints during periods of full employment and tighter constraints during periods of widespread unemployment. My analysis in Chapters 2 and 3 suggest that this is not an optimal policy, as I find that guarantee fees do pass through to credit constraints in a reasonable and predictable fashion and that the deleterious effects of these constraints fall disproportionately on the unemployed. Even setting aside considerations of fairness and borrower welfare, the substantial literature surveyed in Section 1.3.3 suggests that FRM refinancing frictions of just the sort that the GSEs have exacerbated are a meaningful contributor in prolonging recessions and inhibiting monetary transmission to consumption. As an alternative to the realized GSE guarantee fee policy, in this chapter I simulate a policy of high guarantee fees during periods of labor and credit market booms and low fees during recessions and nascent recoveries.

Chapters 2 and 3 of my dissertation provide some indication both of why such a pricing policy could be reasonably expected to be successful and what is required to analyze the effects of such a policy. The results from Chapter 2 indicate that originators pass a substantial fraction of changes in guarantee fees on to mortgage interest rates, with credit conditions tightening as a result. Moreover Chapter 3 suggests that credit supply policies disproportionately benefit the unemployed, in part due to the fact that demand for refinancing credit is strongest among the unemployed. In combination, these results suggest that high guarantee fees will be most costly and low guarantee fees most beneficial during periods of elevated unemployment. However, for two reasons one cannot simply apply the estimates from Chapters 2 and 3 in order to simulate the effects of a countercyclical guarantee fee policy. First, the results from Chapter 2 are based on observations drawn from a period in which the GSEs were the dominant funding channel for mortgage origination. We might expect that at times in which private-label securitization and other funding channels are more attractive credit conditions will be less sensitive to GSE policy, as intuitively any small change in guarantee fees will have no effect on credit conditions if originators never securitize via the GSEs. Second, the results from Chapter 3, even in the best-case scenario, provide evidence solely on the relationship between unemployment and demand for refinancing credit and do not allow for quantitative welfare calculations. While one might suspect that the results in Chapter 3 suggest a consumption-smoothing model of borrower behavior in which the marginal value of liquidity is higher during unemployment spells,[123] but borrowers are constrained in their consumption choices[13,190] and refinance to ease those constraints,[51] such a model is a prerequisite to any welfare calcuation and reduced-form evidence alone is not sufficient.

This chapter designs and estimates a model of the agency mortgage market capable of addressing these concerns and making quantitative statements about the effects of GSE policy on borrower welfare. Simulating such policies requires estimates of borrower valuations for relaxed credit constraints, and in particular how these valuations depend on employment status, as well as regime-invariant estimates of the effect of guarantee fees on credit constraints. The model I develop in this chapter includes both of these features. I estimate a structural model of dynamic borrower behavior in the face of endogenous credit constraints. Borrower preferences for cash-on-hand, and hence for refinancing, depend on individual-specific employment status, but borrowers are constrained in their ability to refinance by the willingness of lenders to extend credit.

Lenders can choose to fund loans via either GSE securitization or alternative channels, and as a result their decisions to approve applicants and the interest rates they offer in equilibrium will depend non-linearly on guarantee fees. I estimate the model's parameters using extensive data on borrower and lender decisions and use the estimated model to simulate the effects of a countercylical guarantee fee policy, which I then compare to the outcomes both from the observed policy and from other proposed GSE policies.

My results suggest that relative to the status quo, borrowers would benefit on average from a countercyclical guarantee fee policy. My structural estimates imply that borrower preferences for cash-on-hand are between 130% and 220% greater for unemployed borrowers, and that as a result an alternative countercyclical policy raising 6.3% greater revenue for the GSEs over the full business cycle with only a 3.9% increase in costs would nonetheless increase net borrower welfare by 2.3% over the same time period. The benefits to borrowers derive from two key sources: first, a 4.1% reduction in home-equity borrowing during the boom period prior to 2009, leading to a 1.1% decline in cumulative default post-2009, and second, a 5% increase in refinancing volume and 19% increase in mortgage equity withdrawal during and immediately after the start of the recession. I find that the effects of this simulated policy change on refinancing activity and borrower welfare following the start of the recession are comparable in magnitude to the effects of a simulated HARP program, while the reduction in home-equity withdrawal prior to the recession is nearly what could be achieved through stricter down-payment requirements and with significantly greater reductions in default. These results suggest that such a dynamic GSE pricing policy addresses multiple policy concerns by restraining borrowing during boom times [137] and providing credit access during recessions, [138] benefiting borrowers by both by reducing their tendency to over-leverage and allowing them to refinance as a consumption-smoothing tool.

This chapter contributes to the existing economic literature in three important ways.

First, I extend the literature on dynamic choice and household refinancing behavior[1] to incorporate endogenous credit constraints, discrete and continuous choices, and persistent time-variant liquidity preference. While I would describe neither my model nor my empirical strategy as a methodological contribution[2], the combination of techniques employed is novel and may prove useful in other settings. Second, this chapter adds to the policy literature on U.S. housing finance reform by proposing a novel GSE pricing policy. Despite the recognized importance of guarantee fees to the current policy debate[57] and the extensive range of housing finance policy proposals, some of which do recommend a countercylical policy stance,[110,180] this paper is the first to suggest a guarantee fee policy that varies with the business cycle. Finally, I build on the results from Chapter 3 by directly estimating how borrower preferences for cash-on-hand co-move with employment status. These estimates contribute directly to the literature on household liquidity preference underlying many computable general equilibrium models[3], as to my knowledge no prior paper has estimated the sensitivity of household liquidity preference to income structurally using micro-level data.

My analysis in this chapter proceeds as follows. In Section 4.2 I enrich the model from Section 1.4 in order to better capture key stylized facts regarding the agency mortgage refinancing market, some of which are motivated directly by either the results from Chapters 2 and 3 of from the previous research surveyed in Section 1.3.1. Section 4.3 outlines my empirical approach to estimating the parameters of the model described in Section 4.2 directly from the micro-data described in Sections 1.5.1 and 1.5.2. Crucially, my estimation strategy does not impose the assumptions from Section 1.4 that are necessary to ensure the validity of those analytical results, although I find that both of key results do in fact hold for the parameter values that I estimate. Finally, in Section 4.4

[1] For examples of papers in this tradition, see Bajari et al (2013),[22] Laufer[145] or Chen et al (2013).[53]

[2] Laufer[145] and Chen et al (2013)[53] incorporate continuous choice variables, while Connault (2014)[58] surveys the substantial existing literature on dynamic latent variables. My empirical methodology draws from a number of sources, which will be discussed in Section 4.3.

[3] Examples include Krusell and Smith (1998),[141] Gourinchas and Parker (2002),[99] and Lustig and Van Nieuwerburgh (2005).[152]

I present the results of several simulation exercises that project the effects of alternative guarantee fee policies out-of-sample. I compare the effects of these policies to several potential alternatives, including a streamlined refinance program similar to HARP and a restriction on maximum LTV ratios. Section 4.5 concludes with a discussion of the relevance of the results of these simulations for GSE policy.

4.2 A Structural Model of Agency Refinancing

Building on the insights developed in Chapter 2 and 3, in this section I describe the full model that I take to the data. In particular, I enrich the simple model presented in Section 1.4 with additional structure that will enable me to make quantitative comparisons of the welfare effects of alternative guarantee fee policies. Recall from Section 1.4 that the goal of my estimation procedure is to recover estimates of V^C, the value of continuation, V^R, the value of refinancing, and lender profits π. In order to do so, I make three additions to the model. First, on the borrower side I posit that V^C is actually the result of another decision of whether to continue and make a mortgage payment or default. Second, I also develop π in a similar fashion by allowing the lender to choose between securitizing and holding the mortgage. Third, I relate the borrower's value function to their flow payoffs using a Bellman equation.

In each period t, households i choose actions $(a_{i,t}, c_{i,t})$ to maximize their value $V_{i,t}$. Discrete action $a_{i,t} \in \{0,1,2\}$ corresponds to their choice of whether to default, continue, or refinance, respectively, and if the household chooses to refinance, they must also choose an LTV ratio $c_{i,t} \in [0, \bar{c}]$ for their application, where $\bar{c}$ is an exogenous upper bound that I assume is dictated by GSE policy[4]. The action-specific value of the borrower's realized action is given by:

$$V(a_{i,t}, c_{i,t}, X_{i,t}, Z_{i,t}) = u(a_{i,t}, c_{i,t}, X_{i,t}, Z_{i,t}) + \delta \mathbb{E}_{X_{i,t+1}, Z_{i,t+1}}[V(X_{i,t+1}, Z_{i,t+1}) | a_{i,t}, c_{i,t}, X_{i,t}, Z_{i,t}]$$

[4] I discuss how I chose a value for $\bar{c}$ based on published GSE documentation in Section 4.3.

where $Z_{i,t}$ is a binary variable denoting the borrower's employment status, $X_{i,t}$ is a collection of all other state variables, both aggregate (e.g. interest rates) and borrower-specific (e.g. home equity), u is a flow payoff which depends on the chosen action, and δ is the per-period discount factor. Building on the insights from Sections 3.2 and 3.3, I allow flow payoffs to vary systematically with employment status, and in particular for the marginal value of cash to differ by employment status. If the household chooses to continue or refinance, their state variables $(X_{i,t}, Z_{i,t})$ update accordingly and they continue into the next period. If instead the household defaults, I assume that the game ends, as the evidence from Hedberg and Krainer (2012) [112] suggests that following a default borrowers are not able to re-enter the market for a considerable length of time[5].

Unlike in standard dynamic optimization models, in this case the unconditional value function, $V(X_{i,t}, Z_{i,t})$, is not the maximum over action-specific values $V(a_{i,t}, c_{i,t}, X_{i,t}, Z_{i,t})$ due to the presence of credit constraints. In cases where the borrower is inhibited by credit constraints, the borrower's realized action will differ from her optimal choice. These credit constraints are endogenously determined via a stage game played between borrowers and lenders in each month t. This stage game takes the form of a static game of complete information. The extensive form of the stage game is shown in Figure 4.1, and as noted above it enriches the model from Section 1.4 with an additional choice on the part of lenders of whether to hold or securitize an originated loan and an additional choice on the part of borrowers of whether to continue or default. Note that I assume, as suggested by the evidence from Woodward and Hall (2012),[188] that borrowers can submit at most one refinancing application each month, so in each period this game is played exactly once per borrower.

The static game of complete information shown in Figure 4.1 is played as follows. At the start of the game, the borrower commits to a mixed-strategy offer rejection policy

[5]While they may be able to borrow again in the distant future, the expected discounted present value of this option can be reasonably approximated as zero.

Figure 4.1: Extensive Form of Structural Model

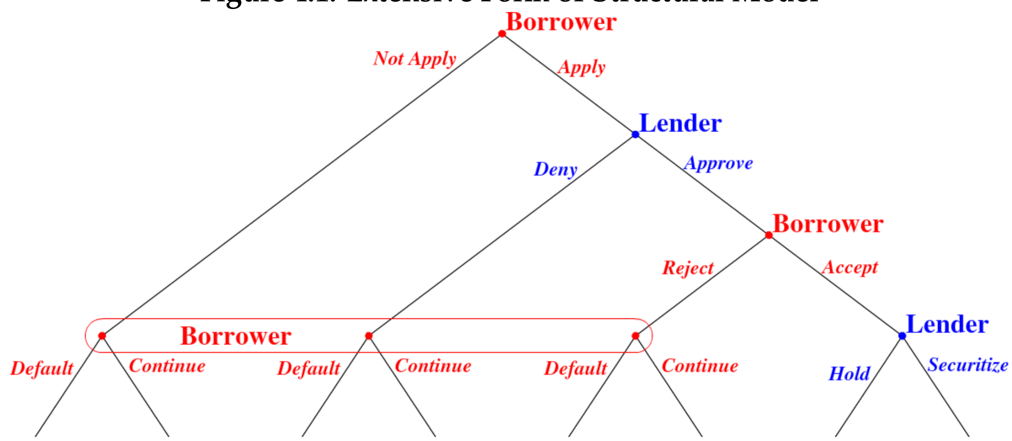

Notes: Timing of actions in stage game of structural model. Red nodes denote borrower decisions, blue nodes denote lender decisions. Information sets indicate nodes econometrician cannot distinguish between.

P_{REJ}, which is a function of the state variables, the offered interest rate $r_{i,j,t}$, and the LTV $c_{i,t}$. The lender plays a pure strategy response in both their decision to approve or deny the loan and their optimal interest rate choice. After the lender chooses an optimal interest rate offer, a cost shock realizes and the initial offer is revised[6]. Lenders face uncertainty generated by the borrower's probabilistic rejection policy[7], while borrowers face uncertainty regarding both the final (revised) interest rate offer and the lender's approval decision, which is based on the realization of a profitability shock that borrowers cannot observe ex-ante. Decisions are made sequentially in the order they appear in Figure 4.1, and different sets of preference shocks realize at each decision node.

I solve the game via backwards induction. Starting from the bottom-right node, lender j chooses the disposition of the loan ($d_{i,j,t}$) in order to maximize static profit $\pi^d(c_{i,t}, r_{i,j,t}, X_{i,t}, Z_{i,t}, S_{j,t}) + \mu^d_{i,j,t}$, where $S_{j,t}$ are credit-supply variables such as prevailing interest rates, guarantee fees, and market competitiveness and $\mu^d_{i,j,t}$ is a choice-specific profitability shock. Under the assumption that $\mu^d_{i,j,t}$ is Standard Gumbel distributed and

[6]Note that revisions to initial interest rate offers are common empirically.

[7]Offer rejection is quite common in the data, hence it is important to model this choice. However, without an ex-ante unobservable shock to interest rates, there would be no role for this choice, as borrowers would never reject offers in the subgame-perfect equilibrium of a game with no uncertainty.

IID across agents and time, the expected ex-ante profitability of an originated loan is:

$$EP_D(c_{i,t}, r_{i,j,t}, X_{i,t}, Z_{i,t}, S_{j,t}) = \log(e^{\pi^0(c_{i,t}, r_{i,j,t}, X_{i,t}, Z_{i,t}, S_{j,t})} + e^{\pi^1(c_{i,t}, r_{i,j,t}, X_{i,t}, Z_{i,t}, S_{j,t})}) + \gamma$$

where γ is Euler's constant and $S_{j,t}$ are credit-supply variables such as prevailing interest rates and guarantee fees. The probability that the loan is securitized is then:

$$P_{SEC}(c_{i,t}, r_{i,j,t}, X_{i,t}, Z_{i,t}, S_{j,t}) = P(d_{i,j,t} = 1) = \frac{e^{\pi^1(c_{i,t}, r_{i,j,t}, X_{i,t}, Z_{i,t}, S_{j,t})}}{e^{\pi^0(c_{i,t}, r_{i,j,t}, X_{i,t}, Z_{i,t}, S_{j,t})} + e^{\pi^1(c_{i,t}, r_{i,j,t}, X_{i,t}, Z_{i,t}, S_{j,t})}}$$

This final choice on the part of lenders allows the effect of guarantee fees to vary systematically with market conditions in the model. At times when securitization is relatively more attractive the effect of changes in guarantee fees on expected profit will be larger, and vice-versa.

I assume that lenders receive refinance applications passively[8], face fixed costs τ_L for making funding offers, and anticipate correctly that the borrower will accept the offer with probability P_{REJ}, which depends on borrower characteristics, the terms of the loan (including the interest rate and LTV), and credit market conditions[9]. At the time they make their approval decision, lenders receive a separate choice-specific profitability shock $\nu_{i,j,t}$, assumed IID Standard Logistic, and choose both whether to fund the loan ($f_{i,j,t}$) and the interest rate to offer ($r_{i,j,t}$). After making their approval decision, they then receive a separate shock $e_{i,j,t}$ to their offered interest rate, which can be thought of as a revision to their initial "good-faith estimate". Under the assumption that $e_{i,j,t}$ is distributed IID Standard Normal, lenders solve:

$$\max_{f \in \{0,1\}, r} f \times [\int_{e=-\infty}^{\infty} EP_D(c_{i,t}, r + e, X_{i,t}, Z_{i,t}, S_{j,t}) P_{REJ}(c_{i,t}, r + e, X_{i,t}, Z_{i,t}, S_{j,t}) d\Phi(e) - \tau_L + \nu_{i,j,t}]$$

Denote the integral term by $EP_F(c_{i,t}, r_{i,j,t}, X_{i,t}, Z_{i,t}, S_{j,t})$ and the optimal interest-rate offer

[8]In other words, that they do not advertise, either directly or via their prices.

[9]Note that this equilibrium borrower strategy is common knowledge.

by $r^*_{i,j,t}$. Borrowers cannot observe $\nu_{i,j,t}$ and thus expect ex-ante that the probability the loan is funded is:

$$P_{FUND}(c_{i,t}, X_{i,t}, Z_{i,t}, S_{i,j,t}) = \frac{e^{EP_F(c_{i,t}, r^*_{i,j,t}, X_{i,t}, Z_{i,t}, S_{j,t}) - \tau_L}}{1 + e^{EP_F(c_{i,t}, r^*_{i,j,t}, X_{i,t}, Z_{i,t}, S_{j,t}) - \tau_L}}$$

Borrowers receive a choice-specific utility shock $\epsilon^a_{i,t}$ based on their discrete actions. This shock is assumed to be IID Standard Gumbel as well, and I further assume that two separate sets of shocks realize at two different times. First, an initial set of shocks $(\epsilon^0_{i,t}, \epsilon^1_{i,t}, \epsilon^2_{i,t})$ realizes in the head node, and borrowers decide simultaneously whether or not to submit an application, continue, or default. Later, if the borrower's application is denied or she rejects a credit offer, a new set of IID shocks $(\epsilon^0_{i,t,\prime}, \epsilon^1_{i,t,\prime})$ realizes for the subsequent default-continue choice. Under this assumption, the ex-ante expected value of the default-continue choice is then:

$$EV_{NR}(X_{i,t}, Z_{i,t}) = \log(e^{V(0,0,X_{i,t},Z_{i,t})} + e^{V(1,0,X_{i,t},Z_{i,t})}) + \gamma$$

After choosing to submit a refinancing application, the borrower receives a multiplicative IID shock $\eta_{i,t}$ to the value of refinancing. Hence the borrower chooses LTV ratio $c_{i,t}$ to solve:

$$\begin{aligned}
\max_{c \in [0,\bar{c}]} \quad & P_{FUND}(c, X_{i,t}, Z_{i,t}, S_{i,j,t}) \int_{e=-\infty}^{\infty} [V(2, c, X_{i,t}, Z_{i,t})(1 - P_{REJ}(c, r^*_{i,j,t} + e, X_{i,t}, Z_{i,t}, S_{i,j,t}))\eta_{i,t} \\
+ \quad & P_{REJ}(c, r^*_{i,j,t} + e, X_{i,t}, Z_{i,t}, S_{i,j,t}) EV_{NR}(X_{i,t}, Z_{i,t})] d\Phi(e) \\
+ \quad & (1 - P_{FUND}(c, X_{i,t}, Z_{i,t}, S_{i,j,t})) EV_{NR}(X_{i,t}, Z_{i,t})
\end{aligned}$$

Denote the optimal LTV choice $c^*(\eta_{i,t})$ and the function to be maximized $EV_R(c_{i,t}, X_{i,t}, Z_{i,t}, S_{i,j,t})$. Then under the assumption that $\eta_{i,t}$ is distributed Standard Normal, and assuming that borrowers must pay fixed costs τ_B to submit a refinancing application, the probability that the borrower submits a refinance application is:

$$P_{APP}(X_{i,t}, Z_{i,t}, S_{i,j,t}) = \frac{e^{\int_{h=-\infty}^{\infty} EV_R(c^*(e^h), X_{i,t}, Z_{i,t}, S_{i,j,t}) d\Phi(h)} - \tau_B}{e^{\int_{h=-\infty}^{\infty} EV_R(c(e^h), X_{i,t}, Z_{i,t}, S_{i,j,t}) d\Phi(h) - \tau_B} + e^{V(0,0,X_{i,t},Z_{i,t})} + e^{V(1,0,X_{i,t},Z_{i,t})}}$$

We can then write the likelihood of any particular observation using these choice probabilities, as shown in Appendix H.

At this stage it may be useful for the purposes of expositing the model to describe how certain model elements are required in order to explain the variation obsered in the data. First, borrowers are assumed to reject high-interest-rate offers, and lenders are assumed to face fixed costs in making loan offers, in order to generate meaningful constraints on lender pricing. Without either the risk that a loan offer is rejected or fixed costs, lenders would always prefer to offer some credit contract at very high prices. Hence these two elements help explain the empirical fact that loan applications are frequently denied and funding offers are frequently rejected. Second, the timing of the shock to the borrower's interest rate e and shock to liquidity preference η are set both so that borrowers face uncertainty in their decision to apply for credit and so that the model is able to fit the observed choices of r and c using these shocks. Without these assumptions regarding timing, borrowers would never choose to apply for loan offers and either be denied or reject the offer as they would be able to predict ex ante the denial or rejection and choose not to apply instead. Endowing the model with this particular structure enables me to generate model-implied probabilities for each action and likelihoods for each continuous choice variable that I can match to the data. In the next section, I describe both how to do so and how to estimate the model directly using observed discrete and continuous choices.

4.3 Empirical Procedure

I estimate the model sequentially in three distinct stages corresponding to, in order, credit supply, borrower value functions, and borrower utility. This estimation procedure is designed to address four key econometric problems that complicate estimation. First, as shown in Figure 4.1, I do not observe refinancing applications, only completed refinances, and thus cannot distinguish between the three potential paths taken to reach the default-continue node. As a result, in many cases I must integrate over the distribution of potential refinancing applications the borrower may have submitted. Second, I observe interest rates only for completed refinancing transactions and not for offers, and hence must treat this as a latent variable. Third, I also do not observe individual employment status, only local-level signals such as the unemployment rate and average length of unemployment, and thus must treat this variable as a latent as well. Fourth and most importantly, as described in Section 1.5 I observe cross-sectional data on refinancing applications ("supply") and panel data on borrower choice ("demand") in two separate datasets, and hence as in Ho (2009) [115] must estimate the model in stages[10]. In what follows, I describe the three stages of estimation.

4.3.1 Stage 1 Estimation

In order to estimate the first-stage of the model I begin by relating credit supply parametrically to the data. I parametrize π^1, the profit from securitizing a loan, as a function of the price of the mortgage and parameters ϕ_1, and π^0, the profit from holding a loan, as a function of applicant credit risk characteristics, interest rate spreads, and parameters ϕ_0. The price of an originated mortgage is itself estimated via a regression of MBS prices on pool characteristics and secondary market conditions, the parameters of

[10] While it is theoretically possible to estimate the model jointly, the computational demands of the model make this approach technically infeasible.

which are shown in Table A40 of Appendix I, and this provides one mechanism for how guarantee fees affect lender profits and credit supply. This model is estimated on a sample consisting of all FNMA and FHLMC 30-year fixed-rate mortgage MBS with a pool size of at least $1 million issued between February 2004 and March 2013, as described in Section 1.5.3. I also parametrize the borrower's offer rejection function P_{REJ} as a logistic function of the loan characteristics, interest rate offer, market-level fixed effects and parameters β. The market-level variables are intended to measure local market competitiveness; if borrowers are less willing to reject high interest rate offers in more concentrated markets, the model will capture this tendency, potentially providing a mechanism for the monopolistically-competitive lender behavior discussed in as in Scharfstein and Sunderam (2013) [178] and estimated in Section 2.3. Finally, I parametrize the lender's interest-rate offer r^* as a function of borrower credit risk characteristics, market interest rates, and parameters α. The set of first-stage parameters to be estimated is thus $(\phi_0, \phi_1, \tau_L, \beta, \alpha)$.

The setup for the first-stage credit supply function is determined in part by the econometric issues described above. First, due to the separation between "demand" and "supply" data and sequential nature of the estimation procedure, I cannot tie the probability of offer rejection P_{REJ} to the value of refinancing. Instead, I must estimate P_{REJ} directly from the data and assume both that the observed offer rejection policy is an equilibrium response and that for the purposes of conducting counterfactual simulations the observed policy is, conditional on observables, invariant to guarantee fee regimes[11]. In this sense P_{REJ} is a lender belief that ends up being true in expectation, as borrowers do in fact reject offers frequently. While this is in some way a simplification of the model from Section 1.4, in that P_{REJ} is a reduced-form proxy for borrower expectations about

[11] While this may seem like a strong assumption, the probability of rejection is related only to credit supply and not to fees themselves, fees are allowed to affect credit supply in a nonlinear fashion that varies across fee regimes, and I observe intertemporal variation in guarantee fees. Hence it is not unreasonable to suppose that my estimates capture borrower offer rejection behavior in a policy-invariant manner.

the future and for local-market competitiveness[12], it would be computationally infeasible to estimate this policy any other way. Second, I estimate a linear model of interest rates as a function of credit risk characteristics so that estimates of the elasticity of interest rates with respect to LTV can be easily imported to Stage 2 for the purposes of solving the borrower's LTV choice problem[13].

Likewise the first-stage estimation procedure for credit supply is also designed to overcome the challenge of missing data on interest rate offers. Without such data, I must construct interest rate offer r^* as the solution to the lender's profit maximization problem, and in estimating α on a set of originated mortgages must correct for sample selection due to the exclusion of denied applicants and rejected offers. To do so, I first estimate $(\phi_0, \phi_1, \tau_L, \beta)$ via maximum likelihood. The maximum-likelihood routine fits the likelihood of four outcomes (application denied, offer rejected, loan originated and securitized, or loan originated and held) as a function of the probabilities P_{SEC}, P_{REJ} and P_{FUND} written above. I estimate this discrete-outcome model using the full HMDA dataset described in 1.5.2, which consists of a sample of 17,109,796 applications for conforming refinance mortgages submitted by borrowers from California between 2000 and 2012. Within the likelihood maximization routine, at each guess of the structural parameters I iterate on the interest-rate first order condition to recover r^*, and I approximate the integral over structural error e using Gauss-Hermite quadrature. Following Jiménez et al (2012, 2014),[125,126] I then use the estimated parameters to construct a selection-correction term and regress the interest rates on originated mortgages on borrower credit characteristics and funding rates to recover α. The dataset used to estimate this Heckman[111]-type model is a version of the full GSE header file described in Section 1.5.1, but restricted to the 3,611,896 refinance originations from California between 2000 and 2012. The full details of this procedure are contained in

[12] Both of which are proxied via covariates included in P_{REJ}.

[13] In order to conduct policy-invariant counterfactuals, this linear interest rate model is used solely for estimation and not for simulations.

Appendix H.

4.3.2 Stage 2 Estimation

In the second stage, I jointly estimate borrower's action-specific value functions and the dynamics of latent state variable $Z_{i,t}$, which captures employment status, using a modified form of the filtered maximum likelihood expectations-maximization (EM) routine discussed in Diebold et al (1999).[66] Following Barwick and Pathak (2011)[26] and Bazdresch et al (2014),[29] I posit that the value function $V(a_{i,t}, c_{i,t}, X_{i,t}, Z_{i,t})$ is a flexible polynomial function of the borrower's action and state. This approach is taken in order to avoid discretizing the continuous state space, and I estimate the value function separately from flow payoffs in order to avoid solving the Bellman equation in a high-dimensional state space. As in Diebold et al (1999)[66] I further posit that latent employment status $Z_{i,t}$ transitions according to a Markov process with time-variant transition densities that are themselves a function of local-area unemployment rates and mean unemployment duration. In what follows, I describe in order the method for estimating the value function and how this estimation procedure fits into the larger EM routine.

My approach to estimating the value function is a modification of the technique of mapping the value function to the observed choice probabilities from Hotz and Miller (1993),[117] adapted in my case to account for credit constraints and a continuous state space. Because the state space is continuous, I cannot estimate choice probabilities fully non-parametrically as in Hotz and Miller (1993),[117] and even if I was able to do so the correspondence they demonstrate between the value function and choice probabilities is invalid in cases, such as mine, where borrowers sometimes makes sub-optimal choices due to credit constraints. As an alternative I parametrize value V as a polynomial function[14] of the state with parameters $\theta_{a,z}$ that vary depending on the observed choice

[14]Specifically, a cubic spline with 10 gridpoints in each of house prices, unpaid balance, income, and

and latent state[15]. This specification allows the mapping from the state to the value function to vary flexibly depending on employment status, and importantly, I make no prior restrictions on how this mapping varies by employment status. As a consequence, how the marginal value of cash-on-hand varies with employment status, a crucial object of interest for the counterfactuals, is not imposed but rather inferred directly from the data. With this setup and given the first-stage estimates $(\hat{\phi}_0, \hat{\phi}_1, \hat{\tau}_L, \hat{\beta}, \hat{\alpha})$, I can construct the likelihood of each of the three observable outcomes (default, continue and refinance) as a function of the probabilities P_{APP}, P_{FUND}, P_{REJ}, and the value functions. As in the first stage, in cases where refinancing (and hence desired LTV) is not observed, I construct latent LTV preference by iterating on the borrower's first-order condition, with any integration over latent LTV choice or interest rate offers approximated via Gauss-Hermite quadrature. I assume a maximum LTV ratio $\bar{c}$ of 100% and constrain the borrower's latent LTV choice to lie below this value. GSE documentation from my sample period suggests that except in certain special circumstances the maximum allowable LTV ratio for conforming loans was 95%, although I observe loans with LTV ratios up to 105% in my sample and thus set the maximum limit slightly higher[16]. In cases where refinancing is observed, I fit the observed continuous LTV choice by adding an additional term based on the likelihood of the continuous preference shock $\eta_{i,t}$ required to rationalize the borrower's choice. The full likelihood function and details on its construction are discussed in Appendix H.

monthly payment, plus interactions between indicators for certain borrower types (e.g. income quartile) and macroeconomic variables (e.g. interest rates) and proxies for the consumption value of the home.

[15] As in Hotz and Miller (1993) [117] I normalize the value of one of the options, in this case default, to zero.

[16] Specifically, at certain times during my sample period GSE rules allowed first-time homebuyers to take out loans with LTV ratios above 95%, HARP borrowers could receive loans with potentially unlimited LTV ratios, and more recently the maximum limit has been extended to 97% for most mortgages. The maximum LTV ratio for first-time homebuyers is not relevant for my purposes, and my sample should not contain HARP refinances or recent originations, hence none of these situations explains the presence of 105% LTV mortgages in my sample and the applicable limit is likely 95% for first mortgages. However GSE standards also set slightly higher maximum LTV ratios for borrowers with secondary financing, enabling them to qualify even with combined LTV ratios as high as 100% during my sample period. Since my model does not feature secondary financing, it seems appropriate to model borrower choices as if they took out a single conforming loan with the limit set at the 100% limit for combined LTV. Please refer to Chapter 23.4 of Freddie Mac's Single-Family Seller/Servicer Guide for additional details.

The action- and latent-state-specific likelihood described above is used as an input to a larger estimation routine that allows me to predict the borrower's latent state and the transition dynamics of that state. As noted in Section 1.5.4, Gyourko and Tracy (2014)[107] note that estimates of the effect of unemployment on refinancing will be strongly biased towards zero if market-level average (unemployment rates) are used as a proxy. To correct for this bias, I posit that $Z_{i,t}$ transitions dynamically according to a function Π of observed unemployment rates, mean unemployment durations, and parameters ρ_z. Using these two variables allow me to distinguish between situations in which unemployment is elevated due to strong persistence of the unemployment state and situations in which unemployment is elevated due to high flows into unemployment. Under this assumption, it is possible to write the full likelihood as a function of the latent state probability $\hat{P}(Z_{i,t})$, the action- and latent-state-specific likelihood, the joint state probability $\hat{P}(Z_{i,t}, Z_{i,t-1})$, and the latent transition density $\Pi(Z_{i,t}|Z_{i,t-1}, X_{i,t}, \rho)$. Following Diebold et al (1999),[66] I estimate the full set of parameters, (θ, τ_B, ρ), via filtered MLE using an EM algorithm. In the E(xpectations) step, I filter out the latent state probability and joint state probability using a forward-backward recursion and the estimates from the prior M step. In the subsequent M(aximization) step, I compute the full likelihood and gradient and update the model parameters, and the EM procedure iterates to convergence. The model is fit to the matched dataset described in Section 1.5.4, which consists of 6,911,395 monthly loan-level panel records for 200,000 FRM borrowers from California from 2000 to 2012, and my estimation procedure relies heavily on two key advantages this dataset affords relative to the basic GSE dataset. First, that I can observe the borrower's choice of LTV when refinancing, and hence target the likelihood of this choice, and second, that I can observe unemployment rates at much finer levels of geographic detail. Additional details are provided in Appendix H.

At this point a discussion of how the second-stage model is identified is warranted. In a Hotz-Miller-type[117] setup, the value function is identified solely from empirical choice probabilities, and thus the parameters θ and τ_B are as well given the functional form restrictions I place on V as a function of parameters. One key difference relative to their model is that the choice to continue is a mixture over cases in which continuation is optimal and cases in which refinance optimal but the borrower is credit-constrained, and likewise for default. As such, the patterns in the data that allow me to identify in particular the covariance between employment status and the value of refinancing are those shown in Section 3.2: specifically, that borrowers in high-unemployment areas are no more likely to be credit constrained if they are good credit risks, and that for such borrowers refinancing is increasing in area unemployment rates. Two other key differences relative to the Hotz-Miller [117] framework are that my model features a continuous LTV choice and persistent unobservable employment status. While it is beyond the scope of this paper to offer a formal proof that the model is identified, I rely on the identification results demonstrated for single-agent discrete-continuous dynamic choice models in Blevins (2010) [40] and for dynamic choices with unobserved state variables in Connault (2014), [58] both of which should apply in my setting, as proof that neither of these two additions to the Hotz-Miller [117] setup, respectively, hinder identification. Note also that unlike as in Chapters 2 and 3 I do not instrument for either house price growth or unemployment rates, as the model allows directly for exactly the sort of latent heterogeneity in the value of the refinancing that these instruments are designed to address.

A final important identification concern in estimating models with latent state variables is the label-switching problem, or the invariance of the likelihood to a permutation of the labels. To give a concrete example, if refinancing is rarely observed relative to continuation then a priori it is equally likely that borrowers in state $Z_{i,t} = 0$ always refinance but the probability that $Z_{i,t} = 0$ is very small and that borrowers in state $Z_{i,t} = 0$ never refinance but the probability that $Z_{i,t} = 0$ is very large. As discussed in Stephens (2000), [182] addressing this concern requires some sort of prior information to break the tie vis a vis the likelihood between these two scenarios. Continuing the example above, I can either introduce a prior on whether $Z_{i,t} = 0$ corresponds to always or never refinancing or on whether $Z_{i,t} = 0$ is rare or common. Most models target the former with what are called artificial identifiability constraints, or inequality constraints on certain parameter values. I instead target the latter, the latent state probability, for three reasons. First, as noted in Stephens (2000),[182] artificial identifiability constraints don't always work. Second, since a key question is whether the marginal value of cash-on-hand is higher for unemployed borrowers, I would prefer not to impose this result via constraints on the parameters. Third, I have reliable data on market-level aggregate probabilities (unemployment rates) that provide a much more suitable prior. Thus, following Chung et al (2004),[55] I introduce prior information on the latent state probability by constraining the place-level unemployment rate implied by $\hat{P}(Z_{i,t} = 1)$ to be within a standard deviation of the observed place-level unemployment rate[17]. The full details of how this prior information is incorporated are discussed in Appendix H.

4.3.3 Stage 3 Estimation

In the third and final stage of estimation, I use the parameters estimated from the first two stages to recover non-parametrically the borrower's flow payoffs and use these payoffs to estimate a linear-in-parameters model of utility. Given the assumed error structure, the ex-ante expected value of being in state $(X_{i,t}, Z_{i,t})$ is given by:

$$\begin{aligned} \mathbb{E}_{\epsilon_{i,t}}[V(X_{i,t}, Z_{i,t}) &= \log(e^{V(0,0,X_{i,t},Z_{i,t})} + e^{V(1,0,X_{i,t},Z_{i,t})} \\ &+ e^{\int_{h=-\infty}^{\infty} EV_R(c^*(e^h),X_{i,t},Z_{i,t},S_{i,j,t})d\Phi(h) - \tau_B}) + \gamma \end{aligned}$$

[17]Chung et al (2004)[55] literally assign one of the observations to the zero state, which is sensible for inferring as they do from a univariate time series. My probabilistic procedure can be considered as a panel-data analog.

Hence I can re-write the Bellman equation as:

$$u(a_{i,t}, c_{i,t}, X_{i,t}, Z_{i,t}) = V(a_{i,t}, c_{i,t}, X_{i,t}, Z_{i,t}) - \delta \mathbb{E}_{X_{i,t+1}, Z_{i,t+1}}[\mathbb{E}_{\epsilon_{i,t+1}}[V(X_{i,t+1}, Z_{i,t+1})]]$$

Borrowers expect $Z_{i,t}$ to transition according to $\Pi(Z_{i,t}|Z_{i,t-1}, X_{i,t}, \hat{\rho})$, and I assume that borrowers have autoregressive expectations regarding the aggregate state. Specifically, I estimate place-level vector autoregression (VAR) models for ten aggregate state variables[18] to use as borrower expectations, where the state $X_{i,t}$ transitions according to a transformation of the past idiosyncratic state and the aggregate state. Using these assumed expectations, I construct the value function using the estimated parameters $(\hat{\theta}, \hat{\tau}_B)$, take simulated draws from the estimated VAR to construct a distribution of expected future states, and calculate utility $\hat{u}(a_{i,t}, c_{i,t}, X_{i,t}, Z_{i,t})$ for each choice combination[19]. In order to invert the Bellman equation I must assume a value for the discount factor β, as the results from Bajari et al (2013) [22] show that the discount factor is not identified unless one observes the full borrower history until the loan is paid off, and most loans in my sample are still active at the time the sample ends in 2013. To recover an estimate of borrower utility I assume a discount rate of .995, or just over 6% per year[20]. Finally I regress $\hat{u}(a_{i,t}, c_{i,t}, X_{i,t}, Z_{i,t})$ on a transformation of $c_{i,t}$ and $X_{i,t}$ in order to estimate marginal utilities and how they vary by employment status. These parametric utility models are also estimated on the matched dataset described in Section 1.5.4, and at this stage I leverage the third advantage this more detailed dataset confers relative to the unmatched GSE dataset by including characteristics of the home not contained in the GSE dataset, such as size and number of bedrooms, in the hedonic utility model. The full details of this estimation procedure are described in Appendix H.

[18] House prices, unemployment rates, mean unemployment duration, and seven yield-curve variables.

[19] I calculate the utility of the refinancing option at the LTV choice corresponding to the mean value of η for each observed borrower-month, even if the borrower did not refinance.

[20] In robustness tests my utility estimates were not significantly affected by choosing a discount rate of .99.

4.3.4 Estimated Parameters

This subsection overviews the parameters estimated according to the empirical procedure described in Section 4.3. The first-stage estimates of the parameters governing lender profits and the borrower's offer rejection probability are shown in Tables A40, A41, A42, and A43 of Appendix I. Table A40 displays the estimates of the sensitivity of MBS prices to certain MBS characteristics. These estimates, obtained via a hedonic regression on a large sample of MBS, are used as an input in estimating the other first-stage parameters. The estimated fit of the model is quite good, with an R^2 above 90%, and we can be reasonably confident that the model predicts MBS prices well-enough for our estimation purposes. Table A41 shows the estimated lender profit parameters. The first group of parameters, ϕ_0, measures the sensitivity of the profit from holding a loan to various loan characteristics, while ϕ_1 does the same for the profit from securitizing a loan, and τ_L is a fixed cost of underwriting. These estimated parameters largely accord with economic intuition. Holding mortgage loans is more profitable when the LTV ratio or unemployment rate is lower and the borrower is an owner-occupant with high income, presumably because the loan is at lower risk of default. Profits follow an inverted-U shape with respect to FICO; at very low credit scores, the high default probability reduces profitability, while at very high credit scores the increased prepayment risk also modestly reduces profitability[21]. Moreover, lenders prefer to hold mortgages with high spreads over average mortgage rates, and prefer to securitize when the price commanded by a mortgage is high[22]. Finally, lenders face substantial fixed costs of origination, including both monetary and other economic costs, equivalent to the value of holding a mortgage with a 2% higher interest rate.

Table A42 sets out the estimated parameters governing the borrower's probability of

[21]Note that the omitted category is the top category, loans with a FICO score between 740 and 850, so all parameter estimates should be compared to a value of zero for the top category.

[22]Note that because π_0 and π_1 measure economic profits and are expressed in terms of utils, the coefficient on price is not 1, as it would be if these profits were measured in dollar terms. It is thus difficult to compare the parameter estimates from ϕ_0 and ϕ_1.

rejecting an interest-rate offer. The estimates indicate that borrowers are more willing to accept an offer if they have high income, apply for a higher-LTV loan, have a low FICO score, are in a high-unemployment area, or receive a low-rate or low-spread offer. The coefficients on FICO score and unemployment suggest that less creditworthy borrowers are more desperate, or have lower bargaining power due to their inability to obtain alternative offers, while the coefficients on LTV and interest rates simply state that borrowers prefer offers for large loans at low prices. I interpret the coefficient on income as suggestive of reduced price-elasticity for higher-income borrowers. Table A43 shows the estimated interest-rate offer model used in the second stage of estimation. The estimates indicate that lenders make higher interest-rate offers to borrowers applying for higher LTV loans, borrowers with lower incomes of lower FICO scores, and non-owner-occupants living in high-unemployment areas. All of these estimates are consistent with lenders charging higher interest rates to borrowers with greater risk of default, or alternatively with prepayment risk being only a second-order driver of pricing. In addition to the signs of the coefficient estimates conforming with intuition, the sign on the selection correction term[23] is negative, suggesting that the selected sample of borrowers who accepted interest-rate offers and were securitized received lower interest rates. This estimate makes sense from both the borrower side, in that Table A42 suggests borrowers tend to accept offers with lower interest rate spreads, and from the lender side, in that Table A41 suggests lenders prefer to hold loans with higher interest rate spreads.

The estimated fit of the first-stage model to the data, as shown in Tables A43 and A45 of Appendix I, is quite accurate. The R^2 in Table A43 suggests that the model explains roughly 84% of the variation in interest rate spreads. While this may seem low, variation in spreads over a benchmark rate[24] are much more difficult to predict than the interest rate itself, as one can achieve fairly high accuracy on the second measure by simply

[23]The coefficient on which is normalized to 1.

[24]In this case, the 10-year swap rate.

predicting that everyone receives the average mortgage rate. Table A45 shows that the model predicts shares for the four outcome probabilities corresponding very closely to the observed shares. Panel A of Table A45 indicates that the predicted and observed shares are effectively identical in the full sample, and within any borrower subtype only differ substantively for the share of securitized loans made to high-DTI-ratio borrowers. Similarly Panel B indicates that the predicted and observed shares are fairly close in most years, although the model over-predicts the likelihood of securitizing a loan relative to holding a loan in 2000 and 2004-2007 and under-predicts in 2012. The 2004 to 2007 period corresponds to a boom in private-label mortgage securitization, largely driven by subprime lending but also involving substantial numbers of conforming loans. Because the profit from holding a loan in portfolio is not modeled separately from the profit from private-label securitization, the credit supply model will tend to under-predict shares from the holding option whenever profits from private-label securitization are high. This likely explains why the model over-predicts the share of loans securitized by the GSEs in these years, although given that the private-label market effectively disintegrated in 2008, I do not believe that this unmodeled feature of the market compromises the simulations I will present in Section 4.4, in which much of the action takes place post-2008. Panel C of Table A45 shows similar metrics by FICO score, and these results indicate that the model predicts outcomes quite accurately for higher-FICO borrowers but not very accurately for borrowers with FICO scores below 700. This may be because my data contains relatively few low-FICO borrowers[25], hence the model is essentially calibrated to predict outcomes for higher-FICO borrowers. Overall, however, the first-stage estimates largely conform to expectations and explain the variation in the data quite well.

The second-stage estimates of the parameters governing borrower's value functions and the probability of transitioning between employment and unemployment are shown

[25]Table A4 in Appendix C shows that the average FICO score in the data is above 700 in every year, and is well-above 700 in the last half of the sample.

in Tables A44 and A46 of Appendix I. Table A44 shows the estimated employment-to-unemployment transition parameters, with the probability of transitioning from employed to unemployed increasing in ρ_0 and the probability of transitioning from unemployed to employed increasing in ρ_1. As we should expect, the estimates suggest that for each income quartile, the probability of remaining unemployed or transitioning to unemployment increases significantly in the local unemployment rate and mean unemployment duration, although this effect is attenuated when either the unemployment rate or mean unemployment duration is either very high or rising rapidly. In lieu of the actual parameters on the borrower's value function, of which there are 3,701 and which cannot be presented in any intelligible format, I present instead the fit of the model to the data in Table A46, as the goal of the second-stage estimates is largely to predict borrower actions accurately throughout the state space. In the sample overall the predicted and observed choice probabilities are very similar, although the model somewhat under-predicts the probability of default, likely because so few defaults are actually observed in the data. The model predicts the probability of refinance accurately within each FICO score category, although it tends to under-predict refinance probabilities in the earlier half of the sample and over-predict them in the later half of the sample. Combined with similar systematic differences in the likelihood of securitization across years shown in Panel B of Table A45, these findings suggest that the model's treatment of the lenders choice between securitization and portfolio retention and the effects on this choice on borrower credit constraints is missing a key factor that varies systematically between the earlier and latter half of the sample. An obvious canditate for such a factor would be the presence of private-label securitization, which is unmodeled[26] and which provided an important source of credit prior to 2008 but vanished almost entirely thereafter. Nevertheless, the in-sample fit of the model is accurate enough to use the model in order

[26]Or more accurately, it is lumped into the held-on-portfolio, along with several other sale-based options.

to derive third-stage estimates of borrower preferences.

The third-stage estimates of the parameters governing borrower per-period utility are shown in Table 4.1. The parameters correspond to the marginal action-specific utility with respect to the variable in question, while $Z = 0$ and $Z = 1$ correspond respectively to the employed and unemployed states. The signs of the estimates indicate that borrowers dislike making high mortgage payments but enjoy receiving cash via home-equity conversion, that borrowers enjoy living in larger houses, and that borrowers generally enjoy making their payments as opposed to defaulting[27] when they have middle-range FICO scores[28]. Moreover the magnitudes of the coefficients vary between the employed and unemployed state in a predictable fashion: for each action, the unemployed have a higher marginal value of cash but care less about the other benefits of not defaulting such as continuing to live in a large house or improving one's credit score. These estimates, consistent with the evidence from Sections 3.2 and 3.3, suggest that the unemployed derive greater benefit from refinancing for consumption-smoothing purposes, as the marginal value of cash is relatively greater for the unemployed, but are more single-minded in that their utility is less sensitive to non-pecuniary aspects of servicing their mortgage. In principle the difference in the marginal value of cash across employment statuses could be measured in three different ways, corresponding the first three parameter estimates, each of which is measured in dollars. These estimates do not always agree with one another; while the measured marginal value of cash is roughly equivalent across states for the value of mortgage payments, the estimate marginal utility of extracted cash is between 4 and 6 times greater than the marginal utility of mortgage payments. This finding in and of itself is not unusual, as many strucural demand models find that preferences for cash streams from different sources are different. Bajari et al (2013) [22] for example find that the marginal

[27]The utility of the default action is normalized to zero, so all marginal utilities should be judged relative to defaulting.

[28]Except for borrowers in the unemployment state who are refinancing.

utilities from income, monthly payments and prepayment penalties are quite different despite all being measured in dollars, and indeed that higher prepayment penalties increase the utility of prepaying. However, using the marginal utility of the mortgage payment as a baseline figure, I estimate that unemployed borrowers derive between 130% and 220% more utility from each dollar than do employed borrowers, which is quite a sensible range. While these figures may seem large, in the case where borrowers have log utility they imply that unemployment spells reduce household income by between 56% and 69%, which seems a reasonable range for middle-class homeowners. Note also that this finding implies that the assumption from Section 1.4 regarding the relationship between unemployment and liquidity preference is true at the estimated parameter values. Importantly, this result is not imposed on the data by assumption, but rather estimated using the label-switching constraints described in Section **??** based on the patterns in the data described in Section 3.2. Finally, the estimated fit of the models, especially for flow payoffs from refinancing, is quite good, and the models explain roughly 72% of the variation in flow payoffs from continuation and 98% of the variation in flow payoffs from refinancing. Due to the procedure used to generate the utilities included as the dependent variable in these models, the dependent variable includes both the mean utility value and the Standard Gumbel econometric error typically included in such models. As such, 1 minus the R^2 places an upper-bound on the fraction of the variation explained by these Standard Gumbel preference shocks, implying that most of the variation is explained by the observed data and structural parameters.

The evidence presented in this subsection shows that the estimated structural model is consistent with a world in which unemployed borrowers derive greater benefits from refinancing for consumption-smoothing purposes than do employed borrowers, but also face greater credit constraints. In this sense the model is also consistent with the evidence presented in Section 3.2, as we should expect given that both models are identified from the same variation in the data. In the next section, I use the model estimated in this

Table 4.1: Estimated Borrower Utility Model

	Continue (Z = 0)	Continue (Z = 1)	Refinance (Z = 0)	Refinance (Z = 1)
Current Payment ($K)	-4.3229***	-13.8238***	-	-
	(0.2347)	(0.6884)	-	-
Refinance Payment ($K)	-	-	-4.2127***	-9.6370***
	-	-	(0.0842)	(0.1302)
Cash Extracted ($ M)	-	-	0.0260***	0.0415***
	-	-	(0.0025)	(0.0042)
FICO 660-700	8.1005***	2.1246***	5.4469***	-0.4528***
	(0.1401)	(0.0354)	(0.1263)	(0.0109)
FICO 700-740	2.6735***	0.4032***	1.8826***	-0.2793***
	(0.1169)	(0.0303)	(0.1065)	(0.0077)
Sq. Feet	17.6413***	13.9842***	4.4750***	1.7728***
	(0.6053)	(0.2557)	(0.4834)	(0.0391)
County FE?	Yes	Yes	Yes	Yes
N	6,711,854	6,711,854	6,711,854	6,711,854
R^2	0.7282	0.7178	0.9822	0.9744

Notes: Dependent variable is estimated action-specific utility, sample is all loans in matched dataset. Standard errors in parentheses. Additional controls include interactions between indicators for origination year and county with purchase price, and indicators for whether the property is a condominium, has a pool, or has a carport. ***/**/* indicates significance at 99%/95%/90%-level respectively.

section to simulate several policies aimed at either increasing or decreasing the supply of credit and analyze how this model predicts such policies will affect borrowers.

4.4 Counterfactual Policy Simulations

In this section I describe how I use the estimated model from Section 4.3 to simulate four different scenarios. The first is a "baseline" simulation to which the other three are compared. In the second scenario, I set an alternative countercyclical guarantee fee policy, raising the average fee to 45 bp between October 2003 and December 2008 and lowering it to 5 bp thereafter.[29] A graph of the average charged guarantee fee in the baseline and alternative simulations is shown in Figure A8 in Appendix J. For the purposes of comparing the results of the alternative guarantee fee policy to other potential refinancing-specific policies, I simulate two other scenarios corresponding to

[29] These dates were chosen because October 2003 ending a long stretch of loosening credit conditions as reported by the Federal Reserve's Senior Loan Officer Survey, making it a reasonable candidate for a time to try to tighten credit; and January 2009 marked the nadir of the recession, as well as roughly matching the implementation of the first HARP program, which began in March.

policies that were either proposed or implemented. In one, I simulate a HARP-type policy by allowing borrowers to refinance without constraints after June 2009. In the second, I simulate the effects of a tighter loan-to-value ratio cap by requiring that borrowers take out loans worth no more than 80% of their home price.

4.4.1 Simulation Results

In order to simulate the model, I step through the observed months in which borrowers make choices, simulating their choices and tracking key variables. I update the borrower's state (e.g. their equity position) depending on their past simulated actions, so that past decisions affect future decisions and benefits. In this sense, the simulation procedure makes two key assumptions. First, as a partial equilibrium model, it assumes that the paths of the exogenous state variables, such as house prices and interest rates, do not change across the various scenarios. Moreover lender conduct is held fixed, so while interest rates and funding offers adjust to the change in guarantee fees, the underlying lender policy functions are unchanged. Second, it assumes that borrowers do not anticipate the policy change and alter their behavior ex-ante. This assumption precludes, for example, a borrower levering up in anticipation of easier credit conditions starting in January 2009. I track a series of key variables used to judge the efficacy of the three policy interventions, including the rates of refinancing and default, the volume of refinancing and home-equity withdrawal, the expected utility and value of borrower's choices, and the GSE's market share, expected outlays, and revenues. The full details of the simulation procedure are explained in Appendix J.

The results of the simulated policies are shown in Table 4.2. Comparing the first and last columns, we see that the alternative guarantee fee policy had its intended effect. During the period of elevated fees, from October 2003 to December 2008, both refinancing volume and cash-out were lower than in the baseline, by 3% for the former and 4.1% for the latter. During the period of low fees, from January 2009 to December

Table 4.2: Simulation Results

Refi. Vol. ($B)	Baseline	HARP	80% Cap	Alternative G-Fee
10/03-12/08	2,923.73	2,923.73	2,902.49	2,835.30
1/09-12/12	4,092.86	4,349.02	3,996.07	4,298.24
Total	7,016.60	7,272.75	6,898.56	7,133.54
Cash Extracted ($B)	**Baseline**	**HARP**	**80% Cap**	**Alternative G-Fee**
10/03-12/08	1,289.92	1,289.92	1,230.91	1,236.72
1/09-12/12	188.16	206.36	165.96	223.84
Total	1,478.08	1,496.28	1,396.87	1,460.56
Cum. Default Prob.(%)	**Baseline**	**HARP**	**80% Cap**	**Alternative G-Fee**
10/03-12/08	1.20%	1.20%	1.21%	1.23%
1/09-12/12	8.70%	8.48%	8.85%	8.61%
Total	9.90%	9.68%	10.06%	9.84%
GSE Rev. ($B)	**Baseline**	**HARP**	**80% Cap**	**Alternative G-Fee**
10/03-12/08	39,755.86	39,755.86	38,631.45	81,774.20
1/09-12/12	46,300.61	78,518.64	43,229.28	9,704.14
Total	86,056.47	118,274.50	81,860.73	91,478.34
GSE Cost ($B)	**Baseline**	**HARP**	**80% Cap**	**Alternative G-Fee**
10/03-12/08	58,454.90	58,454.90	54,966.93	57,235.84
1/09-12/12	164,819.34	284,902.81	152,624.59	174,653.86
Total	223,274.24	343,357.71	207,591.52	231,889.69
GSE Shr. (%)	**Baseline**	**HARP**	**80% Cap**	**Alternative G-Fee**
10/03-12/08	34.05%	34.05%	33.85%	33.18%
1/09-12/12	57.70%	83.51%	56.66%	59.19%
Average	45.87%	58.78%	45.25%	46.18%
Cum. Default Balance ($B)	**Baseline**	**HARP**	**80% Cap**	**Alternative G-Fee**
10/03-12/08	53.40	53.40	52.92	55.63
1/09-12/12	412.43	396.61	411.08	405.85
Total	465.83	450.00	464.00	461.49
Expected Value	**Baseline**	**HARP**	**80% Cap**	**Alternative G-Fee**
10/03-12/08	88.17	88.17	87.89	86.33
1/09-12/12	78.26	82.99	78.07	83.86
Total	166.42	171.15	165.97	170.19

Notes: Results of simulation exercises. "Baseline" denotes simulations with observed guarantee fee path, "HARP" with credit constraints removed after 2009, "80% Cap" with LTV limited to 80% throughout, and "Alternative G-Fee" with the guarantee fee set to 45 from 10/03-12/08 and 5 from 1/09-12/12. Dollar figures denote cumulative values over the period. Value is an average per-person cumulative figure.

2012, the opposite pattern holds. Refinancing volume is 5% greater with the alternative guarantee fee policy relative to the baseline simulation, and net equity withdrawal is 19% greater. The net result of the policy change is an overall reduction home-equity withdrawal over the full 2003 to 2012 sample period of 1.2%, and refinancing volume over the full simulation period is only 1.7% greater than in the baseline scenario. The alternative policy is also revenue-improving; comparing the same columns, we see that in the alternative scenario total GSE fee revenues over the full sample period were 6.3% greater than in the baseline scenario, while outlays are essentially unchanged between the two scenarios. The timing of these revenues is also shifted, with considerably lower revenue during the low-fee period and higher revenue during the high-fee period. The impact of higher or lower guarantee fees on GSE revenue is partially offset in both the high- and low-fee periods by changes in the GSE's market shares. During the high-fee period, the GSE's share of new refinances is 2.5% lower than in the baseline, while in the low-fee period it is 2.6% greater. Cumulative default rates and defaulted volumes are modestly lower in the alternative scenario relative to the baseline over the full sample period, with 3% higher default rates prior to 2009 and 1.1% lower default rates thereafter, while the change relative to the baseline in defaulted volumes are slightly larger in magnitude, indicating that under the alternative guarantee fee policy relatively more high-balance loans default. While the alternative policy reduces default rates, GSE costs actually increase over the full sample period, largely because the GSEs market share is higher during the post-2009 period in which default costs were large. As such, although the alternative policy is revenue-improving in that it results in greater overall fee income over the business cycle, and the 6.3% increase in revenue is greater than the 3.9% increase in costs, the effect on GSE profit is actually negative, as the GSEs are observed in the baseline to suffer significant losses over the period in which GSE market share expands under the alternative simulation. While this result may indicate that the method I use to evaluate the effects of the alternative policy on the GSE's profit and loss is incomplete, I interpret my results as consistent with the alternative policy not amounting to a significant subsidy to borrowers.

Moreover, due to the change in the timing of refinancing and mortgage equity withdrawal, borrowers actually achieve about 2.3% greater expected value under the alternative scenario, with roughly 2.1% lower expected value prior to the recession and 7.1% higher expected value after the recession. This increase in value results both from reducing ex-ante overborrowing, and hence, as the results from Mian and Sufi (2009) [158] suggest, ex-post default, and from increasing the availability of refinancing credit ex-post. Moreover by reducing guarantee fees starting in 2009, the alternative policy not only significantly stimulates cash-out refinancing, which the results from Greenspan and Kennedy (2008) [101] and Zhou and Carroll (2012) [191] indicate is often used to support consumption spending[30], but allows borrowers to obtain credit at lower overall interest rates, further raising borrower welfare. The estimates from Table 4.1 also indicate that during the low-fee period post-2009 borrowers on average derived greater benefit from both extracted cash and the savings on lower interest payments, as more borrowers are estimated to be in the unemployed state where the marginal value of cash-on-hand is greater. Hence by shifting credit supply from a period in which unemployment rates and thus the value of refinancing credit is low to periods in which unemployment rates and the value of credit is high, the alternative policy improves welfare while having only modest effects on total refinancing volumes.

In order to form a basis of comparison, I simulate the effects of two other alternative policies and compare the results. The first scenario, intended to replicate the effects of the HARP program, sets the probability that a loan application is denied to zero starting in January 2009. This essentially allows borrowers to refinance via the GSEs at any time of their choosing[31]. The results of this simulation are shown in the second column of

[30] While my model of borrower utility is agnostic regarding the disposal of extracted cash, implicitly I assume that it is consumed.

[31] Specifically, I set the probability P_{FUND} that the lender grants them credit to 1, although as with the actual HARP program they may still reject high-interest rate offers or choose not to participate. Moreover,

Table 4.2. The simulated HARP policy results in a 6.3% increase in the volume of refinancing and a 9.7% increase in extracted home equity after implementation. I also find that the HARP program substantially increases expected value, with a net increase of 2.8%, or slightly larger than the effects of the alternative fee-pricing policy, over the full sample period, and a slightly smaller 6% increase in welfare during the post-2009 period. The simulated HARP policy also significantly reduces default rates by 2.6% during the implementation period by providing high-LTV borrowers an alternative default. The slightly larger 3.8% reduction in defaulted volumes and less-than-proportional increase in mortgage equity withdrawal relative to the increase in refinancing volume likely reflects the fact that only borrowers with high-LTV ratios, who also tend to have large balances, are granted eligibility under HARP, hence the program stimulates refinancing only for borrowers with little scope to lever up and deters default disproportionately for high-balance loans. As such the difference in the effect on borrower welfare between the simulated HARP policy and alternative fee policy stems from three sources. First, the HARP policy allows certain constrained borrowers to refinance but has no effect on interest rates, leading to lower interest payment savings and relatively more rejected funding offers than would result if guarantee fees were reduced. Second, the HARP policy increases the availability of credit disproportionately for borrowers close to the point at which default becomes optimal. As a consequence, while in the HARP scenario default rates are reduced relatively more than under the alternative fee policy, HARP disproportionately reduces default among borrowers for whom the welfare cost of defaulting is low. Third, as mentioned above the alternative guarantee fee policy results in a significantly greater increase in cash-out refinancing during the post-2009 period, providing greater overall stimulus. As a result, while refinancing volume expands by more following the recession under the HARP policy

interest rate offers are held fixed throughout. I use HARP's true eligibility standards by restricting eligibility to only those whose loans were GSE-securitized and originated prior to June 2009 with current LTV ratios above 75%. Additional details of the implementation of this simulation can be found in Appendix J.

than under the low-fee policy, the alternative fee policy has a greater positive effect on borrower welfare. Moreover my results indicate that the HARP policy also has significant effects on GSE revenues, costs, and market shares. Following implementation, GSE market shares, revenues and costs expand dramatically under the HARP scenario as the GSEs effectively become the refinace lender of last resort. However, unlike under the alternative fee policy costs actually expand by more (72.9%) under the HARP scenario than do revenues (69.6%), leading to significantly greater net GSE losses. As result, my simulations indicate that a targeted refinance program such as HARP would provide only modestly higher benefits to households over the full business cycle relative to a countercyclical guarantee fee policy while entailing much greater GSE outlays.

The second alternative scenario I simulate lowers the exogenous maximum LTV ratio from 100% to 80% for the entire sample period. As shown in the third column of Table 4.2, restraining the maximum LTV ratio in this fashion results in a reduction in both refinancing volume and equity withdrawal, and unlike in Bajari et al (2013),[22] Laufer[145] and Chen et al (2013),[53] I find that borrower welfare is .3% lower with the 80% LTV cap than under the baseline. My simulation results indicate that the alternative guarantee fee policy is nearly as effective as the LTV cap at reducing home-equity withdrawal prior to the recession, with the former generating a 4.1% reduction relative to the baseline and the latter a 4.6% reduction, and that the tighter LTV cap actually leads to an overall increase in the probability of default by 1.6% over the full sample period as borrowers with low home equity find themselves unable to refinance. While default rates are higher with the tighter LTV cap than under the baseline, total defaulted volumes and hence GSE credit losses are actually lower over both sub-periods, as borrowers in this scenario tend to have lower loan balances, so each default results in lower credit losses, and the GSE's market share is lower than in the baseline. As a result GSE revenues also decline relative to the baseline with an 80% LTV cap, although the through-the-cycle decline in GSE revenues of 4.9% is more than made up by reductions in losses of 7%, resulting in lower GSE losses overall. As such, although I find that an 80% LTV limit reduces GSE exposure to losses, by degrading the value of the refinancing option it results in a 1.7% decline in refinancing volume and 5.5% decline in cash-out-refinancing that is not fully offset by a reduction in default rates, leading to an overall decline in borrower welfare.

My estimates from Table 4.2 of the effects of an alternative guarantee fee policy or a HARP-type refinancing stimulus accord reasonably well with the estimated effects of each from Chapters 2 and 3 respectively. My estimates from Chapter 2 suggest that a 10 bp increase in guarantee fees reduces the probability of refinancing by 5.7% to 6% relative to the baseline[32]. Applying this estimated elasticity to the product of the change in guarantee fees relative to the baseline under my simulated counterfactual and the estimated GSE market share[33], which I treat as a proxy for the degree of pass-through, yields an estimated change in refinancing volume of 4%-4.3% for the high-fee period and 8.2% to 8.8% for the low-fee period, roughly in line with my structural estimates of 3% and 5% respectively. Moreover applying my estimate from Chapter 2 of the fraction of deterred refinances that lead to default to the change in refinancing volume generated by the alternative guarantee fee policy implies an in default volumes of 3.1% to 3.8% during the high-fee period and a reduction in default volumes of 1% to 1.2% during the low-fee period, very much in line with my structural estimates of 4.2% and 1.6% respectively[34]. By contrast my estimates from Chapter 3 of the effects of the HARP policy are considerably larger than my structural estimates would suggest, implying an increase in

[32]I obtain these figures by dividing the estimated coefficients from Models 2 and 4 from Table A17 in Appendix E, .0015 and .0016 respectively, by the average probability of refinancing, 2.65% over the applicable period, and converting all figures to CPR.

[33]The average baseline guarantee fee is 23.5 bp in the pre-recession period and 29.5 bp thereafter, relative to 45 and 5 bp respectively in the counterfactual, and market shares of respectively 33.2% and 59.2%.

[34]My estimates from Chapter 2 suggest that between 1.9% and 2.4% of deterred refinances result in default. I apply to the percentage changes in refinancing volume relative to the baseline in the countercyclical guarantee fee simulation (a 3% reduction during the high-fee period and a 5% increase during the low-fee period) to predict the change in the volume of mortgage refinances, multiply this volume by the estimated elasticities of 1.9% and 2.4%, and divide by the predicted volumes of defaults in the counterfactual scenario to arrive at these figures.

refinancing probabilities between 26.9% and 31.7% rather than the simulated increase of 6.3%. However, the sample I use to simulate the effects of the HARP program does not consist entirely of HARP-eligible borrowers, and indeed many borrowers in the matched dataset have loans originated after June 2009 or current LTV ratios below the 75% cutoff I use for eligibility. If I apply the estimated effect of the HARP program from Chapter 2 to the fraction of borrowers in the matched dataset who have loans originated prior to June 2009 with current LTV ratios above 75%[35] I obtain an estimated increase in refinancing volume of betwen 5.8% and 6.9%, very much in line with my structural estimates. While this back-of-the-envelope calculation does not adjust for the potential effects of burnout[36] or differences in sample composition[37] it nonetheless offers some confidence in the quantitative magnitude of my structural estimates of the effect of the HARP program. While for policy analysis I place greater weight on the results of my counterfactual simulations, as they account structurally for fundamental issues of regime invariance, it is reassuring that the align well quantitatively with the reduced-form evidence presented in Chapters 2 and 3.

4.4.2 Simulation Caveats

While the results of the the simulation study from Section 4.4.1 are indicative of the potential effects of various alternative GSE policies, both the quantitative and qualitative results are subject to several important caveats regarding their direct applicability to actual housing finance policy. First, the as a partial equilibrium model the simulation

[35] Such borrowers represent 21.7% of the the post-June 2009 monhtly observations in the matched dataset.

[36] Because borrowers can only refinance once under HARP, the long-term impact of the program on refinancing volume is considerably less than the short-term impact, as eligible borrowers will refinance soon after implementation but will thereafter be ineligible, causing my estimates of short-term takeup in Chapter 3 to be large relative the long-term structural estimates.

[37] The sample used to estimate takeup in Chapter 3 consists of borrowers with loans originated between mid-2008 and mid-2010, and in many respects these borrowers had quite different credit risk characteristics from the full sample of loans originated between 2000 and 2012 in the matched dataset. In particular the average LTV ratio in the sample used in Chapter 3 is about 68, while the average for loans in the matched dataset used for estimation in Chapter 4 is 63, hence relatively fewer loans in the matched dataset will be treated as HARP eligible.

does not provide for feedback effects between policy, refinancing activity, and other exogenous variables such as house prices and unemployment rates. Moreover, within the simulation it is assumed that the GSE's policy response is unanticipated, potentially complicating the application of these policies in the real world. Second, the simulation model makes little or no adjustment for competitive conditions in lending markets, and indeed provides no estimate of the welfare of lenders and MBS investors, implying that the true welfare impact of each policy should be closer to zero than the estimated impact. In this section, I discuss both of these issues in detail.

The first caveat with the results from Section 4.4.1 concerns the partial equilibrium nature of my analysis. In my baseline simulations, I assume that house prices, employment, and interest rates do not respond to changes in refinancing activity of mortgage equity withdrawal driven by supply-side policy interventions. In reality, there is substantial evidence that increases in credit supply have a positive effect on house prices [4,78] and employment [166] and vice-versa. To the extent that this additional second-order channel goes unaccounted for in my simulations, I will tend to understate the effects of policies that affect credit supply in both the positive and negative directions. As such, we might expect a HARP-type policy to have more stimulative effects and a tighter LTV ratio cap to have more inhibatory effects than what I measure in Table 4.2, and likewise for raising or lowering guarantee fees. However, because the measured effects of each of these policies on refinancing activity is relatively modest, I would expect these second-order effects through general equilibrium channels to be modest in magnitude. Moreover, throughout my simulations I assume that all policy interventions are fully unanticipated. If these policy measures were announced in advance, we might expect a gradual adjustment to the new equilibrium and thus little long-term impact on borrower or lender behavior. However, in reality the actual housing finance policies that were implemented during this time period, summarized in Figure 1.4 in Section 1.2.3, were by and large unanticipated and were announced only just prior to their implementation. As such, I do not believe that the assumption of no prior anticipation is unreasonable. Of greater concern for my simulation of a countercyclical guarantee fee policy is the issue that such an ongoing shift in the GSE's policy stance might prompt a permanent equilibrium response; for example, borrowers might anticipate receiving lower refinance interest rates during recessions and therefore delay refinancing until a time at which guarantee fees decline. While it is beyond the scope of my analysis to model such dynamic incentives, it is worth noting that my simulation results may for this reason overstate the effect of a permanent shift in GSE policy, and actual policy changes would have to deal with a variety of practical considerations that I cannot address.

The second caveat with the results from Section 4.4.1 deals with the exclusion of structural effects on lending market. As discussed in Section 2.3.1 the structure of the mortgage market, and in particular the degree of market concentration, exert a significant influence on lender conduct and thus on mortgage interest rates and refinancing activity. In my simulations I assume that lender conduct is well-characterized by the model and that the parameters governing such conduct are held fixed across various policy scenarios. However, competition among lenders is not explicitly microfounded in the model; as discussed in Sections 4.2 and 4.3.2 I assume that lenders compete against themselves in setting interest rates, constrained only by the ability of borrowers to reject high-interest-rate loan offers. The propensity of borrowers to reject loan offers is taken directly from the data, and while I include county-level fixed effects in order to proxy for the degree of market competition under the assumption that borrowers in more competitive markets will be more willing on the margin to reject an offer and apply to a different lender, as a consequence competition among lenders is modeled in a reduced-form fashion. Scharfstein and Sunderam (2013)[178] show that competition in mortgage lending varies substantially over time, and to the extent that changes in housing finance policy have a noticeable effect on market structure my simulations will fail to capture the resulting effects on lender behavior and borrower welfare. To give a concrete example, as shown in Table 2.5 in Section 2.3.1 non-bank mortgage lenders tend to respond more to changes in guarantee fees, as they tend to be less able to hold originated loans in portfolio and are hence more reliant on GSE securitization. We might expect that changes in guarantee fees would thus have a greater impact on non-bank mortgage originators than on banks, and as these lenders tend to be smaller than banks on average, changes in guarantee fees may have important effects on the size distribution of mortgage lending. Due to the data limitations discussed in Sections 4.2 and 4.3.2 I am unable to address this concern directly, but I expect the effect of small changes in guarantee fees on mortgage market structure to be modest. Of more direct concern is the fact that my welfare calculations take no account of producer welfare, measured as the sum of originator and MBS investor profits. FRM refinancing involves a direct transfer from the MBS holder to the mortgage borrower, hence a measure of the welfare impact of policy changes on borrowers that operate through the refinancing channel will tend to over-state the effects both positive and negative on total welfare. Moreover any changes in guarantee fees that are not fully passed-through to borrowers will result in lower originator profits overall, and I do not account for this effect in my welfare calculations. Because MBS investors tend to be wealthier overall, it is likely that a transfer from investors to borrowers results in a net gain in welfare, although it is worth bearing in mind that my estimate of the effect of policy interventions on welfare measures only the impact on borrowers.

The results of the simulations presented in Section 4.4.1 suggest a clear role for countercyclical guarantee GSE fee policies. A revenue-enhancing alternative policy can achieve greater welfare for borrowers through two mechanisms. First, by shifting the timing of changes in refinancing credit supply, a countercyclical fee policy can result in greater overall borrower welfare despite lower total refinancing volume by making credit more available at times when it is especially valuable. Second, by reducing over-borrowing during expansionary periods, a countercyclical guarantee fee policy reduces default rates and leaves borrowers better positioned to refinance when interest rates decline and unemployment rates rise. The result is that the deleveraging associated with recessions is mitigated by the alternative policy. The simulations suggest that a countercyclical GSE fee policy would be roughly as effective as the HARP program at improving borrower welfare through the former mechanism, and substantially more effective than a tighter LTV ratio cap at improving borrower welfare through the latter mechanism. Indeed, these simulations indicate that a countercylical guarantee fee policy would manage to combine the benefits of both policies while increasing total GSE revenue. Moreover the key limitation of my counterfactual analysis, that it features no general equilibrium effects, likely implies that my estimates of the effect of a countercyclical fee policy represent a lower bound on the true effects. In equilibrium, a significant shift in credit supply and refinancing activity achieved through changes in guarantee fees would also likely lead to changes in house prices [4,78] and employment [166] that would amplify the true costs and benefits. Moreover my structural model is able by design to simulate the effects other potential alternative policies, particularly policies that alter secondary market conditions such as changes in GSE lending standards, large-scale asset purchase programs intended to increase the prices of mortgage securities, or even the elimination of the GSEs themselves[38]. While it is beyond the scope of this analysis to simulate the effects of such policies, the framework developed in this chapter may prove useful for future policy evaluations.

[38]Each of these policies have clear analogues in my model. One could simulate changes to GSE credit standards, for example refusing to accept loans from borrowers with FICO scores below a certain threshold, simply by removing the GSE securitization option for originators extending loans to such borrowers, and the effect of eliminating the GSEs entirely would be analogous for all borrowers. MBS prices factor directly into originator securitization profits, hence it is straightforward to simulate the effects of increasing them by some fixed amount.

4.5 Conclusion

This chapter analyzes the effects of the GSE's guarantee fee policies on credit provision and borrower welfare across the business cycle. In Section 4.2 I extend the dynamic choice literature by designing a novel partial equilibrium model of the agency refinancing market featuring dynamic borrower behavior and endogenous credit constraints, which I then estimate in Section 4.3. By incorporating latent liquidity preference, I show that borrowers derive 130% to 220% greater utility from cash-on-hand, and hence from refinancing, when they are likely to be income-constrained. In Section 4.4 I then analyze via simulation the effects of an alternative countercyclical fee policy that reduces credit supply during market expansions and increases credit supply during recessions. I show that although the policy improves GSE revenue by 6.3% and results in a 1.2% lower total volume of home equity withdrawal, borrower welfare is still 2.3% higher under the alternative than under the baseline. The increase in borrower welfare derives from two sources. First, by reducing credit supply during the expansionary period, the alternative policy reduces ex-ante overborrowing and thus future default. Second, by improving credit supply during the recession, the program stimulates borrowing at exactly the time when it is most valuable, leading to a 5% increase in refinancing volume during and after the recession. By comparing this policy with other proposed alternative, I find that countercylical guarantee fee pricing would combine the benefits of a policy of tighter LTV limits and an affordable refinance policy, having benefits roughly equivalent to the former in reducing over-leveraging and the latter in stimulating refinancing after the recession.

These findings complement the results from Chapters 2 and 3 in two key ways. First, the model of credit supply I develop endogenizes and microfounds the pass-through rate from guarantee fees to interest rates and credit approval decisions that I estimate in

Chapter 2. These structural estimates allow for out-of-sample predictions about the effects of changes in guarantee fees under counterfactual GSE policy regimes. Second, the borrower preference parameters I estimate provide a model for behavior behind the suggestive evidence I present in Chapter 3 regarding the interaction between unemployment, credit constraints, and refinancing. These preference parameters allow for a quantitative analysis of the welfare implications of alternative GSE policies.

To my knowledge the question of how best to set GSE guarantee fees across the business cycle has not been considered previously in either the economics or policy literature. The collapse of the private-label market in 2008 and subsequent ascendance of the GSEs revealed how important their role remains in backstopping the refinancing market during period of credit market turmoil. With housing finance reform currently being debated in congress, the question of how the GSEs should vary their policies with the business cycle has never been more relevant. The alternative countercylical guarantee fee policy analyzed in this paper functions as a means to achieve macroprudential controls on credit growth during credit expansions while still functioning as a "refinancing lender of last resort". Moreover because of their low impact on non-mortgage interest rates or other economic sectors aside from housing and mortgage finance, guarantee fees are in some ways an ideal "second instrument" for regulating credit macroprudentially. Policy makers may wish to consider revenue-neutral policies of the sort analyzed in this chapter that complement the agencies' traditional role while preserving key elements of the existing housing finance system.

Bibliography

[1] Abdallah, C. S., & Lastrapes, W. D. (2012). Home Equity Lending and Retail Spending: Evidence From a Natural Experiment in Texas. *American Economic Journal: Macroeconomics*, 4(4), 94-125.

[2] Abruzzo, N., & Park, Y. H. (2014). An Empirical Analysis of Futures Margin Changes: Determinants and Policy Implications. Federal Reserve Board FEDS Working Paper (2014-86).

[3] Acharya, V. V., Richardson, M., Van Nieuwerburgh, S., & White, L. J. (2011). *Guaranteed to Fail: Fannie Mae, Freddie Mac, and the Debacle of Mortgage Finance*. Princeton University Press.

[4] Adelino, M., Schoar, A., & Severino, F. (2012). Credit Supply and House Prices: Evidence from Mortgage Market Segmentation. National Bureau of Economic Research (No. w17832).

[5] Agarwal, S., Chang, Y., & Yavas, A. (2012). Adverse Selection in Mortgage Securitization. *Journal of Financial Economics*, 105(3), 640-660.

[6] Agarwal, S., Amromin, G., Ben-David, I., Chomsisengphet, S., Piskorski, T., & Seru, A. (2012). Policy Intervention in Debt Renegotiation: Evidence from the Home Affordable Modification Program. National Bureau of Economic Research (No. w18311).

[7] Agarwal, S., Rosen, R. J., & Yao, V. (2013). Why Do Borrowers Make Mortgage Refinancing Mistakes? Federal Reserve Bank of Chicago Working Paper (2013-02).

[8] Agarwal, S., & Evanoff, D. D. (2013). Loan Product Steering in Mortgage Markets.

[9] Agarwal, S., Driscoll, J. C., & Laibson, D. I. (2013). Optimal Mortgage Refinancing: A Closed-Form Solution. Journal of Money, Credit and Banking, 45(4), 591-622.

[10] Agarwal, S., & Qian, W. (2014). Access to Home Equity and Consumption: Evidence from a Policy Experiment.

[11] Agarwal, S., Chomsisengphet, S., Mahoney, N., & Stroebel, J. (2015). Do Banks Pass Through Credit Expansions? The Marginal Profitability of Consumer Lending During the Great Recession. National Bureau of Economic Research (No. w21567).

[12] Agarwal, S., Amromin, G., Ben-David, I., Chomsisengphet, S., Piskorski, T., Seru, A. & Yao, V. (2015). Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinancing Program. National Bureau of Economic Research (No. w21512).

[13] Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, 109(3), 659-684.

[14] Allen, J., Clark, R., & Houde, J. F. (2013). The Effect of Mergers in Search Market: Evidence from the Canadian Mortgage Industry. National Bureau of Economic Research (No. w19126).

[15] Allen, J., Clark, R., & Houde, J. F. (2014). Price Dispersion in Mortgage Markets. *The Journal of Industrial Economics*, 62(3), 377-416..

[16] Ambrose, B. W., LaCour-Little, M., & Sanders, A. B. (2004). The Effect of Conforming Loan Status on Mortgage Yield Spreads: A Loan Level Analysis. *Real Estate Economics*, 32(4), 541-569.

[17] Amromin, G., & Kearns, C. (2014). Access to Refinancing and Mortgage Interest Rates: HARPing on the Importance of Competition. Federal Reserve Bank of Chicago Working Paper (2014-25).

[18] An, X., Clapp, J. M., & Deng, Y. (2010). Omitted Mobility Characteristics and Property Market Dynamics: Application to Mortgage Termination. *The Journal of*

Real Estate Finance and Economics, 41(3), 245-271.

[19] Antoniades, A. (2015). Liquidity Risk and the Credit Crunch of 2007-2008: Evidence From Micro-Level Data on Mortgage Loan Applications. *Journal of Financial and Quantitative Analysis*, forthcoming

[20] Andersen, S., Campbell, J. Y., Meisner-Nielsen, K., & Ramadorai, T. (2014). Inattention and Inertia in Household Finance: Evidence from the Danish Mortgage Market.

[21] Avery, R. B., & Brevoort, K. P. (2015). The Subprime Crisis: Is Government Housing Policy to Blame? *Review of Economics and Statistics*, 97(2), 352-363.

[22] Bajari, P., Chu, C. S., Nekipelov, D., & Park, M. (2013). A Dynamic Model of Subprime Mortgage Default: Estimation and Policy Implications. National Bureau of Economic Research (No. w18850).

[23] Balla, E., & McKenna, A. B. (2009). Dynamic Provisioning: A Countercyclical Tool for Loan Loss Reserves. Federal Reserve Bank of Richmond Economic Quarterly, 95(4), 383-418.

[24] Bank of England. (2011). Instruments of Macroprudential Policy. Bank of England Discussion Paper.

[25] Bartik, Timothy J. (1991). *Who Benefits from State and Local Economic Development Policies?* Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

[26] Barwick, P. J., & Pathak, P. A. (2011). The Costs of Free Entry: An Empirical Study of Real Estate Agents in Greater Boston. National Bureau of Economic Research (No. w17227).

[27] Baum, L. E., Petrie, T., Soules, G., & Weiss, N. (1970). A Maximization Technique Occurring in the Statistical Analysis of Probabilistic Functions of Markov Chains. *The Annals of Mathematical Statistics*, 41(1), 164-171.

[28] Bayer, P., Ferreira, F., & Ross, S. L. (2013). The Vulnerability of Minority Homeowners in the Housing Boom and Bust. National Bureau of Economic

Research (No. w19020).

[29] Bazdresch, S., Kahn, R. J., & Whited, T. M. (2014). Empirical Policy Function Benchmarks for Evaluation and Estimation of Dynamic Models.

[30] Beltran, D. O., & Thomas, C. P. (2010). Could Asymmetric Information Alone Have Caused the Collapse of Private-Label Securitization? Federal Reserve Board International Finance Discussion Paper Series (2010-1010).

[31] Benito, A. (2009). Who Withdraws Housing Equity and Why? *Economica*, 76(301), 51-70.

[32] Bennett, P., Peach, R., & Peristiani, S. (2001). Structural Change in the Mortgage Market and the Propensity to Refinance. *Journal of Money, Credit and Banking*, 33(4), 955-975.

[33] Beraja, M., Fuster, A., Hurst, E., & Vavra, J. (2015). Regional Heterogeneity and Monetary Policy. Federal Reserve Bank of New York Staff Report (No. 731).

[34] Bernanke, B. (2007, June 15). *The Financial Accelerator and the Credit Channel*. Speech presented at The Credit Channel of Monetary Policy in the Twenty-First Century Conference, Federal Reserve Bank of Atlanta, Atlanta, GA.

[35] Bethune, Z. (2014). Consumer Credit, Unemployment, and Aggregate Labor Market Dynamics.

[36] Bhutta, N., Dokko, J., & Shan, H. (2010). The Depth of Negative Equity and Mortgage Default Decisions. Federal Reserve Board FEDS Working Paper (2010-35).

[37] Bhutta, N. (2012). Mortgage Debt and Household Deleveraging: Accounting for the Decline in Mortgage Debt Using Consumer Credit Record Data. Federal Reserve Board FEDS Working Paper (2012-14).

[38] Bhutta, N., & Keys, B. J. (2014). Interest Rates and Equity Extraction During the Housing Boom.

[39] Black, L., Hancock, D., & Passmore, W. (2010). The Bank Lending Channel of

Monetary Policy and Its Effect on Mortgage Lending. Federal Reserve Board FEDS Working Paper (2010-39).

[40] Blevins, J. R. (2010). Nonparametric Identification of Dynamic Games with Discrete and Continuous Choices.

[41] Bolotnyy, V. (2014). The Government-Sponsored Enterprises and the Mortgage Crisis: The Role of the Affordable Housing Goals. *Real Estate Economics*, 42(3), 724-755.

[42] Bostic, R. W., & Gabriel, S. A. (2006). Do the GSEs Matter to Low-Income Housing Markets? An Assessment of the Effects of the GSE Loan Purchase Goals on California Housing Outcomes. *Journal of Urban Economics*, 59(3), 458-475.

[43] Boyce, A., Hubbard, G., Mayer, C., & Witkin, J. (2012). Streamlined Refinancings for Up To 14 Million Borrowers. Columbia Business School Working Paper.

[44] Brayton, F., & Tinsley, P. A. (1996). A Guide to FRB/US: A Macroeconomic Model of the United States. Federal Reserve Board FEDS Working Paper (1996-42).

[45] Brown, M., Haughwout, A., Lee, D., & Van der Klaauw, W. (2013). The Financial Crisis at the Kitchen Table: Trends in Household Debt and Credit. *Current Issues in Economics and Finance*, 19(2), 1-10.

[46] Bucks, B., & Pence, K. (2008). Do Borrowers Know Their Mortgage Terms? *Journal of Urban Economics*, 64(2), 218-233.

[47] Cahill, M. E., D'Amico, S., Li, C., & Sears, J. S. (2013). Duration Risk Versus Local Supply Channel in Treasury Yields: Evidence from the Federal Reserve's Asset Purchase Announcements. Federal Reserve Board FEDS Working Paper (2013-35).

[48] Calomiris, C. W., & Haber, S. H. (2015). *Fragile by Design: The Political Origins of Banking Crises and Scarce Credit*. Princeton University Press.

[49] Campbell, J. Y. (2013). Mortgage Market Design. *Review of Finance*, 17(1), 1-33.

[50] Carroll, C. D., Dynan, K. E., & Krane, S. D. (2003). Unemployment Risk and Precautionary Wealth: Evidence from Households' Balance Sheets. *Review of*

Economics and Statistics, 85(3), 586-604.

[51] Carroll, C., Slacalek, J., Tokuoka, K., & White, M. N. (2014). The Distribution of Wealth and the Marginal Propensity to Consume.

[52] Congressional Budget Office (2010). Fannie Mae, Freddie Mac, and the Federal Role in the Secondary Mortgage Market. CBO Pub. No. 4021, December 2010.

[53] Chen, H., Michaux, M., & Roussanov, N. (2013). Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty. National Bureau of Economic Research (No. w19421).

[54] Chetty, R., Looney, A., & Kroft K. (2009). Salience and Taxation: Theory and Evidence. *American Economic Review*, 99(4), 1145-1177.

[55] Chung, H., Loken, E., & Schafer, J. L. (2004). Difficulties in Drawing Inferences with Finite-Mixture Models. *The American Statistician*, 58(2), 152-158.

[56] Collins, B. "Why 'G-Fees' Have Taken Center Stage in Housing Finance Debate." *National Mortgage News*, 19 Aug. 2014. Print.

[57] Collins, B. (2014, August 14). Why 'G-Fees' Have Taken Center Stage in Housing Finance Debate. *National Mortgage News*. Retrieved from www.nationalmortgagenews.com

[58] Connault, B. (2014). Hidden Rust Models.

[59] Cooper, D. (2013). House Price Fluctuations: The Role of Housing Wealth as Borrowing Collateral. *Review of Economics and Statistics*, 95(4), 1183-1197.

[60] Crawford, G. S., & Yurukoglu, A. (2012). The Welfare Effects of Bundling in Multichannel Television Markets. *American Economic Review*, 102(2), 643-85.

[61] Crawford, G. S., Pavanini, N., & Schivardi, F. (2014). Asymmetric Information and Imperfect Competition in Lending Markets. Centre for Economic Policy Research Discussion Paper Series (No. 10473).

[62] Crossley, T. F., & Low, H. W. (2014). Job Loss, Credit Constraints, and Consumption Growth. *Review of Economics and Statistics*, 96(5), 876-884.

[63] Dechario, T., Mosser, P., Tracy, J., Vickery, J., & Wright, J. (2011). A Private Lender Cooperative Model for Residential Mortgage Finance. in Marvin M. Smith & Susan M. Wachter (Eds.), *The American Mortgage System: Crisis and Reform* (pp. 286-304). University of Pennsylvania Press.

[64] DeFusco, A. A., & Paciorek, A. (2014). The Interest Rate Elasticity of Mortgage Demand: Evidence From Bunching at the Conforming Loan Limit. Federal Reserve Board FEDS Working Paper (2014-11).

[65] Di Maggio, M., Kermani, A., & Ramcharan, R. (2014). Monetary Policy Pass-Through: Household Consumption and Voluntary Deleveraging.

[66] Diebold, F. X., Lee, J. H., & Weinbach, G. C. (1999). Regime Switching with Time-Varying Transition Probabilities. Francis X. Diebold & Glenn D. Rudebusch (Eds.), *Business Cycles: Durations, Dynamics and Forecasting* (pp. 144-168). Princeton University Press.

[67] Disney, R., & Gathergood, J. (2011). Financial Literacy and Indebtedness: New Evidence for U.K. Consumers.

[68] Downing, C., Jaffee, D., & Wallace, N. (2009). Is the Market for Mortgage-Backed Securities a Market for Lemons? *Review of Financial Studies*, 22(7), 2457-2494.

[69] Duca, J. V., & Kumar, A. (2014). Financial Literacy and Mortgage Equity Withdrawals. *Journal of Urban Economics*, 80(1), 62-75.

[70] Duca, J. V., Muellbauer, J., & Murphy, A. (2011). House Prices and Credit Constraints: Making Sense of the U.S. Experience. *The Economic Journal*, 121(552), 533-551.

[71] Dynan, K., & Gayer, T. (2011). The Government's Role in the Housing Finance System: Where Do We Go from Here? Brookings Institution Paper, February 2011.

[72] Dynarski, S., Gruber, J., Moffitt, R. A., & Burtless, G. (1997). Can Families Smooth Variable Earnings? *Brookings Papers on Economic Activity*, 1997(1), 229-303.

[73] Eggertsson, G. B., & Krugman, P. (2012). Debt, Deleveraging, and the Liquidity

Trap: A Fisher-Minsky-Koo Approach. *The Quarterly Journal of Economics*, 127(3), 1469-1513.

[74] Einav, L., Jenkins, M., & Levin, J. (2012). Contract Pricing in Consumer Credit Markets. *Econometrica*, 80(4), 1387-1432.

[75] Elenev, V., Landvoigt, T., & Van Nieuwerburgh, S. (2015). Phasing Out the GSEs. National Bureau of Economic Research (No. w21626).

[76] Ellen, I. G., Tye, J.N. & Willis, M.A. (2010). Improving US Housing Finance Through Reform of Fannie Mae and Freddie Mac: Assessing the Options. Urban Institute Housing Finance Policy Center Commentary (May 2010).

[77] Elul, R., Souleles, N. S., Chomsisengphet, S., Glennon, D., & Hunt, R. (2010). What "Triggers" Mortgage Default? *The American Economic Review*, 100(2), 490-94.

[78] Favara, G., & Imbs, J. (2015). Credit Supply and the Price of Housing. *The American Economic Review*, 105(3), 958-992.

[79] Favilukis, J., Kohn, D., Ludvigson, S. C., & Van Nieuwerburgh, S. (2012). International Capital Flows and House Prices: Theory and Evidence. National Bureau of Economic Research (No. w17751).

[80] Federal Housing Finance Agency Office of Inspector General. (2015). The Continued Profitability of Fannie Mae and Freddie Mac Is Not Assured. Washington, DC: FHFA OIG (White Paper Report 2015-001).

[81] Finberg, J. (2003). Financial Abuse of the Elderly in California. *Loyola of Los Angeles Law Review*, 36(2), 667-1753.

[82] Flood, G., Fischetti, P., and Swift, D. (2014). The Return of Private Capital. Fannie Mae Housing Insights, Vol. 4, Issue 7.

[83] Frame, W. S., Fuster, A., Tracy, J., & Vickery, J. (2015). The Rescue of Fannie Mae and Freddie Mac. Federal Reserve Bank of New York Staff Report (No. 719).

[84] Fuster, A., Goodman, L., Lucca, D., Madar, L., Molloy, L., & Willen, P. (2013) The Rising Gap between Primary and Secondary Mortgage Rates. *Federal Reserve Bank*

of New York Economic Policy Review, 19(2), 17-40.

[85] Fuster, A., & Vickery, J. (2013). Securitization and the Fixed-Rate Mortgage. Federal Reserve Bank of New York Staff Report (No. 594).

[86] Gabriel, S. A., & Rosenthal, S. S. (2007). Secondary Markets, Risk, and Access to Credit: Evidence From the Mortgage Market.

[87] Gabriel, S. A., & Rosenthal, S. S. (2010). Do the GSEs Expand the Supply of Mortgage Credit? New Evidence of Crowd Out in the Secondary Mortgage Market. *Journal of Public Economics*, 94(11), 975-986.

[88] Gale, D., & Shapley, L. S. (1962). College Admissions and the Stability of Marriage. *American Mathematical Monthly*, 120(5), 386-391.

[89] Gan, J. (2010). Housing Wealth and Consumption Growth: Evidence from a Large Panel of Households. *Review of Financial Studies*, 23(6), 2229-2267.

[90] Government Accountability Office (2009). Fannie Mae and Freddie Mac: Analysis of Options for Revising the Housing Enterprises' Long-term Structures. GAO-09-782, September 2009.

[91] Gerardi, K., Shapiro, A. H., & Willen, P. S. (2008). Subprime Outcomes: Risky Mortgages, Homeownership Experiences, and Foreclosures. Federal Reserve Bank of Boston Working Paper (No. 07-15).

[92] Gerardi, K., Herkenhoff, K. F., Ohanian, L. E., & Willen, P. S. (2013). Unemployment, Negative Equity, and Strategic Default. Federal Reserve Bank of Atlanta Working Paper Series (2013-4).

[93] Gerardi, K., Herkenhoff, K. F., Ohanian, L. E., & Willen, P. S. (2015). Can't Pay or Won't Pay? Unemployment, Negative Equity, and Strategic Default. National Bureau of Economic Research (No. w21630).

[94] Gete, P., & Zecchetto, F. (2015). Distributional Implications of Government Guarantees in Mortgage Markets.

[95] Goodman, L. S., Li, W. (2014). A Better Measure of Mortgage Application Denial

Rates. Urban Institute Housing Finance Policy Center Commentary (December 2014).

[96] Goodman, L. S., Seidman, E., Parrott, J., & Zhu, J. (2014). Guarantee Fees - An Art, Not a Science. Urban Institute Housing Finance Policy Center Commentary (August 2014).

[97] Goodman, L. S. (2014). A Realistic Assessment of Housing Finance Reform. Urban Institute Housing Finance Policy Center Commentary (August 2014).

[98] Gordy, M. B., & Howells, B. (2006). Procyclicality in Basel II: Can We Treat the Disease Without Killing the Patient? *Journal of Financial Intermediation*, 15(3), 395-417.

[99] Gourinchas, P. O., & Parker, J. A. (2002). Consumption Over the Life Cycle. *Econometrica*, 70(1), 47-89.

[100] Green, D. (2013, May 16). FHA Mortgage Rates: New MIP Costs Harm "Effective" FHA Loan Rates. *The Mortgage Reports*. Retrieved from www.themortgagereports.com

[101] Greenspan, A., & Kennedy, J. (2008). Sources and Uses of Equity Extracted from Homes. *Oxford Review of Economic Policy*, 24(1), 120-144.

[102] Gropp, R., Krainer, J., & Laderman, E. (2014). Did Consumers Want Less Debt? Consumer Credit Demand Versus Supply in the Wake of the 2008-2009 Financial Crisis. Federal Reserve Bank of San Francisco Working Paper (2014-08).

[103] Gross, D. B., & Souleles, N. S. (2002). Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data. *The Quarterly Journal of Economics*, 117(1), 149-185.

[104] Guerrieri, V., & Lorenzoni, G. (2011). Credit Crises, Precautionary Savings, and the Liquidity Trap. National Bureau of Economic Research (No. w17583).

[105] Gurun, U. G., Matvos, G., & Seru, A. (2013). Advertising Expensive Mortgages. National Bureau of Economic Research (No. w18910).

[106] Gyourko, J., Saiz, A., & Summers, A. (2008). A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index. *Urban Studies*, 45(3), 693-729.

[107] Gyourko, J., & Tracy, J. (2014). Reconciling Theory and Empirics on the Role of Unemployment in Mortgage Default. *Journal of Urban Economics*, 80(1), 87-96.

[108] Hancock, D., & Passmore, W. (2011). Catastrophic Mortgage Insurance and the Reform of Fannie Mae and Freddie Mac. In Martin Neil Baily (Ed.), *The Future of Housing Finance: Restructuring the U.S. Residential Mortgage Market* (pp. 111-145). Brookings Institution Press.

[109] Hancock, D., & Passmore, W. (2014). How the Federal Reserve's Large-Scale Asset Purchases (LSAPs) Influence Mortgage-Backed Securities (MBS) Yields and US Mortgage Rates. Federal Reserve Board FEDS Working Paper (2014-12).

[110] Hancock, D., & Passmore, W. (2011). Macroprudential Mortgage-Backed Securitization: Can It Work? Working Paper.

[111] Heckman, J. J. (1979). Sample Selection Bias as a Specification Error. *Econometrica*, 47(1), 153-161.

[112] Hedberg, W., & Krainer, J. (2012). Borrower Access to Credit Markets Following a Mortgage Default. Federal Reserve Bank of San Francisco Economic Letter (2012-32).

[113] Hembre, E. (2014). HAMP, Home Attachment, and Mortgage Default.

[114] Hill, J. A. (2010). Bailouts and Credit Cycles: Fannie, Freddie, and the Farm Credit System. *Wisconsin Law Review*, 2010(1), 1-77.

[115] Ho, K. (2009). Insurer-Provider Networks in the Medical Care Market. *American Economic Review*, 99(1), 393-430.

[116] Hochguertel, S., Weber, G., & Alessie, R. (2005). Consumer Credit: Evidence from Italian Micro Data. *Journal of the European Economic Association*, 3(1), 144-178.

[117] Hotz, V. J., & Miller, R. A. (1993). Conditional Choice Probabilities and the

Estimation of Dynamic Models. *The Review of Economic Studies*, 60(3), 497-529.

[118] Hsu, J. W., Matsa, D. A., & Melzer, B. T. (2014). Positive Externalities of Social Insurance: Unemployment Insurance and Consumer Credit. National Bureau of Economic Research (No. w20353).

[119] Hurd, M. D., & Rohwedder, S. (2010). Effects of the Financial Crisis and Great Recession on American Households. National Bureau of Economic Research (No. w16407).

[120] Hurst, E., & Stafford, F. (2004). Home is Where the Equity Is: Mortgage Refinancing and Household Consumption. *Journal of Money, Credit and Banking*, 36(6), 985-1014.

[121] Hurst, E., Keys, B., Seru, A., & Vavra, S. (2014). Regional Risk Sharing Through the U.S. Mortgage Market.

[122] Jaffee, D., & Quigley, J. M. (2011). The Future of the Government Sponsored Enterprises: The Role for Government in the U.S. Mortgage Market. National Bureau of Economic Research (No. w17685).

[123] Jappelli, T., & Pistaferri, L. (2014). Fiscal Policy and MPC Heterogeneity. *American Economic Journal: Macroeconomics*, 6(4), 107-136.

[124] Jeske, K., Krueger, D., & Mitman, K. (2013). Housing, Mortgage Bailout Guarantees and the Macro Economy. *Journal of Monetary Economics*, 60(8), 917-935.

[125] Jiménez, G., Ongena, S., Peydró, J. L., & Saurina, J. (2012). Credit Supply and Monetary Policy: Identifying the Bank Balance-Sheet Channel with Loan Applications. *American Economic Review*, 102(5), 2301-26.

[126] Jiménez, G., Ongena, S., Peydró, J. L., & Saurina, J. (2014). Hazardous Times for Monetary Policy: What Do Twenty-Three Million Bank Loans Say About the Effects of Monetary Policy on Credit Risk-Taking? *Econometrica*, 82(2), 463-505.

[127] Jiménez, G., Mian, A. R., Peydró, J. L., & Saurina Salas, J. (2011). Local Versus Aggregate Lending Channels: The Effects of Securitization on Corporate Credit

Supply. National Bureau of Economic Research (No. w16595).

[128] Johnson, E., Meier, S., & Toubia, O. (2015). Money Left on the Kitchen Table: Exploring Sluggish Mortgage Refinancing using Administrative Data, Surveys, and Field Experiments.

[129] Johnson, K. and Li, G. (2007). Do High Debt Payments Hinder Household Consumption Smoothing? Federal Reserve Board FEDS Working Paper (2007-52).

[130] Judd, K. L. (1998). *Numerical Methods in Economics*. MIT Press.

[131] Karlan, D., & Zinman, J. (2009). Observing Unobservables: Identifying Information Asymmetries with a Consumer Credit Field Experiment. *Econometrica*, 77(6), 1993-2008.

[132] Kaufman, A. (2014). The Influence of Fannie and Freddie on Mortgage Loan Terms. *Real Estate Economics*, 42(2), 472-496.

[133] Keys, B. J., Mukherjee, T., Seru, A., & Vig, V. (2010). Did Securitization Lead to Lax Screening? Evidence from Subprime Loans. *The Quarterly Journal of Economics*, 125(1), 307-362.

[134] Keys, B. J., Seru, A., & Vig, V. (2012). Lender Screening and the Role of Securitization: Evidence from Prime and Subprime Mortgage Markets. *Review of Financial Studies*, 25(7), 2071-2108.

[135] Keys, B. J., Piskorski, T., Seru, A., & Vig, V. (2012). Mortgage Financing in the Housing Boom and Bust. In Edward L. Glaeser and Todd Sinai (Eds.), *Housing and the Financial Crisis* (pp. 143-204). University of Chicago Press.

[136] Keys, B. J., Piskorski, T., Seru, A., & Yao, V. (2014). Mortgage Rates, Household Balance Sheets, and the Real Economy. National Bureau of Economic Research (No. w20561).

[137] Khandani, A. E., Lo, A. W., & Merton, R. C. (2013). Systemic Risk and the Refinancing Ratchet Effect. *Journal of Financial Economics*, 108(1), 29-45.

[138] Kimball, M. S. (2012). Getting the Biggest Bang for the Buck in Fiscal Policy.

National Bureau of Economic Research (No. w18142).

[139] Krainer, J., & Laderman, E. (2011). Prepayment and Delinquency in the Mortgage Crisis Period. Federal Reserve Bank of San Francisco Working Paper (2011-25).

[140] Kroft, K., & Notowidigdo, M. J. (2011). Should Unemployment Insurance Vary with the Unemployment Rate? Theory and Evidence. National Bureau of Economic Research (No. w17173).

[141] Krusell, P., & Smith, Jr, A. A. (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy*, 106(5), 867-896.

[142] Lacko, J. M., & Pappalardo, J. K. (2010). The Failure and Promise of Mandated Consumer Mortgage Disclosures: Evidence from Qualitative Interviews and a Controlled Experiment with Mortgage Borrowers. *American Economic Review*, 100(2), 516-21.

[143] Laderman, E. (2012). Mortgage Prepayment: An Avenue Foreclosed? Federal Reserve Bank of San Francisco Economic Letter (2012-05).

[144] Landais, C., Michaillat, P., & Saez, E. (2014). A Macroeconomic Theory of Optimal Unemployment Insurance.

[145] Laufer, S. (2013). Equity Extraction and Mortgage Default. Federal Reserve Board FEDS Working Paper (2013-30).

[146] Lee, Y. H. A., & Ko, K. J. (2011). Consumer Mistakes in the Mortgage Market: Choosing Unwisely Versus Not Switching Wisely. *University of Pennsylvania Journal of Business Law*, 14(2), 417-462.

[147] Lehnert, A., Passmore, W., & Sherlund, S. M. (2008). GSEs, Mortgage Rates, and Secondary Market Activities. *The Journal of Real Estate Finance and Economics*, 36(3), 343-363.

[148] Leth-Petersen, S. (2010). Intertemporal Consumption and Credit Constraints: Does Total Expenditure Respond to an Exogenous Shock to Credit? *The American Economic Review*, 100(3), 1080-1103.

[149] Loutskina, E., & Strahan, P. E. (2009). Securitization and the Declining Impact of Bank Finance on Loan Supply: Evidence from Mortgage Originations. *The Journal of Finance*, 64(2), 861-889.

[150] Loutskina, E., & Strahan, P. E. (2011). Informed and Uninformed Investment in Housing: The Downside of Diversification. *Review of Financial Studies*, 24(5), 1447-1480.

[151] Loutskina, E. (2011). The Role of Securitization in Bank Liquidity and Funding Management. *Journal of Financial Economics*, 100(3), 663-684.

[152] Lustig, H. N., & Van Nieuwerburgh, S. G. (2005). Housing Collateral, Consumption Insurance, and Risk Premia: An Empirical Perspective. *The Journal of Finance*, 60(3), 1167-1219.

[153] Lustig, H., & Van Nieuwerburgh, S. (2010). How Much Does Household Collateral Constrain Regional Risk Sharing? *Review of Economic Dynamics*, 13(2), 265-294.

[154] Ma, C. (2014). Estimating a Single Agent Dynamic Discrete Choice Model for Household Mortgage Default and Prepayment Decisions.

[155] Markwardt, K., Martinello, A., Sandor, L. (2014). Liquidity Substitutes for Public Unemployment Insurance: Evidence from the Introduction of Home Equity Loans in Denmark.

[156] Mayer, C., Morrison, E., Piskorski, T., & Gupta, A. (2014). Mortgage Modification and Strategic Behavior: Evidence from a Legal Settlement with Countrywide. *American Economic Review*, 104(9), 2830-2857.

[157] Merrill, C., Nadauld, T., & Strahan, P. (2014). Final Demand for Structured Finance Securities.

[158] Mian, A. R., & Sufi, A. (2009). The Consequences of Mortgage Credit Expansion: Evidence from the US Mortgage Default Crisis. *The Quarterly Journal of Economics*, 124(4), 1449-1496.

[159] Mian, A., & Sufi, A. (2011). House Prices, Home Equity-Based Borrowing, and the

US Household Leverage Crisis. *American Economic Review*, 101(5), 2132-56.

[160] Mian, A. R., & Sufi, A. (2012). What Explains High Unemployment? The Aggregate Demand Channel. National Bureau of Economic Research (No. w17830).

[161] Mian, A., Rao, K., & Sufi, A. (2013). Household Balance Sheets, Consumption, and the Economic Slump. *The Quarterly Journal of Economics*, 128(4), 1687-1726.

[162] Mian, A., & Sufi, A. (2014). House Price Gains and US Household Spending from 2002 to 2006. National Bureau of Economic Research (No. w20152).

[163] Michel-Kerjan, E., & Kunreuther, H. (2011). Redesigning Flood Insurance. *Science*, 333(6041), 408-409.

[164] Miller, N. H., Osborne, M., & Sheu, G. (2015). Pass-Through in a Concentrated Industry: Empirical Evidence and Regulatory Implications.

[165] Mishkin, F. S. (2007). Housing and the Monetary Transmission Mechanism. National Bureau of Economic Research (No. w13518).

[166] Mondragon, J. (2014). Household Credit and Employment in the Great Recession.

[167] Mora, N. (2014). The Weakened Transmission of Monetary Policy to Consumer Loan Rates. Federal Reserve Bank of Kansas City Economic Review (Q1 2014), 94-117.

[168] Mosser, P. C., Tracy, J., & Wright, J. (2013). The Capital Structure and Governance of a Mortgage Securitization Utility. Federal Reserve Bank of New York Staff Report (No. 644).

[169] Murphy, D., Vasios, M., & Vause, N. (2014). An Investigation Into the Procyclicality of Risk-Based Initial Margin Models. Bank of England Financial Stability Paper (No. 29).

[170] Passmore, W., Sherlund, S. M., & Burgess, G. (2005). The Effect of Housing Government-Sponsored Enterprises on Mortgage Rates. *Real Estate Economics*, 33(3), 427-463.

[171] Peek, J., & Wilcox, J. A. (2003). Secondary Mortgage Markets, GSEs, and the Changing Cyclicality of Mortgage Flows. *Research in Finance*, 20, 61-80.

[172] Pennacchi, G. G. (2009). Deposit Insurance. In *AEI Conference on Private Markets and Public Insurance Programs*, March 2009.

[173] Midrigan, V. & Philippon, T., (2011). Household Leverage and the Recession. National Bureau of Economic Research (No. w16965).

[174] Quercia, R. G., Pennington-Cross, A., & Yue Tian, C. (2012). Mortgage Default and Prepayment Risks among Moderate- and Low-Income Households. *Real Estate Economics*, 40(1), 159-198.

[175] Quigley, J. M. (2006). Federal Credit and Insurance Programs: Housing. Federal Reserve Bank of St. Louis Review, 88(4).

[176] Remy, M., Lucas, D., & Moore, D. (2011). An Evaluation of Large-Scale Mortgage Refinancing Programs. Congressional Budget Office Working Paper (2011-04).

[177] Scharfstein, D., & Sunderam, A. (2011). The Economics of Housing Finance Reform: Privatizing, Regulating and Backstopping Mortgage Markets. In Martin Neil Baily (Ed.), *The Future of Housing Finance: Restructuring the U.S. Residential Mortgage Market* (pp. 146-198). Brookings Institution Press.

[178] Scharfstein, D. S., & Sunderam, A. (2013). Concentration in Mortgage Lending, Refinancing Activity and Mortgage Rates. National Bureau of Economic Research (No. w19156).

[179] Shan, H. (2011). Reversing the Trend: The Recent Expansion of the Reverse Mortgage Market. *Real Estate Economics*, 39(4), 743-768.

[180] Smith, S., & Weiher, J. C. (2012). Countercyclical Capital Regime: A Proposed Design and Empirical Evaluation. Federal Housing Finance Agency Working Paper, (12-2).

[181] Stanton, R., Walden, J., & Wallace, N. (2014). The Industrial Organization of the US Residential Mortgage Market. *Annual Review of Financial Economics*, 6(1), 259-288.

[182] Stephens, M. (2000). Dealing with Label Switching in Mixture Models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(4), 795-809.

[183] Sullivan, J. X. (2008). Borrowing During Unemployment: Unsecured Debt as a Safety Net. *Journal of Human Resources*, 43(2), 383-412.

[184] Tenreyro, S., & Thwaites, G. (2013). Pushing on a String: U.S. Monetary Policy is Less Powerful in Recessions.

[185] Tracy, J. S., & Wright, J. (2012). Payment Changes and Default Risk: The Impact of Refinancing on Expected Credit Losses. Federal Reserve Bank of New York Staff Report (No. 562).

[186] Vickery, J., & Wright, J. (2013). TBA Trading and Liquidity in the Agency MBS Market. *Federal Reserve Bank of New York Economic Policy Review*, 19(1), 1-18.

[187] Winter, R. A. (1991). Solvency Regulation and the Property Liability "Insurance Cycle". *Economic Inquiry*, 29(3), 458-471.

[188] Woodward, S. E., & Hall, R. E. (2012). Diagnosing Consumer Confusion and Sub-optimal Shopping Effort: Theory and Mortgage-Market Evidence. *American Economic Review*, 102(7), 3249-76.

[189] Zandi, M., & Deritis, C. (2014). Housing Finance Reform Steps Forward. Moody's Analytics Economic & Consumer Credit Analytics.

[190] Zeldes, S. P. (1989). Consumption and Liquidity Constraints: An Empirical Investigation. *The Journal of Political Economy'*, 97(2), 305-346.

[191] Zhou, X., & Carroll, C. D. (2012). Dynamics of Wealth and Consumption: New and Improved Measures for U.S. States. *The BE Journal of Macroeconomics*, 12(2).

[192] Zhu, J. (2012). Refinance and Mortgage Default: An Empirical Analysis of the HARP's Impact on Default Rates.

Appendix

Appendix A: Stylized Model Proofs

This section shows that the model from Section 1.4 has a unique equilibrium and proves the two propositions outlined in that section. Note that in what follows, all time- and agent subscripts are suppressed for clarity. These proofs rely on five key assumptions:

Assumption 1: Distributions

$(\eta, V^C) \sim F(\eta, V^C)$ and $k \sim G(k)$, and both F and G are exchangeable

Assumption 2: Concavity

(a) $\frac{\partial^2 EP(c,r,p,k)}{\partial r^2} \leq 0$

(b) $\frac{\partial^2 EV(c,p,\eta,V^C,\tau_L)}{\partial c^2} \leq 0$

Assumption 3: Monotonic Choice

(a) $\frac{\partial^2 EP(c,r,p,k)}{\partial r \partial k} > 0$

(b) $\frac{\partial^2 EV(c,p,\eta,V^C,\tau_L)}{\partial c \partial \eta} > 0$

Assumption 4: Lender Preferences

(a) $\frac{\partial \pi(c,r,p,k)}{\partial r} \geq 0 \geq \frac{\partial \pi(c,r,p,k)}{\partial p}$

(b) $\frac{\partial^2 \pi(c,r,p,k)}{\partial r \partial p} \geq 0$ and $|\frac{\partial \pi(c,r,p,k)}{\partial r}| \geq |\frac{\partial \pi(c,r,p,k)}{\partial p}|$

(c) $\pi(c,r,p,k) f_{V^C}(V^R(c,r,\eta)|c,\eta) \frac{\partial V^R(c,r,\eta)}{\partial c} \geq$
$| \int_0^{V^R(c,r,\eta)} \frac{\partial \pi(c,r,p,k)}{\partial c} f_{V^C}(V^R(c,r,\eta)|c,\eta) \, \partial V^C|$

Assumption 5: Borrower Preferences

(a) $\frac{\partial V^R(c,r,\eta)}{\partial c} \geq 0$

(b) $\frac{\partial V^R(c,r,\eta)}{\partial r} < 0$

(c) $\frac{\partial V^R(c,r,\eta)}{\partial \eta} \geq 0$

where expected value and expected profit are given by:

$$\begin{aligned} EP(c,r,p,k) &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} (1 - R(c,r,\eta,V^C))\pi(c,r,p,k)\, f(\eta,V^C|c)d\eta dV^C \\ EV(c,p,\eta,V^C,\tau_L) &= \int_{-\infty}^{\infty} [D(c,p,k,\tau_L) + (1 - D(c,p,k,\tau_L))R(c,r,\eta,V^C)]V^C \\ &+ \ (1 - D(c,p,k,\tau_L))(1 - R(c,r,\eta,V^C))V^R(c,r,\eta)\, g(k)dk \end{aligned}$$

The first assumption is purely technical. Crucially, the distribution F must still be exchangeable when conditioned on the observed c^*, although members of the exponential family satisfy this requirement. Assumptions 2 and 3 are also technical, and are required to ensure, respectively, that the LTV and interest rate choice problems are well-behaved and that the optimal choice is increasing in the borrowers or lender's type. This amounts to assuming that the solution r^* to the lender's interest-rate-choice problem satisfies certain requirements that lead to concavity and positive cross-partials in c and η for expected value. These assumptions also allow me to ignore issues of incentive compatibility. The fourth assumption places some minimal structure on lender payoffs; 4a and 4b state that interest rates are weakly more profitable than guarantee fees are costly, while 4c states that without a compensating increase in interest rates, the lender's benefit from lower rejection rates at higher LTV ratios does not outweigh the cost of lower profit. The former must be true, as profit on a securitized loan depends on $r - p$ while profit on an un-securitized loan depends only r; hence the cross-partial is zero for a securitized loan and positive for an un-securitized loan. The latter is a reasonable assumption, and if false, would imply that lenders prefer high-LTV applicants, all else equal. The final assumption places similarly reasonable structure on borrower payoffs.

Proof of Unique Equilibrium:

1. Let $h(c,r,\eta,V^C) = V^R(c,r,\eta) - V^C$. Because h is monotonically decreasing in r by Assumption 5b, I can write the condition $V^C \geq V^r(c,r,\eta)$ as equivalently $h(c,r,\eta,V^C) \leq 0$ or $r \geq h_r^{-1}(c,\eta,V^c)$. Hence there is a unique interest rate $h_r^{-1}(c,\eta,V^c)$ above which offers are rejected, so for each type (η,V^c) the borrower's rejection policy R is uniquely determined.
2. Because by Assumption 2a the expected profit function is concave, there exists a unique $r^*(c,p,k)$ that maximizes $EP(c,r,p,k)$. By Assumption 3a, $r^*(c,p,k)$ is increasing monotonically in k. Since $r^*(c,p,k)$ is unique for each type k, then $EP(c,r^*(c,p,k),p,k)$ is also uniquely determined, hence the lender's denial policy D is also unique for each type k.
3. Because by Assumption 2b the expected value function is concave, there exists a unique $c^*(p,\eta,V^C,\tau_L)$ that maximizes $EV(c,p,\eta,V^C,\tau_L)$. By assumption 3b, $c^*(p,\eta,V^C,\tau_L)$ is increasing monotonically in η. Since $c^*(p,\eta,V^C,\tau_L)$ is unique for each type (η,V^C), then $EV(c^*(p,\eta,V^C,\tau_L),p,\eta,V^C,\tau_L)$ is also uniquely determined, hence the borrower's application policy A is also unique for each type (η,V^C).
4. Since the policies R, D, A, r^* and c^* are all unique for each set of borrower and lender types (k,η,V^C), the game features a unique equilibrium.

Proof of Proposition 1:

r^* maximizes $EP(c,r,p,k)$, so $\frac{dEP(c,r^*,p,k)}{dr^*} = \frac{\partial EP(c,r^*,p,k)}{\partial r^*} = 0 \geq \frac{\partial^2 EP(c,r^*,p,k)}{\partial r^{*2}} = \frac{d^2 EP(c,r^*,p,k)}{dr^{*2}}$. Since $\frac{\partial EP(c,r^*,p,k)}{\partial r^*}$ is identically zero at r^*, $\frac{d}{dp}\frac{\partial EP(c,r^*,p,k)}{\partial r^*} = \frac{\partial^2 EP(c,r^*,p,k)}{\partial r^* \partial p} + \frac{\partial^2 EP(c,r^*,p,k)}{\partial r^{*2}}\frac{\partial r^*(c,p,k)}{\partial p} = 0$. Then since $\frac{dr^*}{dp} = \frac{\partial r^*}{\partial p} = -\frac{\frac{\partial^2 EP(c,r^*,p,k)}{\partial r^* \partial p}}{\frac{\partial^2 EP(c,r^*,p,k)}{\partial r^{*2}}}$, therefore $Sgn(\frac{dr^*}{dp}) = Sgn(\frac{\partial^2 EP(c,r^*,p,k)}{\partial r^* \partial p})$. Note that I can write expected profit at r^* by conditioning on the event that the offer is accepted:

$$EP(c,r^*,p,k) = \int_{-\infty}^{\infty}[\int_{-\infty}^{V^R(c,r^*,\eta)} \pi(c,r^*,p,k) f_{V^C}(V^C|c,\eta)dV^C] f_\eta(\eta|c)d\eta$$

then using the exchangeability from Assumption 1,

$$\begin{aligned}
\frac{\partial^2 EP(c,r^*,p,k)}{\partial r^* \partial p} &= \frac{\partial}{\partial p}\int_{-\infty}^{\infty}[\pi(c,r^*,p,k)f_{V^C}(V^R(c,r^*,\eta)|c,\eta) \\
&+ \int_{-\infty}^{V^R(c,r^*,\eta)} \frac{\partial \pi(c,r^*,p,k)}{\partial r^*} f_{V^C}(V^C|c,\eta)dV^C]f_\eta(\eta|c)d\eta \\
&= \int_{-\infty}^{\infty}[\underbrace{\frac{\partial \pi(c,r^*,p,k)}{\partial p}}_{\leq 0 \text{ by Asm. 4a}} \underbrace{f_{V^C}(V^R(c,r^*,\eta)|c,\eta)}_{\geq 0, \text{ Density}} \underbrace{\frac{\partial V^R(c,r^*,\eta)}{\partial r^*}}_{\leq 0 \text{ by Asm. 5b}} \\
&+ \int_{-\infty}^{V^R(c,r^*,\eta)} \underbrace{\frac{\partial^2 \pi(c,r^*,p,k)}{\partial r^* \partial p}}_{\geq 0 \text{ by Asm. 4b}} \underbrace{f_{V^C}(V^C|c,\eta)}_{\geq 0, \text{ Density}} dV^C] \underbrace{f_\eta(\eta|c)}_{\geq 0, \text{ Density}} d\eta
\end{aligned}$$

because $V^R(c,r,\eta)$ is not a function of p, hence $\frac{\partial V^R(c,r^*,\eta)}{\partial p} = 0$. The entire integrand of the inner integral term is positive, so the integral itself is positive, and therefore the integrand of the outer integral term is also positive, so the integral itself is positive. Thus $\frac{dr^*}{dp} \geq 0$, which completes the first part of the proof. For the second part, notice that since $D(c,p,k,t) = \mathbf{1}[EP(c,r^*(c,p,k),p,k) \leq \tau_L]$ is a step function, it is sufficient to show that $\frac{dEP(c,r^*(c,p,k),p,k)}{dp} \leq 0 \; \forall k$. As before I can write:

$$\begin{aligned}
\frac{dEP(c,r^*(c,p,k),p,k)}{dp} &= \frac{\partial EP(c,r^*,p,k)}{\partial p} + \underbrace{\frac{\partial EP(c,r^*,p,k)}{\partial r^*}}_{=0 \text{ at Optimum}} \frac{\partial r^*}{\partial p} \\
&= \int_{-\infty}^{\infty}\int_{-\infty}^{V^R(c,r^*,\eta)} \underbrace{\frac{\partial \pi(c,r^*,p,k)}{\partial p}}_{\leq 0 \text{ by Asm. 4a}} \underbrace{f(\eta,V^C|c)}_{\geq 0, \text{ Density}} dV^C d\eta
\end{aligned}$$

because the integral over f is exchangeable by Assumption 1 and the dependence of bound $V^R(c,r^*,\eta)$ on p comes only through r^* and is thus absent from the partial derivative. Therefore the entire integrand is negative, so the integral itself is negative, so $\frac{d\int_{-\infty}^{\infty} D(c^*(p_t,\eta_{i,t},V^C_{i,t},\tau_L),p_t,k,\tau_L)g(k)dk}{dp_t} \geq 0$.

Proof of Proposition 2:

I begin with the latter of the two claims. As above, since $D(c,p,k,t)$ is a step function, it is sufficient to show that $\frac{dEP(c^*(p,\eta,V^C,\tau_L),r^*(c^*(p,\eta,V^C,\tau_L),p,k),p,k)}{d\eta} \leq 0$. Using the fact that $r^*(c,p,k)$ maximizes $EP(c,r,p,k)$, I can write:

$$
\begin{aligned}
\frac{dEP(c^*(p,\eta,V^C,\tau_L),r^*(c^*(p,\eta,V^C,\tau_L),p,k),p,k)}{d\eta} &= (\frac{\partial EP(c^*(p,\eta,V^C,\tau_L),r^*,p,k)}{\partial c^*(p,\eta,V^C,\tau_L)} \\
+ \underbrace{\frac{\partial EP(c^*(p,\eta,V^C,\tau_L),r^*(c^*(p,\eta,V^C,\tau_L),p,k),p,k)}{\partial r^*(c^*(p,\eta,V^C,\tau_L),p,k)}}_{=0\text{ at Optimum}} & \frac{\partial r^*(c^*(p,\eta,V^C,\tau_L),p,k)}{\partial c^*(p,\eta,V^C,\tau_L)}) \underbrace{\frac{\partial c^*(p,\eta,V^C,\tau_L)}{\partial \eta}}_{>0\text{ by Asm. 3b}}
\end{aligned}
$$

where as discussed previously the positive cross-partial of $EV(c,p,\eta,V^C,\tau_L)$ with respect to c and η guarantees that $\frac{\partial c^*(p,\eta,V^C,\tau_L)}{\partial \eta} > 0$. Therefore since:

$$
\begin{aligned}
\frac{\partial EP(c^*,r^*,p,k)}{\partial c^*} &= \int_{-\infty}^{\infty} [\frac{\partial \pi(c^*,r^*,p,k)}{\partial c^*} f_{V^C}(V^R(c^*,r^*,\eta)|c,\eta) \frac{\partial V^R(c,r^*,\eta)}{\partial c^*} \\
&+ \int_{-\infty}^{V^R(c^*,r^*,\eta)} \frac{\partial \pi(c^*,r^*,p,k)}{\partial c^*} f_{V^C}(V^C|c,\eta)dV^C] \underbrace{f_\eta(\eta|c)}_{\geq 0,\text{ Density}} d\eta
\end{aligned}
$$

where the entire inner integrand (bracketed) is weakly negative by Assumption 4c, then $\frac{\partial EP(c^*,r^*,p,k)}{\partial c^*}$ is itself weakly negative, so $\frac{dEP(c^*(p,\eta,V^C,\tau_L),r^*(c^*(p,\eta,V^C,\tau_L),p,k),p,k)}{d\eta} \leq 0$, which completes the proof. For the former of the two claims, note that:

$$
\begin{aligned}
\frac{dEV(c^*(p,\eta,V^C,\tau_L),p,\eta,V^C,\tau_L)}{d\eta} &= \underbrace{\frac{\partial EV(c^*(p,\eta,V^C,\tau_L),p,\eta,V^C,\tau_L)}{\partial c^*(p,\eta,V^C,\tau_L)}}_{=0\text{ at Optimum}} \frac{\partial c^*(p,\eta,V^C,\tau_L)}{\partial \eta} \\
&+ \frac{\partial EV(c^*,p,\eta,V^C,\tau_L)}{\partial \eta}
\end{aligned}
$$

Thus I need only consider the partial derivative. Recall that under the maintained assumptions, I can express the borrower's interest rate offer rejection decision as a cutoff rule in the interest rate $r \geq h_r^{-1}(c,\eta,V^C)$. I can analogously express denial decision $EP(c,r^*(c,p,k),p,k) \leq \tau_L$ as $M(c,p,k,\tau_L) \leq 0$, where $M(c,p,k,\tau_L) = EP(c,r^*(c,p,k),p,k) - \tau_L$. Having shown previously in the proof of Proposition 1 that $\frac{dEP(c,r^*(c,p,k),p,k)}{dp} \leq 0$, therefore M is monotonically decreasing in k. As a consequence I can invert M in k and express the lender's denial decision as $k \geq M_k^{-1}(c, p, \tau_L)$. I can also combine the borrower's rejection decision rule with the lender's optimal interest rate offer to express the rejection decision as $r^*(c, p, k) \geq h_r^{-1}(V^C, c, \eta)$. I can express this rule as $N(c, p, k, \eta, V^C) \geq 0$, where $N(c, p, k, \eta, V^C) = r^*(c, p, k) - h_r^{-1}(V^C, c, \eta)$. Since as discussed previously Assumption 3 guarantees that $r^*(c, p, k)$ is monotonically increasing in k, therefore N is monotonically increasing in k as well. I can thus invert N in k and write the rejection decision as $k \geq N_k^{-1}(c, p, \eta, V^C)$. We can see from $EV(c, p, \eta, V^C, \tau_L)$ that there are effectively three regions:

Region 1: At low levels of k, the lender offers an interest rate that the borrower accepts

Region 2: At moderate levels of k, the lender offers an interest rate that the borrower rejects

Region 3: At high levels of k, the lender denies the borrower's application

where it is possible, for certain high values of η or low values of V^C, that Region 2 has zero mass (i.e. the borrower accepts all offers the lender is willing to make). Therefore the borrower's application results in a completed refinancing if $k < Q$, where $Q = \min\{M_k^{-1}(c, p, \tau_L), N_k^{-1}(c, p, \eta, V^C)\}$. Hence I can rewrite the borrower's expected value as:

$$EV(c, p, \eta, V^C, \tau_L) = \int_{-\infty}^{Q} V^R(c, r^*(c, p, k), \eta) g(k) dk + \int_{Q}^{\infty} V^C g(k) dk$$

Thus:

$$\begin{aligned} \frac{\partial EV(c^*, p, \eta, V^C, \tau_L)}{\partial \eta} &= \int_{-\infty}^{Q} \underbrace{\frac{\partial V^R(c, r^*(c, p, k), \eta)}{\partial \eta}}_{\geq 0 \text{ by Asm. 5c}} \underbrace{g(k)}_{\geq 0, \text{ Density}} dk \\ &+ (V^R(c, r^*(c, p, k), \eta) - V^C) \underbrace{g(Q)}_{\geq 0, \text{ Density}} \frac{\partial Q}{\partial \eta} \end{aligned}$$

While the first term must be positive because the integrand itself is positive, for the second term there are two cases to consider. In the first case, η and V^C (and hence c^*) are such that the rejection region has zero mass, or $N_k^{-1}(c, p, \eta, V^C) \geq M_k^{-1}(c, p, \tau_L)$. In this case, $\frac{\partial Q}{\partial \eta} = 0$ and $V^R(c, r^*(c, p, k), \eta) \geq V^C$, because the borrower prefers refinancing at the lender's equilibrium interest rate offer. As a result, the second term is weakly positive, so the whole expression is weakly positive. In the second case, the rejection region has positive mass, or $N_k^{-1}(c, p, \eta, V^C) < M_k^{-1}(c, p, \tau_L)$. In this case, $Q = N_k^{-1}(c, p, \eta, V^C)$, and thus at $k = Q$ we have $V^R(c, r^*(c, p, k), \eta) = V^C$, since the borrower is just indifferent at this point. In this case, the second term is zero, so the whole expression is weakly positive. Thus we have $\frac{\partial EV(c^*, p, \eta, V^C, \tau_L)}{\partial \eta} \geq 0$, which completes the proof.

Appendix B: Details on Dataset Assembly

Primary Market Datasets

The first step in the primary-market analysis is to generate the GSE dataset. Both the raw FHLMC Single-Family Loan-Level Dataset and FNMA Single-Family Loan Performance Dataset contain a header file, including characteristics of the loan at origination, and a separate dynamic file, including the choices of the borrower each month. For all analysis with interest rates as the independent variable in Chapters 2 and 4, I use only the header dataset, as this includes credit risk characteristics and interest rates at origination for 21,435,660 refinancing loans issued between 2000 and early 2013. For all analysis with monthly choices (either refinance or default) as the independent variable in Chapter 2, I use a subset of the full merged header-dynamic file GSE dataset containing only loans originated after June 2009, as these loans were ineligible for HARP. This sample consists of 6,973,525 loans observed whenever active between July 2009 and June 2013, totaling 150,759,569 loan-month-level observations in all. Both of these datasets are nationwide. For most of the analysis in Chapter 3, I use a subset of the full merged header-dynamic GSE dataset containing only loans originated within 12 months of the deadline for HARP eligibility (June 2009). This entire sample consists of 4,664,219 distinct fixed-rate mortgage loans observed whenever active between June 2008 and December 2012, totaling 177,174,761 loan-month-level observations in all. This dataset is nationwide and covers a substantial fraction of all mortgage loans originated over the period in question. I supplement these datasets by importing a number of aggregate and sub-aggregate variables and using them to impute certain data elements. I obtain house price indices from Zillow and applied them at the MSA-month- or state-month-level, whichever is the finest level of detail available. The house price variable itself is a combination of the observed price at origination, in turn imputed from the balance and LTV at origination, and the percentage change in the local-area house price index. Servicing decisions, monthly payments, home equity, and in some cases ending balances are imputed from the GSE data via the appropriate formula, in the case of home equity using the imputed house price.[39] I also import MSA-month or state-month-level unemployment rates from the Census Bureau, again using the finest-available level of geographic detail, month-level average mortgage interest rates provided by Freddie Mac, month-level average interest rates for a variety of instruments from Federal Reserve Table H.15, and secondary market MBS yields and the Merrill Lynch Options Volatility Index (MOVE), averaged up to the month-level, from Bloomberg.

I construct instruments for house price indices using the WRLURI index, described in detail in Gyourko et al (2008).[106] I regress the house price index on time- and geography-level dummies as well as the constituent variables making up the index, which measure various supply restrictions (e.g. whether local approval is needed for building decisions, or whether homebuilders have to pay to put in their own sewer lines) to generate a predicted variable. When importing instruments at the MSA- or

[39] Ending balances are censored in the data for the first six months the loan is active. For FRMs, monthly payments are a fixed nonlinear function of the initial balance, interest rate, and term, and can thus be reconstructed manually.

state-level as in Chapters 2 and 3, I use MSA- or state-level averages over all towns in the WRLURI dataset constituting that MSA or state, and when importing instruments at the town-level in Chapter 3 I use the town-level indices.I generate instruments for unemployment rates in a similar fashion following the procedure from Bartik (1991).[25] These unemployment instruments measure the change in employment that would be predicted if the industry mix of employment remained fixed and each industry grew at its corresponding national average rate. I use state-level employment figures from the 2000 census for 12 North American Industry Classification System categories to generate base-year employment shares and apply national month-level employment growth figures from the Bureau of Labor Statistics. This generates a predicted employment figure, which I then include in a first-stage regression for unemployment rates at the finest observed level of geographic detail (either town, MSA, or state) in order to form a prediction for unemployment rates. Finally, for the jumbo-conforming spread regressions shown in Figure 6, I use daily average jumbo and conforming refinance mortgage rates from Bloomberg, ultimately sourced fro Bankrate.com.

Secondary Market Datasets

To assemble the short-panel secondary market dataset, I first download all FHLMC's Historical Daily Loan Level Fixed-Rate Disclosures for every trading day in a two-week window of the April 1st, 2012 and December 1st, 2012 guarantee fee increase implementation dates. FNMA's comparable tool, PoolTalk, does not list MBS by the trading day they were issued but rather by the month; hence I download all FNMA single-family FRM-backed MBS issued in either March, April, November or December of 2012. These disclosures contain all of the MBS coupon and risk characteristics used in the analysis. I merge them with daily mid prices from Bloomberg by merging on the CUSIP. The prices reported by Bloomberg are the "default" option, which blends transaction prices, quoted prices, and model estimates in order to provide the best possible source for each query. Wherever possible, the price is the CBBT price, a weighted average of 2^{nd}- and 3^{rd}-best bid and ask prices from actual completed transactions (the mid price is the midpoint of the bid and ask). If there are insufficient transactions, the price is the BGN price, which is a similar weighted average of quotes from broker-dealers that do not correspond to transactions. If there are insufficient quotes, Bloomberg supplies the BVAL price, an estimated price from a proprietary option-adjusted spread model, the standard pricing model in the industry. For the purposes of determining the exact issuance date for FNMA securities, I assume that the first reported date with a trading price is the issuance date. The resulting pattern of issuance dates for FNMA resembled the pattern for FHLMC MBS, for which I observe the issuance date, suggesting that this assumption is valid. This yields a dataset of 59,438 MBS-trading day observations over the two 4-week windows. For the longer-panel secondary market dataset, I first download all monthly FNMA single-family FRM-backed MBS files for the months July 2011 to June 2013. I merge them with weekly mid prices from Bloomberg by merging on the CUSIP of all MBS for which prices are available.

I also construct a mortgage-backed securities dataset for the purposes of estimating two primitives as inputs to the structural model in Chapter 4, specifically, an estimate of the elasticity of the guarantee fee with respect to loan-to-value ratio and a pricing model for mortgages sold on the secondary market. I downloaded data from Bloomberg on all FNMA and FHLMC 30-year FRM-backed MBS with at least a $1 million origination volume issued between February 2004 and March 2013. This dataset consists of 40,307 MBS with their coupon rate, weighted-average coupon, weighted-average credit score, weighted-average LTV, origination balance, principal geography, and price, at origination. I measure the guarantee fee as the difference between the weighted-average coupon and the coupon rate. I estimate a regression model to obtain the elasticity of the guarantee fee with respect to LTV, which I estimate as .5992 bp for each 1% increase in LTV. I also estimate a pricing model for MBS via linear regression on the same data; this pricing model is shown in Table A40 of Appendix I.

HMDA Records

In order to assemble the HMDA dataset, I first remove observations missing key variables or not corresponding to agency refinancing. I remove loans applications from outside of California, FHA- or VA-backed loans, second-lien mortgages, home purchase or home improvement loans, loans above the conforming loan limit, or loans originated by subprime lenders. Subprime lenders are identified using the Deparment of Housing and Urban Development's subprime lender list; this list is only available through 2006, so for years 2007-2012 I identify lenders as subprime if their fraction of high-interest loans originated is greater than the median fraction (roughly 27%) for known subprime originators. This sample construction is designed to make the HMDA dataset resemble the agency refinancing market, in order to infer what credit conditions were like for that market specifically. I next impute the values of certain borrower credit risk characteristics, in particular LTV, FICO and DTI, using a technique similar to that in Bayer et al (2013).[28] Using the set of all refinances observed in the GSE dataset, I calculate the mean of these variables within a year, lender, Sectional Center Facility (SCF) code, decile of income and decile of origination balance. I then apply that mean to observations in the HMDA dataset.[40] I then randomly assign observations to a month by year and import aggregate data at the county-month level on mortgage interest rates and unemployment rates, and at the month level on average charged guarantee fees and various base interest rates (including 1-, 3-, 5-, 7-, 10- and 30-year treasury rates and 10- and 30-year swap rates). Finally, to complete the dataset assembly process I generate indicators for one of four categories of action taken on the loan; either denied application, rejected credit offer, originated loan held on balance sheet, or originated

[40] There is considerable variation in these mean values; across the 127,113 bins, the mean FICO varies between 364 and 832, while the mean LTV and DTI vary between 1 and 100 and 65, respectively.

loan sold to GSE. I include in the “denial” category applications where a pre-approval request was denied as well as those where the file was closed for incompleteness (i.e. the borrower did not provide all the requested information). I include as “offer rejected” cases where the borrower withdrew their application prior to receiving an offer or rejected a preapproval offer. The “held” category contains all other dispositions of originated loans, including over-the-counter sales to a commercial bank or to a nonbank financial company, private-label securitization, and securitization via Ginnie Mae and Farmer Mac as well as actually retained portfolio loans. This categorization tends to under-represent GSE securitization, as the loan is labeled as “held” if the loan is not transferred to a GSE by the end of the reporting year. This process yields a dataset containing 17,109,796 over 13 years from 2000 to 2012.

Matched GSE-Deeds Sample

The first step in assembling the matched dataset is to generate the deeds dataset. Starting with the universe of all loans secured by residential real estate originated between 2000 and 2012, I first restrict the sample to only loans made in California and remove all home-equity loans, lines of credit, and second-lien loans (leaving only first-lien mortgages). I also remove all observations for any properties with observations that are missing key data elements, or that have a very high number of observed transactions.[41] I then generate a sequence of transactions on the same property. I classify a loan as a refinance if the borrower’s last name is the same as the previous loan for that property and the previous loan was made more than two months prior.[42] All other loans I classify as sales. I then reformat each series of loans on the same property into matched pairs, with one origination loan and one termination. I classify terminations as either refinance, sale, default (in cases where the name of the borrower is a bank or other

[41] These are a relatively small fraction of the sample and appear to be data entry errors

[42] In cases where the previous loan was made 2 or fewer months prior, I classify the loan as a second mortgage and exclude it from the sample.

financial company) or none (continuation). I then remove all pairs in which the termination is via sale or the origination would not appear in the GSE dataset; specifically, I drop all ARM originations or originations with a balance over the conforming loan limit.[43] These steps are taken so that the sample of loan originations is the same in both the deeds and GSE dataset.[44] I then merge these pairs of observations with assessor records at the property-level. This generates a sample of 1,484,481 origination-termination pairs, which I refer to henceforth as the deeds dataset.

The second step is to assemble the GSE dataset. Each record in the raw dataset contains a header file, including characteristics of the loan at origination, and a separate dynamic file, including the choices of the borrower each month. Using the dynamic file, I classify each loan by termination reason; either a prepayment (including sales and refinances), default, or if still active at the end of 2012, continuation. I merge these termination records onto the header file in order to generate matched origination-termination pairs. I then remove all observations from outside of California or missing key variables. This leaves a sample of 5,003,826 origination-termination pairs, which I refer to in what follows as the GSE dataset.

The third step is to merge these two dataset. I first structure them in a common format by changing certain codes. This includes recasting the ZIP Code in the deeds dataset as a SCF Code, a 3-digit version of a ZIP Code, placing lender names in a common format, and recoding small lenders not listed in the GSE dataset as "Other" in the deeds dataset. I then perform a many-to-many merge, classifying each origination-termination pair in both datasets into a bin depending on the origination loan purpose (purchase or refinance), occupancy status (owner-occupied or investment property), lender (one of a possible 32, or all other), termination code (continue, default or refinance), and SCF code. Within each bin, for each possible pair of deeds dataset and GSE dataset observations, I generate a distance metric based on the sum of three

[43] I use the conforming loan limit for that year and city, as provided by the FHFA.
[44] I discuss the role of excluding sales shortly.

discrepancies: their origination balances as a percentage of the balance recorded in the deeds record, their origination dates, and their termination dates.[45] The purpose of allowing for some discrepancy in these figures is that differences in record-keeping standards can generate minor differences; for example, closing costs may be recorded as part of the balance in one record but not the other, or origination dates can record the date the contract was signed in one case and the first payment date in another. Then within each bin I implement a deeds-optimal Gale-Shapley algorithm to generate a stable matching, where preferences are symmetric and equal to the negative of the distance in each case (so that all observations "prefer" a closer match). This procedure involves iteratively forming potential matches between "nearby" observation pairs, which are then broken if an unmatched observation is closer; see Gale and Shapley (1962) [88] for more detail on implementation. I then remove all unmatched observations or matched pairs in which the distance is greater than .1 (corresponding to either a 10% difference in loan balances or two months total discrepancy in timing). This generates a set of 481,984 matched pairs containing the necessary information: risk characteristics and monthly decisions from the GSE dataset, and property characteristics and the refinancing balance from the deeds dataset.

In order to enhance the external validity of this study and enable the results to be reasonably applied to the universe of GSE FRMs, I resample the set of matched paired observations to make it resemble aggregate figures. This procedure avoids the over-sampling of "unique" circumstances, such as loans made in rural areas, that would otherwise result from matching these two datasets. I generate sampling probabilities based on characteristics at both origination and refinance. Using the master GSE dataset, I generate observed probabilities of bins defined by an origination year, SCF, termination code, termination year, loan purpose, and occupancy status. I appy these probabilities to observations in the matched dataset corresponding to the same bin. Using the FHLMC

[45] Each 1-month difference in origination or termination date was specified to be equivalent in distance to a 8.33% difference in origination balance.

Refinance Report, I also generate sampling probabilities for observations terminating in refinance. Each quarter, I estimate via Method of Moments the mean and variance of a normal distribution of cash-out as a percentage of refinancing balance. I target three average quarterly moments; average cash-out as a percentage of refinancing balance, percentage reducing balance, and percentage increasing balance.[46] Using this distribution, I generate a likelihood for the observed cash-out percentage of each refinance observed in the matched dataset, and adjust the sampling probability accordingly.[47] I then draw 200,000 observations, with replacement, according to these sampling probabilities. As shown in Figure 1.5 in Section 1.5.4 and Figures A2, A3, A4, and A5 in Appendix C, the matched dataset corresponds fairly well with aggregate figures.

To finish constructing the matched dataset, in the final step I import a number of aggregate and sub-aggregate variables and impute certain data elements. Certain variable definitions warrant discussion in greater detail. House price indices are obtained from Zillow and applied at the place-level[48], as are local-area unemployment rates from the Census Bureau.[49] House prices themselves are a combination of the observed price at origination (itself imputed from the origination balance and LTV) and the percentage change in the local-area house price index. Servicing decisions, monthly payments, home equity, and in some cases ending balances are imputed from the GSE data via the appropriate formula.[50] Other aggregate variables, such as mortgage rates

[46] In the FHLMC definition, "reducing balance" means cash-out less than zero, while "increasing balance" means cash-out greater than 5% of the refinancing balance, so the sum of the two is not 100%.

[47] I also adjust the sampling probabilities of non-refinanced observations by a corresponding factor in order to preserve their relative weight.

[48] Zillow provides home price indices at the ZIP Code-level, which I observe, but the records are less complete than those at the city level.

[49] Places (either an incorporated place or a census-designated place (CDP)) are roughly the size of towns, and most large cities contain many. For example, I can separately observe homeowners from Beverly Hills, Compton, Culver City, Hollywood, Huntington Park, Inglewood, Montebello, Monterey Park, Santa Monica, and a variety of other places colloquially considered to be part of Los Angeles. There are over 900 places represented in the matched dataset, and they represent smaller populations that are much more demographically homogeneous than either cities or counties.

[50] Ending balances are censored in the data for the first six months the loan is active. For FRMs, monthly payments are a fixed nonlinear function of the initial balance, interest rate, and term.

and mean unemployment duration, are imported at slightly higher levels of aggregation.

I construct several aggregate-level datasets for the purposes of imputing certain borrower-level covariates and for defining borrower expectations. I have data on house price indices from Zillow, defined at the city, county and MSA level for all of California. I also have analogous data on unemployment rates at the same levels from the BLS. I import this data at the finest-available level of geographic detail. If a borrower's observed city does not match to a city with aggregate data, I attempt to match at the county-month, and so forth. Similarly if data is not available at the city-level for a particular month in which I have borrower-level observations, I use county-level data. Certain individual-level covariates, such as home equity, are defined as a function of the aggregate home price index and borrower-level data on purchase price at origination and ending balance. I also have mortgage rate data at the MSA level from HSH Associates and national-level interest rates from the Federal Reserve, and import these data in a similar fashion. For counties outside of MSAs, I use the average for the state of California from HSH Associates. Finally, I construct measures of mean unemployment duration by county using data from the Current Population Survey. I take averages within a county-month of continuous weeks unemployed for currently unemployed persons, weighted by CPS-defined statistical sampling weights. I then smooth the resulting county-month-level panel series, placing additional weight on any observations for March, the month for which an additional economic supplementary survey is conducted by the CPS and for which more observations are available. A handful of rural counties in California have insufficient data to estimate a reasonable county-level average[51], and for these counties I use the average for the state of California as a whole. These ten variables (House Price Indices, Mean Unemployment Duration, Unemployment Rate, Mortgage Rates, 30-Year Rates, 10-Year Rates, 7-Year Rates, 5-Year Rates, 3-Year Rates and 1-Year Rates) defined at the county-level form the panel dataset

[51]Specifically, Amador, Butte, Del Norte, Kings, Lassen, Madera, Mariposa, Sutter, Tehama, and Yuba counties.

for the vector autoregression models I estimate to specify borrower preferences.

APPENDIX C: ADDITIONAL DATASET SUMMARY STATISTICS

Table A1: GSE Dataset Summary Statistics
Panel A: Complete Origination Dataset

Variable	Mean	Std. Dev.	Min.	Max.
Orig. Balance	187,578	103,301	4,000	1,470,000
Int. Rate	5.91	1.11	1.875	13.5
FICO	730.71	55.51	300	850
LTV	72.35	15.99	1	105
DTI	35.00	12.61	0	65
Rate-Term Refi.	0.24	0.43	0	1
Cash-Out Refi.	0.19	0.39	0	1
Investor	0.10	0.30	0	1
TPO	0.56	0.50	0	1
N				35,701,979

Panel B: Refinance Originations

Variable	Mean	Std. Dev.	Min.	Max.
Orig. Balance	193381.8	105030.4	4,000	1,403,000
Int. Rate	5.79	1.04	1.875	12.63
FICO	728.60	56.61	300	850
LTV	67.96	15.74	1	105
DTI	34.35	12.76	0	65
Rate-Term Refi.	0.40	0.49	0	1
Cash-Out Refi.	0.32	0.47	0	1
Investor	0.08	0.27	0	1
TPO	0.56	0.50	0	1
N				21,435,660

Panel C: HARP-Ineligible Refinance Originations

Variable	Mean	Std. Dev.	Min.	Max.
Orig. Balance	240749.8	134517.7	8,000	1,470,000
Int. Rate	4.45	0.63	2.5	8.625
FICO	762.04	39.73	439	850
LTV	70.70	15.94	1	105
DTI	32.69	10.18	0	65
Rate-Term Refi.	0.26	0.44	0	1
Cash-Out Refi.	0.14	0.34	0	1
Investor	0.08	0.27	0	1
TPO	0.55	0.50	0	1
N				6,973,525

Notes: Average characteristics of mortgages at origination in GSE dataset. Panel A is all mortgages from Q1 2000-Q2 2013, Panel B is all refinance mortgages from Q1 2000-Q2 2013, and Panel C is all mortgages from Q3 2009-Q2 2013.

Figure A1: Originations by Month in GSE Dataset

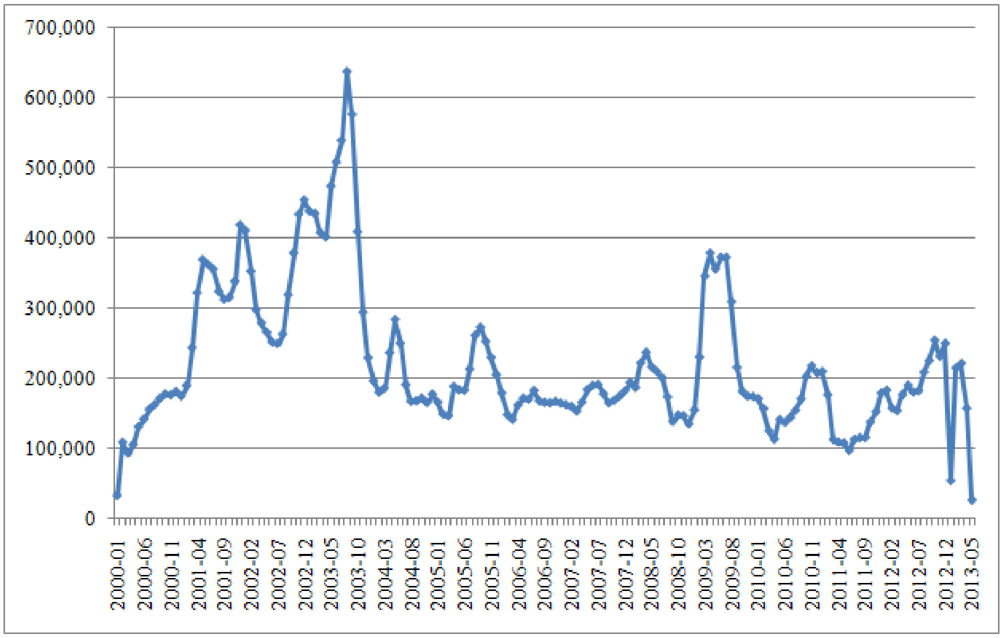

Notes: Monthly Origination Volume in GSE dataset.

Table A2: Count of Mortgage Originations by Month in HARP Dataset

Pre-HARP		Post-HARP	
Month	**# Originations**	**Month**	**# Originations**
2008-06	136,928	2009-06	296,716
2008-07	105,094	2009-07	207,276
2008-08	114,085	2009-08	174,923
2008-09	119,994	2009-09	168,540
2008-10	113,360	2009-10	167,701
2008-11	135,936	2009-11	165,077
2008-12	214,947	2009-12	151,407
2009-01	329,514	2010-01	120,959
2009-02	362,719	2010-02	110,024
2009-03	342,493	2010-03	137,321
2009-04	358,264	2010-04	132,821
2009-05	358,234	2010-05	139,886
Total	2,691,568	Total	1,972,651

Notes: Number of unique mortgages originated by month in dataset.

Table A3: Secondary Market Dataset Summary Statistics

Panel A: April 2012 Fee Increase Implementation

Variable	Mean	Std. Dev.	Min.	Max.
Pool Size	2.54E+07	1.09E+08	25,564	3.25E+09
Coupon	3.74	0.51	2.5	7.0
# Loans	119.77	418.02	1	10570
Avg. Loan Size	203,900	89,400	25,600	785,000
Orig. WAC	4.22	0.50	2.96	8.37
Orig. WAM	324.6	69.2	107	361
Avg. FICO	751.34	26.17	621.38	808.06
Avg. LTV	75.36	12.27	38.84	104.37
% Refi.	80.34	23.60	0	100
% Occ.	80.84	24.85	0	100
FNMA	0.52	0.50	0	1
G-Fee	0.49	0.11	0.25	1.37
N	-	-	-	2,696

Panel B: December 2012 Fee Increase Implementation

Variable	Mean	Std. Dev.	Min.	Max.
Pool Size	2.09E+07	9.20E+07	31,403	4.08E+09
Coupon	3.30	0.52	2.5	5.5
# Loans	98.67	350.30	1	14559
Avg. Loan Size	198,615	82,101	31,450	757,250
Orig. WAC	3.88	0.52	2.78	6.24
Orig. WAM	329.6	61.6	117	361
Avg. FICO	747.92	26.06	622.00	816.70
Avg. LTV	81.41	20.70	37.14	218.41
% Refi.	80.61	24.27	0	100
% Occ.	80.94	23.16	0	100
FNMA	0.53	0.50	0	1
G-Fee	0.59	0.15	0.25	1.47
N	-	-	-	3,646

Panel C: Full FNMA Issuance Dataset

Variable	Mean	Std. Dev.	Min.	Max.
Pool Size	2.16E+07	8.55E+07	22,734	4.66E+09
Coupon	3.44	0.63	1.5	7.0
# Loans	102.75	350.68	1	16451
Avg. Loan Size	197,175	87,945	22,795	877,098
Orig. WAC	3.97	0.62	2.25	8.37
Orig. WAM	297.9	82.9	38	480
Avg. FICO	750.57	25.69	595	820
Avg. LTV	76.64	18.14	21	226
% Refi.	77.94	26.56	0	100
% Occ.	80.89	24.13	0	100
FNMA	1.00	0.00	1	1
G-Fee	0.53	0.14	0.25	1.68
N	-	-	-	58,106

Notes: Panels A and B show average characteristics of all FHLMC and FNMA MBS issued in 2-week window around April 1st, 2012 and December 1st, 2012, respectively. Panel C shows average characteristics of all FNMA MBS issued in 2-year period from July 2011 to June 2013.

Table A4: Average Loan Applicant Risk Characteristics by Year

Year	Income ($)	FICO	LTV	Interest Rate	Guarantee Fee (bp)
2000	74,434	710	74	8.13	19.3
2001	82,297	712	70	7.15	19.3
2002	84,863	717	66	6.65	19.3
2003	86,920	725	62	5.94	21.0
2004	80,443	716	60	5.79	20.8
2005	82,184	718	55	5.87	21.0
2006	95,755	721	54	6.60	21.8
2007	105,199	725	57	6.61	26.6
2008	115,772	740	66	6.60	28.0
2009	130,066	764	62	5.41	23.8
2010	135,181	766	64	4.99	25.7
2011	135,346	766	65	4.66	28.8
2012	124,791	770	65	3.90	39.9

Notes: Average characteristics of applications in HMDA dataset by year.

Table A5: Loan Application Outcomes by Year

Year	Denied (%)	Rejected (%)	Held (%)	Securitized (%)	Total
2000	33.88	23.67	33.56	8.89	546,472
2001	21.63	21.52	36.43	20.41	1,557,992
2002	18.44	21.56	34.84	25.15	2,019,569
2003	16.63	21.37	32.20	29.80	3,046,025
2004	22.55	24.93	36.14	16.38	1,870,789
2005	24.14	25.44	37.70	12.71	1,559,074
2006	25.59	26.75	37.67	10.00	1,290,774
2007	32.68	23.49	30.85	12.98	983,978
2008	34.76	23.60	20.00	21.64	625,338
2009	22.10	18.67	25.30	33.93	867,548
2010	21.65	15.62	30.03	32.70	889,065
2011	22.20	15.55	28.78	33.47	814,439
2012	19.23	14.06	21.99	44.71	1,038,733
Total	22.40	21.69	32.52	23.39	17,109,796

Notes: Percent of applications by disposition in HMDA dataset by year.

Table A6: Merged Dataset Average Borrower Risk Characteristics by Origination Year

Year	Price ($)	Income ($)	UPB ($)	Interest Rate	Payment ($)	LTV	FICO	DTI
2000	290,916	45,853	188,691	7.54	1,315	72	706	38.04
2001	308,341	48,591	195,151	6.88	1,283	69	717	36.31
2002	344,132	51,288	210,012	6.36	1,308	66	719	35.77
2003	363,884	49,396	213,324	5.79	1,250	63	723	35.74
2004	385,459	49,557	227,657	5.80	1,334	63	713	37.44
2005	459,430	50,997	243,616	5.77	1,422	57	721	38.12
2006	500,898	55,456	257,428	6.34	1,600	55	728	39.11
2007	537,583	60,771	287,353	6.22	1,762	57	733	39.53
2008	568,698	67,212	319,766	5.80	1,872	62	750	38.65
2009	638,570	72,373	340,083	4.98	1,821	60	766	34.80
2010	611,083	71,885	348,652	4.82	1,834	62	766	34.69
2011	580,626	66,360	332,245	4.53	1,687	63	767	34.54
2012	565,859	62,543	316,632	4.06	1,522	61	771	33.05

Notes: Average characteristics of loans in matched dataset at origination.

Figure A2: Comparison of Loan Purpose Distribution by Origination Year

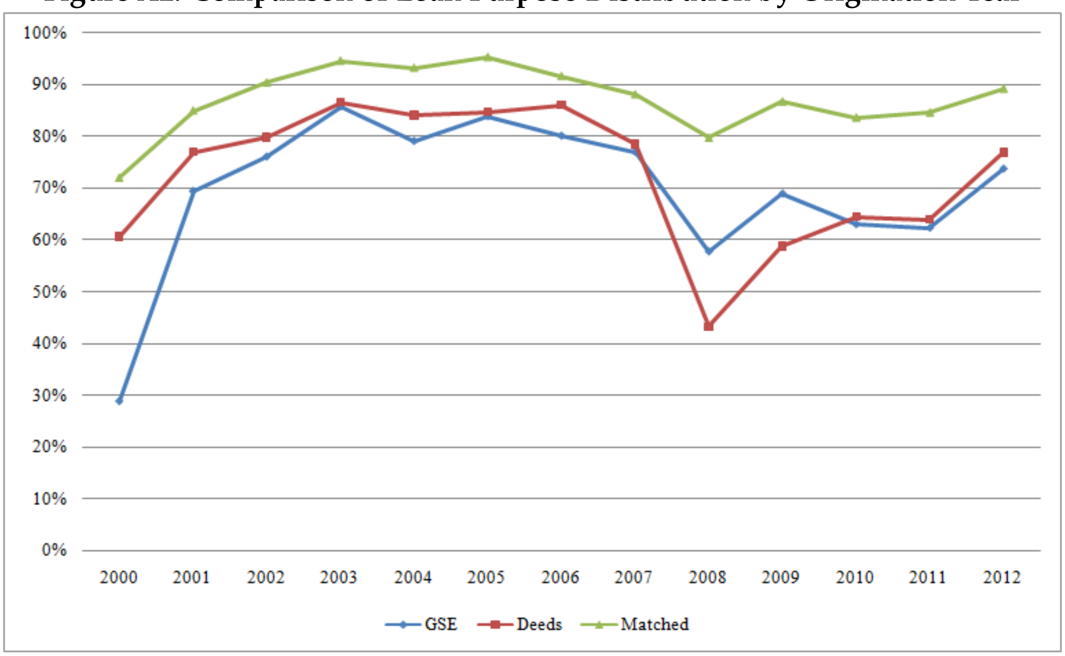

Notes: % of observations originated as refinances by origination year in GSE, deeds, and matched datasets.

Figure A3: Comparison of % of Refinances Increasing Loan Balance

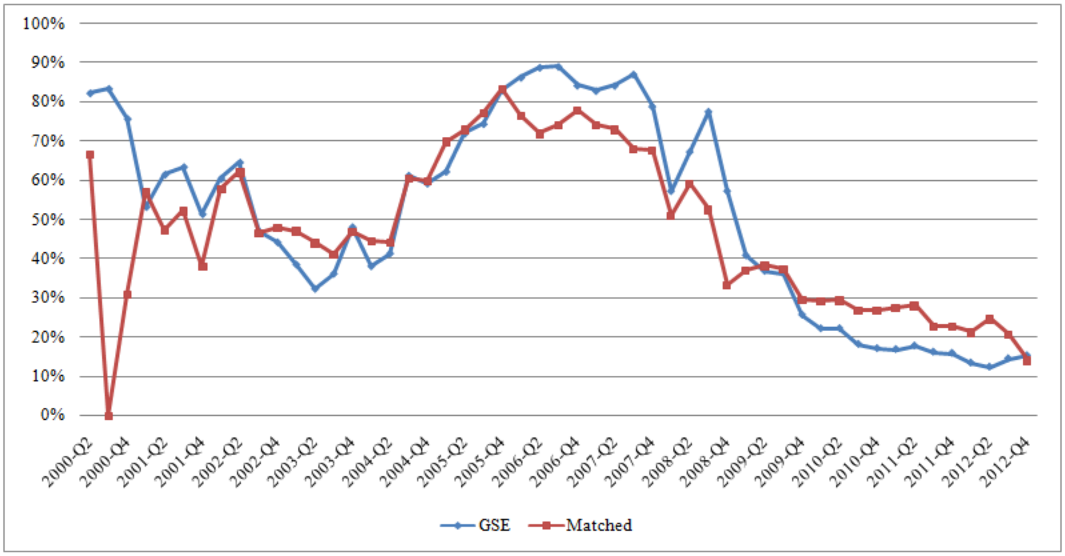

Notes: Unweighted % of observed refinances in matched dataset increasing loan balance by greater than 5% (in accordance with FHLMC definition of "increasing balance") by quarter compared with aggregate figure reported in FHLMC Refinance Report.

Figure A4: Comparison of Median Appreciation of Refinanced Homes

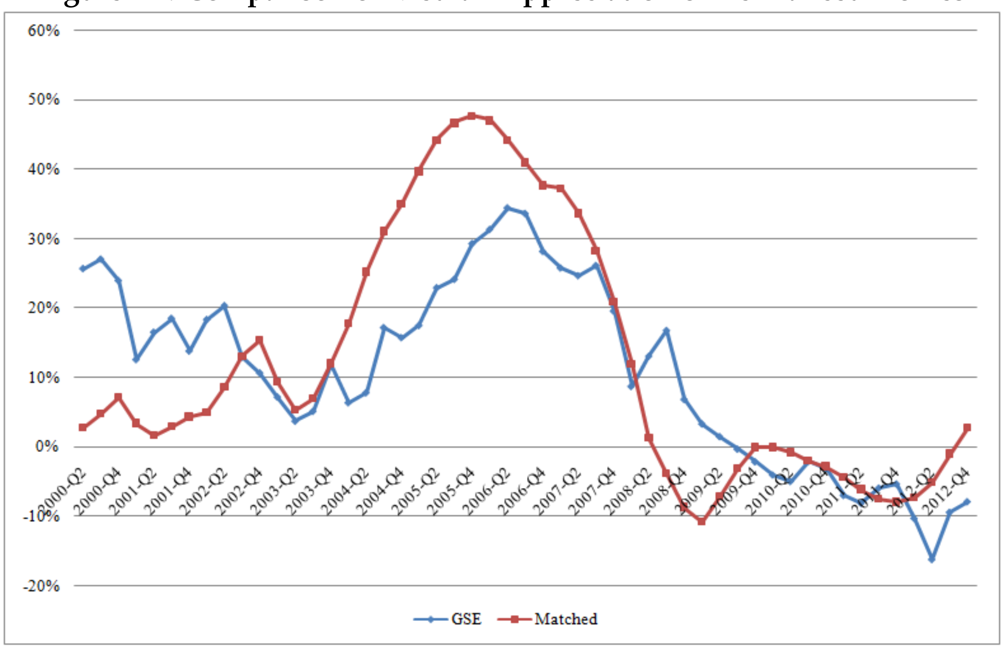

Notes: Unweighted median appreciation of observed refinances in matched dataset by quarter compared with aggregate figure reported in FHLMC Refinance Report.

Figure A5: Comparison of Cash Out as % of Refinanced Balances

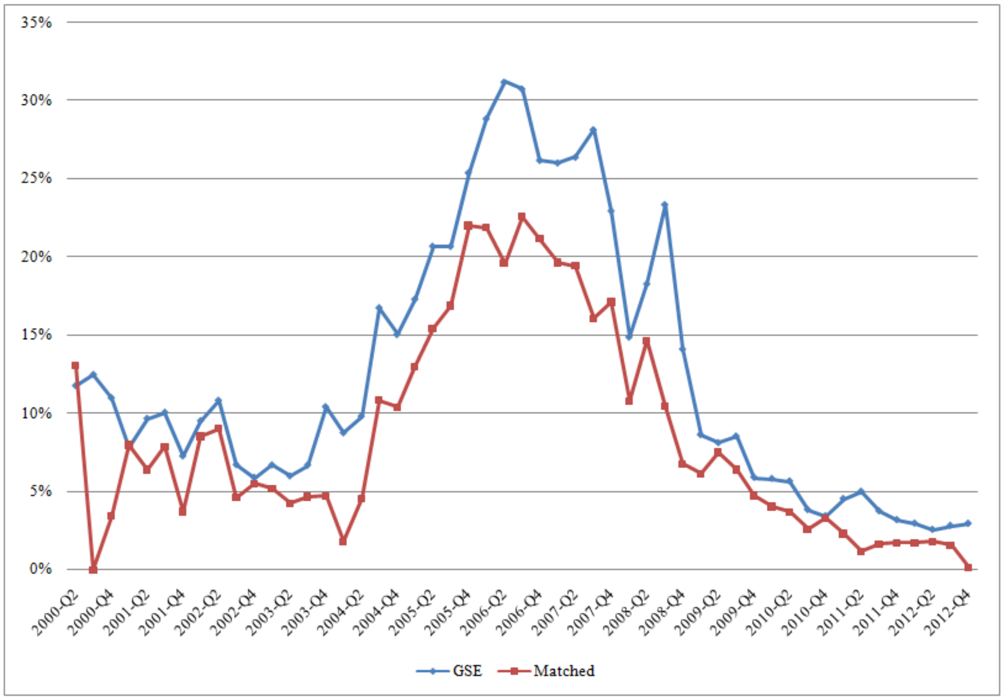

Notes: Total withdrawn equity over total refinancing volume observed in matched dataset by quarter compared with aggregate figure reported in FHLMC Refinance Report.

Appendix D: Alternate Secondary-Market Model Specifications

Table A7: Average Characteristics of MBS

Month	Avg. Size	Avg. FICO	Avg. LTV	% Refi.	% TPO	% Occ.
2011-09	201,357	749.1	74.1	67.4	32.2	81.9
2011-10	210,224	753.3	74.0	71.4	35.9	83.0
2011-11	201,654	754.4	72.9	74.6	35.6	83.7
2011-12	191,543	750.6	71.7	74.1	35.6	82.0
2012-01	192,247	752.5	73.7	78.0	33.2	81.7
2012-02	194,737	753.4	73.0	79.9	31.5	81.2
2012-03	195,475	752.8	73.8	80.7	33.2	81.1
2012-04	196,139	753.7	73.9	78.4	32.0	80.8
2012-05	190,296	749.6	75.3	76.2	31.5	81.1
2012-06	194,988	746.9	79.4	76.2	32.0	80.9
2012-07	198,178	749.6	79.2	75.3	32.3	80.7
2012-08	201,895	751.8	77.7	75.1	35.6	82.3
2012-09	199,556	749.5	78.7	75.4	32.6	81.7
2012-10	197,837	750.8	78.3	79.7	35.6	82.6
2012-11	200,935	751.7	77.6	80.1	34.1	80.7
2012-12	194,196	751.5	77.2	81.0	36.2	80.0
2013-01	198,091	749.7	77.6	81.2	35.3	80.8
2013-02	199,673	748.2	77.5	83.4	38.0	79.8
2013-03	192,707	745.0	79.0	86.6	34.7	80.0

Notes: Average characteristics of FNMA MBS by month of issuance.

Table A8: Guarantee Fee Increase Event Study with All Controls

	Model 1	Model 2	Model 3	Model 4	Model 5
Issued Post-April 2012	0.0532***	0.0838***	-	-	0.0748***
	(0.0042)	(0.0138)	-	-	(0.0020)
Issued Post-December 2012	-	-	0.0379***	0.0450***	0.0430***
	-	-	(0.0047)	(0.0059)	(0.0019)
Q3 LTV	0.0014***	0.0036***	0.0009	-0.0001	0.0021***
	(0.0004)	(0.0006)	(0.0006)	(0.0007)	(0.0001)
Q2 LTV	-0.0001	-0.0009	-0.0016*	-0.0006	-0.0014***
	(0.0004)	(0.0006)	(0.0009)	(0.0012)	(0.0002)
Q1 LTV	-0.0012***	-0.0020***	-0.0005	-0.0006	-0.0008***
	(0.0003)	(0.0004)	(0.0005)	(0.0006)	(0.0001)
Q3 FICO	-0.0002	-0.0009***	0.0000	0.0001	-0.0003***
	(0.0002)	(0.0002)	(0.0003)	(0.0003)	(0.0000)
Q2 FICO	0.0000	0.0009***	0.0002	-0.0001	0.0000
	(0.0002)	(0.0002)	(0.0003)	(0.0003)	(0.0001)
Q1 FICO	0.0003***	0.0003***	0.0000	0.0000	0.0003***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0000)
Q3 Int. Rt.	0.0145	0.1210***	0.0808***	0.0539	0.1018***
	(0.0238)	(0.0294)	(0.0277)	(0.0366)	(0.0069)
Q2 Int. Rt.	0.1251***	0.0129	0.0190	0.0438	0.1612***
	(0.0256)	(0.0311)	(0.0324)	(0.0412)	(0.0091)
Q1 Int. Rt.	-0.0656***	-0.0347*	-0.0601***	-0.0703***	-0.1900***
	(0.0151)	(0.0182)	(0.0154)	(0.0203)	(0.0077)
Orig. WAM	-0.0002***	-0.0001	0.0004***	0.0006***	-0.0001***
	(0.0000)	(0.0001)	(0.0000)	(0.0001)	(0.0000)
% Refi.	-0.0001	0.0000	-0.0001	0.0001	-0.0004***
	(0.0001)	(0.0001)	(0.0001)	(0.0002)	(0.0000)
% Occ.	0.0004***	0.0004***	0.0005***	0.0004***	0.0005***
	(0.0001)	(0.0001)	(0.0001)	(0.0002)	(0.0000)
Implementation Date	April 2012	April 2012	December 2012	December 2012	Both
Sample Window	2 Weeks	1 Week	2 Weeks	1 Week	6 Months
Fixed Effects	State	State	State	State	State, Orig. Year
N	2,696	1,554	3,646	2,338	58,106
R^2	0.1975	0.2045	0.1488	0.1615	0.1459

Notes: Dependent variable is difference between MBS coupon and WAC. Standard errors in parentheses. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A9: Guarantee Fee Increase Event Study with All Controls

	Est. FE	SE
Oct. 2011	-0.027***	(0.0040)
Nov. 2011	-0.038***	(0.0039)
Dec. 2011	-0.020***	(0.0038)
Jan. 2012	-0.024***	(0.0039)
Feb. 2012	-0.012**	(0.0050)
Mar. 2012	-0.015***	(0.0053)
Apr. 2012	0.013**	(0.0057)
May 2012	0.031***	(0.0057)
Jun. 2012	0.067***	(0.0057)
Jul. 2012	0.060***	(0.0058)
Aug. 2012	0.043***	(0.0058)
Sep. 2012	0.055***	(0.0060)
Oct. 2012	0.086***	(0.0061)
Nov. 2012	0.073***	(0.0061)
Dec. 2012	0.096***	(0.0064)
Jan. 2013	0.087***	(0.0062)
Feb. 2013	0.079***	(0.0059)
Mar. 2013	0.121***	(0.0060)
Apr. 2013	0.159***	(0.0058)
Implementation Date	Both	
Sample Window	6 Months	
Fixed Effects	State, Orig. Year	
N	58,106	
R^2	0.1679	

Notes: Dependent variable is difference between MBS coupon and WAC. Standard errors in parentheses. Additional controls include quartiles of LTV, FICO and interest rate, weighted-average maturity, and % refinance and owner-occupied. Bold months denote policy intervention dates. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A10: 2SLS MBS Coupon Event Study

	Model 1	Model 2	Model 3	Model 4	Model 5
Predicted G-Fee	-1.2294***	-1.2138***	-0.9264***	-0.9900***	-1.0112***
	(0.0850)	(0.2058)	(0.1399)	(0.1336)	(0.0206)
Implementation Date	April 2012	April 2012	December 2012	December 2012	Both
Sample Window	2 Weeks	1 Week	2 Weeks	1 Week	6 Months
Fixed Effects	State	State	State	State	State/Orig. Year
N	2,696	1,554	3,646	2,338	58,106
R^2	0.9447	0.9518	0.9258	0.9326	0.9528

Notes: Dependent variable is MBS coupon. Standard errors in parentheses. Additional controls include quartiles of LTV, FICO and interest rate, weighted-average maturity, and % refinance and owner-occupied. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A11: MBS Coupon Event Study with All Controls

	Model 1	Model 2	Model 3	Model 4	Model 5
Issued Post-April 2012	-0.0614***	-0.0637***	-	-	-0.0759***
	(0.0042)	(0.0108)	-	-	(0.0022)
Issued Post-December 2012	-	-	-0.0276***	-0.0443***	-0.0432***
	-	-	(0.0042)	(0.0060)	(0.0020)
Q3 LTV	-0.0027***	-0.0045***	-0.0015**	-0.0004	-0.0023***
	(0.0005)	(0.0007)	(0.0006)	(0.0008)	(0.0001)
Q2 LTV	-0.0001	0.0003	0.0025**	0.0014	0.0014***
	(0.0005)	(0.0007)	(0.0010)	(0.0013)	(0.0002)
Q1 LTV	0.0018***	0.0028***	0.0001	0.0002	0.0009***
	(0.0003)	(0.0004)	(0.0005)	(0.0006)	(0.0001)
Q3 FICO	0.0002	0.0012***	-0.0001	-0.0002	0.0003***
	(0.0002)	(0.0002)	(0.0003)	(0.0003)	(0.0000)
Q2 FICO	-0.0002	-0.0011***	-0.0001	0.0002	0.0000
	(0.0002)	(0.0003)	(0.0003)	(0.0004)	(0.0001)
Q1 FICO	-0.0005***	-0.0006***	-0.0002*	-0.0002	-0.0003***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0000)
Q3 Int. Rt.	0.4249***	0.3296***	0.2888***	0.3511***	0.2596***
	(0.0274)	(0.0336)	(0.0299)	(0.0393)	(0.0073)
Q2 Int. Rt.	0.2894***	0.4162***	0.4836***	0.4497***	0.1157***
	(0.0295)	(0.0356)	(0.0350)	(0.0443)	(0.0097)
Q1 Int. Rt.	0.1839***	0.1250***	0.1800***	0.1692***	0.5545***
	(0.0174)	(0.0208)	(0.0166)	(0.0218)	(0.0082)
Orig. WAM	0.0003***	0.0002**	-0.0004***	-0.0006***	0.0001***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0000)
% Refi.	0.0002	0.0001	0.0002	0.0000	0.0004***
	(0.0001)	(0.0002)	(0.0001)	(0.0002)	(0.0000)
% Occ.	-0.0006***	-0.0006***	-0.0005***	-0.0004**	-0.0006***
	(0.0001)	(0.0001)	(0.0001)	(0.0002)	(0.0000)
Implementation Date	April 2012	April 2012	December 2012	December 2012	Both
Sample Window	2 Weeks	1 Week	2 Weeks	1 Week	6 Months
Fixed Effects	State	State	State	State	State/Orig. Year
N	2,696	1,554	3,646	2,338	58,106
R^2	0.9447	0.9518	0.9258	0.9326	0.9528

Notes: Dependent variable is MBS coupon. Standard errors in parentheses. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A12: 2SLS MBS Price Event Study

	Model 1	Model 2	Model 3	Model 4	Model 5
Predicted G-Fee	-4.5063***	-6.2330***	-14.9385***	-8.2187***	-2.9213***
	(0.3034)	(0.5456)	(0.7409)	(1.0492)	(0.0912)
Implementation Date	April 2012	April 2012	December 2012	December 2012	Both
Sample Window	2 Weeks	1 Week	2 Weeks	1 Week	6 Months
Fixed Effects	State/Day	State/Day	State/Day	State/Day	State/Day/Orig. Year
N	20,512	6,582	23,777	7,922	499,321
$\mathbf{R^2}$	0.8718	0.8958	0.7453	0.7529	0.8029

Notes: Dependent variable is MBS price. Standard errors in parentheses. Additional controls include quartiles of LTV, FICO and interest rate, weighted-average maturity, days from issuance, and % refinance, third-party and owner-occupied. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A13: MBS Price Event Study with All Controls

	Model 1	Model 2	Model 3	Model 4	Model 5
Issued Post-April 2012	-0.2251***	-0.3272***	-	-	-0.2058***
	(0.0152)	(0.0286)	-	-	(0.0045)
Issued Post-December 2012	-	-	-0.4447***	-0.3676***	-0.3515***
	-	-	(0.0221)	(0.0469)	(0.0057)
Q3 LTV	-0.0156***	-0.0324***	0.0158***	0.0240***	0.0081***
	(0.0010)	(0.0020)	(0.0015)	(0.0025)	(0.0004)
Q2 LTV	-0.0002	0.0040**	0.0297***	0.0221***	0.0032***
	(0.0010)	(0.0020)	(0.0023)	(0.0040)	(0.0004)
Q1 LTV	0.0114***	0.0176***	-0.0029**	-0.0023	0.0063***
	(0.0006)	(0.0012)	(0.0012)	(0.0020)	(0.0003)
Q3 FICO	0.0030***	0.0104***	-0.0027***	-0.0035***	0.0002*
	(0.0004)	(0.0007)	(0.0007)	(0.0011)	(0.0001)
Q2 FICO	-0.0027***	-0.0061***	-0.0022***	-0.0019	0.0021***
	(0.0004)	(0.0008)	(0.0007)	(0.0011)	(0.0002)
Q1 FICO	-0.0036***	-0.0088***	-0.0008***	-0.0012***	-0.0042***
	(0.0002)	(0.0004)	(0.0002)	(0.0004)	(0.0001)
Q3 Int. Rt.	1.4293***	1.0774***	1.9460***	1.4995***	0.6006***
	(0.0526)	(0.0941)	(0.0744)	(0.1191)	(0.0153)
Q2 Int. Rt.	1.4860***	1.3035***	-0.5861***	-0.3439***	0.2837***
	(0.0564)	(0.0970)	(0.0825)	(0.1332)	(0.0198)
Q1 Int. Rt.	0.4819***	1.1504***	0.2018***	0.2108***	1.1820***
	(0.0351)	(0.0694)	(0.0380)	(0.0655)	(0.0168)
Orig. WAM	-0.0165***	-0.0179***	-0.0035***	-0.0013***	-0.0052***
	(0.0001)	(0.0002)	(0.0001)	(0.0002)	(0.0000)
% Refi.	-0.0005**	-0.0005	0.0120***	0.0134***	0.0034***
	(0.0002)	(0.0004)	(0.0003)	(0.0005)	(0.0001)
% Occ.	-0.0115***	-0.0117***	-0.0135***	-0.0124***	-0.0053***
	(0.0002)	(0.0004)	(0.0004)	(0.0006)	(0.0001)
Days Since Issue	0.0041*	-0.0271***	-0.0335***	-0.0390***	0.0005***
	(0.0024)	(0.0078)	(0.0032)	(0.0116)	(0.0000)
Implementation Date	April 2012	April 2012	December 2012	December 2012	Both
Sample Window	2 Weeks	1 Week	2 Weeks	1 Week	6 Months
Fixed Effects	State/Day	State/Day	State/Day	State/Day	State/Day/Orig. Year
N	20,512	6,582	23,777	7,922	499,321
$\mathbf{R^2}$	0.8718	0.8958	0.7453	0.7529	0.8036

Notes: Dependent variable is MBS price. Standard errors in parentheses. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Appendix E: Alternate Primary-Market Model Specifications

Table A14: Average Characteristics of Mortgages

Month	Avg. Size	Avg. FICO	Avg. LTV	Avg. DTI	% Cash	% TPO	% Occ.
2011-09	289,952	766.5	66.3	31.6	9.7%	60.4%	8.3%
2011-10	263,521	767.2	66.1	31.4	9.3%	61.3%	8.7%
2011-11	261,753	767.0	65.7	31.4	16.3%	59.0%	8.7%
2011-12	252,414	765.6	65.7	31.5	28.6%	53.6%	9.9%
2012-01	256,033	766.7	66.0	31.3	29.5%	54.8%	9.5%
2012-02	258,782	767.3	66.0	31.2	28.8%	52.4%	9.8%
2012-03	252,374	766.6	66.1	31.4	30.6%	54.4%	10.0%
2012-04	250,676	764.8	66.2	31.7	33.8%	51.9%	9.9%
2012-05	253,363	763.7	66.7	31.8	33.4%	53.8%	10.5%
2012-06	262,992	766.4	66.6	31.2	30.4%	53.7%	9.8%
2012-07	266,983	767.4	66.4	30.9	28.7%	54.1%	9.3%
2012-08	267,388	767.5	66.7	30.8	27.5%	54.2%	8.9%
2012-09	260,917	765.9	67.0	31.0	28.3%	51.2%	9.4%
2012-10	263,267	766.0	67.4	30.9	26.4%	53.6%	9.6%
2012-11	271,436	766.0	67.2	30.6	28.6%	58.0%	11.3%
2012-12	262,148	765.7	67.0	30.9	27.3%	55.2%	10.1%
2013-01	255,615	764.7	67.0	31.0	27.8%	56.4%	10.5%
2013-02	244,350	761.7	67.3	31.6	29.6%	60.1%	11.7%
2013-03	238,822	759.5	67.6	31.7	29.0%	58.8%	11.2%

Notes: Average characteristics of originated refinance mortgages by month. Size denotes loan balance, Cash/TPO/Occ. respectively the fraction of cash-out refinances, third-party originations, and owner-occupants.

Figure A6: Correlation Between Secondary-Swap Spread and MOVE Index

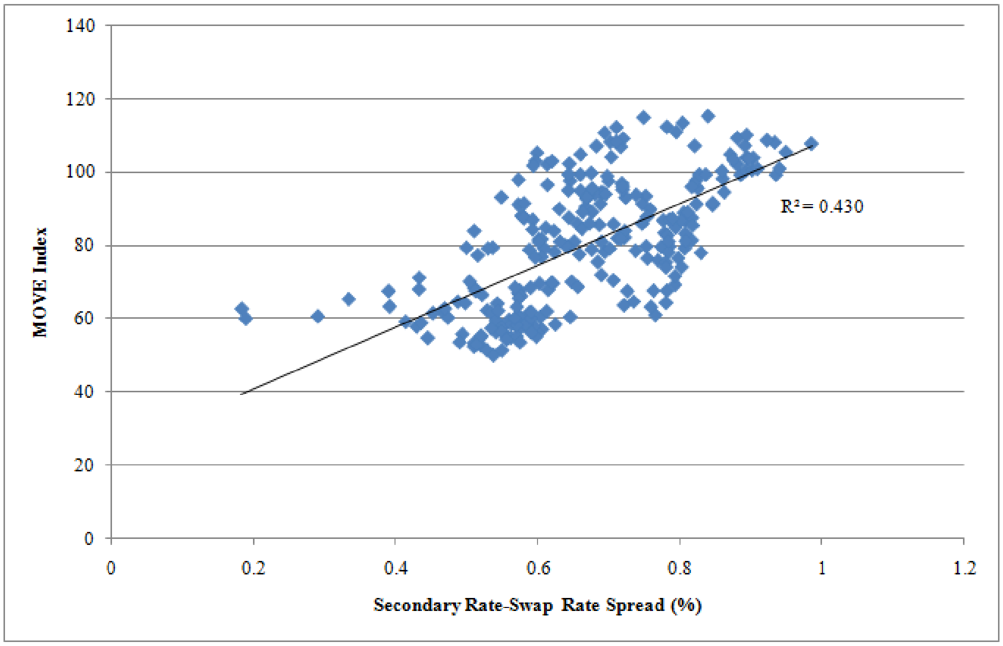

Notes: Correlation between spread of secondary market rate over option-implied interest rate volatility index. Observations at week-level using 3-week centered moving averages.

Table A15: Mortgage Interest Rate Event Study Robustness Tests

Post-Jan. 2012	.0917***	.0861***	.0751***	.0885***
	(1.25E-3)	(1.26E-3)	(1.94E-3)	(1.25E-3)
Post-Sept. 2012	.0834***	.0694***	.1522***	.0815***
	(4.75E-4)	(4.85E-4)	(1.38E-3)	(4.92E-4)
FICO	-.0019***	-.0019***	-.0019***	-.0019***
	(4.74E-6)	(4.72E-6)	(4.74E-6)	(4.74E-6)
LTV	.0031***	.0031***	.0031***	.0031***
	(1.12E-5)	(1.12E-5)	(1.12E-5)	(1.12E-5)
DTI	.0012***	.0012***	.0012***	.0012***
	(1.71E-5)	(1.70E-5)	(1.71E-5)	(1.71E-5)
Cash-Out Refi.	.0752***	-.0181***	.0750***	.0752***
	(5.29E-4)	(1.21E-3)	(5.29E-4)	(5.28E-4)
Post-January 2012 x Cash-Out	-	.0720***	-	-
	-	(1.35E-3)	-	-
Post-Jan. 2012 x Cash-Out	-	.1564***	-	-
	-	(1.20E-3)	-	-
Post-Jan. 2012 x HHI	-	-	.0594***	-
	-	-	(6.03E-3)	-
Post-Sept. 2012 x HHI	-	-	-.2715***	-
	-	-	(5.15E-3)	-
Post-Jan. 2012 x Non-Bank	-	-	-	.0246***
	-	-	-	(7.69E-4)
Post-Sept. 2012 x Non-Bank	-	-	-	.0092***
	-	-	-	(1.08E-3)
N	3,329,478	3,329,478	3,329,478	3,329,478
R^2	0.2974	0.3024	0.2980	0.2980

Notes: Dependent variable is spread of primary mortgage rate over secondary mortgage rate, sample is all loans originated between September 2011 and March 2013 in GSE dataset. Standard errors in parentheses. Additional controls include origination balance, interest rate volatility, state house price indices and unemployment rates, and indicators for investor-owned, third-party origination, and condominium properties. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A16: Jumbo-Conforming Loan Spread Event Study

	Model 1	Model 2	Model 3
Conforming Refi. Rate	0.6569***	0.6382***	0.5705***
	(0.0152)	(0.0314)	(0.0165)
Post-Jan. 2012	-0.1069***	-0.1001***	-
	(0.0052)	(0.0058)	-
Post-Sept. 2012	-0.1149***	-	-0.1126***
	(0.0127)	-	(0.0104)
Sample Window	6 Months	6 Months	6 Months
State FE?	Yes	Yes	Yes
Population Weights?	Yes	Yes	Yes
N	20,145	11,730	11,628
R^2	0.9272	0.8945	0.7576

Notes: Dependent variable is jumbo mortgage refinance rate at State-Day Level. Jumbo and conforming mortgage rates measured as 5-day centered moving averages. Standard errors in parentheses. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A17: Refinancing Probability Event Study with Additional Controls

	Model 1	**Model 2**	**Model 3**	**Model 4**
Post-Announcement	-0.0017***	-0.0015***	-0.0024***	-0.0016***
	(0.0001)	(0.0000)	(0.0001)	(0.0001)
FICO	8.13E-05***	6.67E-05***	8.40E-05***	8.09E-05***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\widehat{\textbf{Equity}}$	2.45E-07***	2.15E-07***	2.91E-07***	2.76E-07***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Loan Age	0.0003***	0.0004***	0.0004***	0.0004***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Initial LTV	0.0014***	0.0013***	0.0018***	0.0017***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\widehat{\textbf{Unemp. Rate}}$	-0.0002***	0.0047***	-0.0032***	0.0038***
	(0.0002)	(0.0001)	(0.0002)	(0.0001)
Initial Int. Rt. Spread	0.0156***	0.0132***	0.0114***	0.0116***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Δ Mtg. Rt.	-0.0293***	-0.0214***	-0.0212***	-0.0193***
	(0.0002)	(0.0001)	(0.0001)	(0.0001)
State FE?	Yes	Yes	Yes	Yes
Lender FE?	Yes	Yes	Yes	Yes
Announcement Date	Jan. 2012	Jan. 2012	Sept. 2012	Sept. 2012
Sample Window	3 Months	6 Months	3 Months	6 Months
N	29,872,035	52,340,967	36,057,155	63,184,479
$\mathbf{R^2}$	0.0763	0.0666	0.0900	0.0857

Notes: Dependent variable is indicator for prepayment, sample is all loans originated after June 2009 in GSE dataset. Standard errors clustered at individual-level in parentheses. Hats denote instrumented variables. Additional controls include indicators for owner-occupied, third-party origination, initial cash-out refinance, initial purchase loan, condominiums, manufactured housing, and planned-unit developments. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A18: Refinancing Probability Event Study at Implementation Date

	Model 1	Model 2	Model 3	Model 4
Post-Implementation	-0.0032***	-0.0042***	0.0000	-0.0020***
	(0.0001)	(0.0000)	(0.0001)	(0.0001)
FICO	7.51E-05***	8.10E-05***	8.46E-05***	7.23E-05***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\widehat{\textbf{Equity}}$	2.43E-07***	2.62E-07***	2.96E-07***	2.67E-07***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Loan Age	0.0005***	0.0004***	0.0003***	0.0005***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Initial LTV	0.0014***	0.0016***	0.0018***	0.0017***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\widehat{\textbf{Unemp. Rate}}$	0.0050***	0.0026***	0.0085***	0.0013***
	(0.0002)	(0.0001)	(0.0002)	(0.0001)
Initial Int. Rt. Spread	0.0137***	0.0137***	0.0102***	0.0103***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Δ Mtg. Rt.	-0.0184***	-0.0214***	-0.0212***	-0.0152***
	(0.0002)	(0.0001)	(0.0001)	(0.0001)
State FE?	Yes	Yes	Yes	Yes
Lender FE?	Yes	Yes	Yes	Yes
Implementation Date	Apr. 2012	Apr. 2012	Dec. 2012	Dec. 2012
Sample Window	3 Months	6 Months	3 Months	6 Months
N	32,080,570	56,274,415	38,586,254	66,539,436
$\mathbf{R^2}$	0.0738	0.0814	0.0926	0.0825

Notes: Dependent variable is indicator for prepayment, sample is all loans originated after June 2009 in GSE dataset. Standard errors clustered at individual-level in parentheses. Hats denote instrumented variables. Additional controls include indicators for owner-occupied, third-party origination, initial cash-out refinance, initial purchase loan, condominiums, manufactured housing, and planned-unit developments. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A19: Refinancing Probability Event Study Without Instruments

	Model 1	Model 2	Model 3	Model 4
Post-Announcement	-0.0018***	-0.0031***	-0.0026***	-0.0033***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
FICO	7.87E-05***	6.39E-05***	8.07E-05***	7.78E-05***
	(6.49E-07)	(5.06E-07)	(6.67E-07)	(5.48E-07)
Equity	2.73E-07***	2.39E-07***	3.25E-07***	3.08E-07***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Loan Age	0.0004***	0.0004***	0.0004***	0.0004***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Initial LTV	0.0016***	0.0014***	0.0020***	0.0019***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Unemp. Rate	0.0000	0.0002***	-0.0013***	0.0006***
	(0.0001)	(0.0000)	(0.0001)	(0.0000)
Initial Int. Rt. Spread	0.0151***	0.0125***	0.0107***	0.0109***
	(0.0001)	(0.0001)	(0.0002)	(0.0001)
Δ Mtg. Rt.	-0.0295***	-0.0194***	-0.0208***	-0.0191***
	(0.0002)	(0.0001)	(0.0002)	(0.0001)
State FE?	Yes	Yes	Yes	Yes
Lender FE?	Yes	Yes	Yes	Yes
Announcement Date	Jan. 2012	Jan. 2012	Sept. 2012	Sept. 2012
Sample Window	3 Months	6 Months	3 Months	6 Months
N	29,872,035	52,340,967	36,057,155	63,184,479
R^2	0.0845	0.0737	0.0998	0.0945

Notes: Dependent variable is indicator for prepayment, sample is all loans originated after June 2009 in GSE dataset. Standard errors clustered at individual-level in parentheses. Additional controls include indicators for owner-occupied, third-party origination, initial cash-out refinance, initial purchase loan, condominiums, manufactured housing, and planned-unit developments. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A20: Default Probability Event Study Without Instruments

	Model 1	**Model 2**	**Model 3**	**Model 4**
Post-Announcement	8.08E-05***	1.20E-04***	3.71E-06	3.59E-05***
	(8.25E-06)	(5.53E-06)	(1.27E-05)	(6.36E-06)
FICO	-2.69E-06***	-2.48E-06***	-2.57E-06***	-2.81E-06***
	(8.67E-08)	(6.61E-08)	(7.70E-08)	(6.02E-08)
Home Price App.	-2.76E-04	3.47E-05	-9.40E-04***	-7.56E-04***
	(2.63E-04)	(1.49E-04)	(1.07E-04)	(6.46E-05)
Loan Age	8.14E-06***	9.01E-06***	4.57E-06***	5.99E-06***
	(3.62E-07)	(2.68E-07)	(3.22E-07)	(2.50E-07)
Initial LTV	4.77E-06***	4.32E-06***	4.09E-06***	4.58E-06***
	(1.73E-07)	(1.35E-07)	(1.57E-07)	(1.21E-07)
Unemp. Rate	-9.37E-05***	-7.67E-05***	-1.02E-04***	-4.39E-05***
	(7.21E-06)	(4.78E-06)	(9.82E-06)	(4.15E-06)
Initial Int. Rt.	5.49E-05***	5.68E-05***	8.19E-05***	5.89E-05***
	(8.88E-06)	(6.79E-06)	(8.07E-06)	(6.12E-06)
Δ Mtg. Rt.	8.98E-05*	2.96E-04***	3.80E-04***	2.59E-04***
	(5.33E-05)	(1.22E-05)	(3.34E-05)	(1.46E-05)
State FE?	Yes	Yes	Yes	Yes
Lender FE?	Yes	Yes	Yes	Yes
Announcement Date	Jan. 2012	Jan. 2012	Sept. 2012	Sept. 2012
Sample Window	3 Months	6 Months	3 Months	6 Months
N	29,872,035	52,340,967	36,057,155	63,184,479
R^2	0.0006	0.0007	0.0005	0.0006

Notes: Dependent variable is indicator for default, sample is all loans originated after June 2009 in GSE dataset. Standard errors clustered at individual-level in parentheses. Additional controls include initial DTI ratio, state-level home value indices, indicators for owner-occupied, third-party origination, initial cash-out refinance, initial purchase loan, condominiums, manufactured housing, and planned-unit developments. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Appendix F: Alternate Supply-Demand Interaction Model Specifications

Table A21: Baseline Credit Constraint Probit Model

	Model 1	Model 2	Model 3	Model 4
Income	-2.34E-7***	-2.32E-7***	-2.31E-7***	-2.36E-7***
	(4.90E-9)	(4.87E-9)	(4.84E-9)	(4.92E-9)
Orig. Balance	-6.57E-8***	-7.96E-8***	-1.41E-7***	-3.71E-8***
	(3.73E-9)	(3.75E-9)	(3.79E-9)	(3.76E-9)
Owner-Occ.	3.11E-2***	3.12E-2***	3.15E-2***	2.93E-2***
	(1.15E-3)	(1.15E-3)	(1.15E-3)	(1.16E-3)
FICO	-9.28E-4***	-1.88E-4***	-8.43E-4***	-9.74E-4***
	(2.45E-5)	(3.16E-5)	(2.44E-5)	(2.45E-5)
LTV	-4.20E-3***	-4.14E-3***	-3.72E-3***	-7.64E-3***
	(3.22E-5)	(3.25E-5)	(3.27E-5)	(6.36E-5)
DTI	0.0205***	0.0202***	0.0048***	0.0206***
	(9.95E-5)	(9.97E-5)	(1.67E-4)	(9.96E-5)
Mtg. Spread	-0.0111***	-0.0095***	-0.0124***	-0.0072***
	(7.68E-4)	(7.69E-4)	(7.68E-4)	(7.69E-4)
G-Fee	0.0007***	0.0003***	0.0008***	0.0013***
	(9.35E-5)	(9.38E-5)	(9.58E-5)	(9.38E-5)
$\widehat{\text{Unemp.}}$ Rate	0.0015***	-	-	-
	(3.45E-4)	-	-	-
$\widehat{\text{Unemp.}}$ x Bin 1	-	0.0234***	0.0272***	0.0112***
	-	(2.83E-3)	(4.11E-4)	(3.61E-4)
$\widehat{\text{Unemp.}}$ x Bin 2	-	0.0022***	0.0159***	0.0027***
	-	(6.47E-4)	(3.80E-4)	(3.54E-4)
$\widehat{\text{Unemp.}}$ x Bin 3	-	0.0189***	0.0104***	-0.0029***
	-	(4.31E-4)	(3.65E-4)	(3.54E-4)
$\widehat{\text{Unemp.}}$ x Bin 4	-	0.0073***	0.0041***	-0.0061***
	-	(3.56E-4)	(3.49E-4)	(3.63E-4)
$\widehat{\text{Unemp.}}$ x Bin 5	-	-0.0060***	-0.0022***	-0.0086***
	-	(9.24E-4)	(3.56E-4)	(4.09E-4)
Year FE?	No	No	No	No
County FE?	Yes	Yes	Yes	Yes
Bin Type	None	FICO	DTI	LTV
N	16,613,070	16,613,070	16,613,070	16,613,070
Pseudo-R^2	0.0155	0.0158	0.0161	0.0160

Notes: Dependent variable is indicator for either denial or offer rejection, sample is all loans in HMDA dataset. Robust standard errors in parentheses. Hats denote instrumented variables. Interaction terms use equally-sized quintiles of the corresponding variable (higher bins correspond to lower credit risk). ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A22: Alternate Credit Constraint Probit Model Without Instruments

	Model 1	Model 2	Model 3	Model 4
Income	-2.48E-7***	-2.45E-7***	-2.44E-7***	-2.49E-7***
	(5.13E-9)	(5.09E-9)	(5.08E-9)	(5.14E-9)
Orig. Balance	-3.49E-7***	-3.24E-7***	-3.86E-7***	-3.09E-7***
	(3.57E-9)	(3.58E-9)	(3.59E-9)	(3.62E-9)
Owner-Occ.	2.87E-2***	2.76E-2***	2.82E-2***	2.74E-2***
	(1.15E-3)	(1.15E-3)	(1.15E-3)	(1.15E-3)
FICO	-3.37E-3***	-2.85E-3***	-2.88E-3***	-3.26E-3***
	(2.08E-5)	(2.99E-5)	(2.11E-5)	(2.08E-5)
LTV	-2.14E-3***	-2.29E-3***	-1.76E-3***	-6.07E-3***
	(3.04E-5)	(3.07E-5)	(3.06E-5)	(5.43E-5)
DTI	0.0247***	0.0243***	0.0105***	0.0244***
	(9.81E-5)	(9.83E-5)	(1.56E-4)	(9.81E-5)
Mtg. Spread	-0.0364***	-0.0332***	-0.0357***	-0.0286***
	(7.44E-4)	(7.45E-4)	(7.43E-4)	(7.46E-4)
G-Fee	0.0047***	0.0055***	0.0046***	0.0046***
	(7.25E-5)	(7.30E-5)	(7.42E-5)	(7.26E-5)
Unemp. Rate	0.0181***	-	-	-
	(1.31E-4)	-	-	-
Unemp. x Bin 1	-	0.0164***	0.0374***	0.0296***
	-	(8.82E-4)	(2.03E-4)	(1.63E-4)
Unemp. x Bin 2	-	0.0161***	0.0219***	0.0200***
	-	(1.48E-4)	(1.69E-4)	(1.51E-4)
Unemp. x Bin 3	-	0.0238***	0.0177***	0.0131***
	-	(1.61E-4)	(1.56E-4)	(1.54E-4)
Unemp. x Bin 4	-	-0.0075***	0.0134***	0.0091***
	-	(5.46E-4)	(1.48E-4)	(1.63E-4)
Unemp. x Bin 5	-	-0.0149***	0.0091***	0.0069***
	-	(2.88E-3)	(1.72E-4)	(2.30E-4)
Year FE?	No	No	No	No
County FE?	Yes	Yes	Yes	Yes
Bin Type	None	FICO	DTI	LTV
N	16,613,070	16,613,070	16,613,070	16,613,070
Pseudo-R^2	0.0162	0.0163	0.0164	0.0165

Notes: Dependent variable is indicator for either denial or offer rejection, sample is all loans in HMDA dataset. Robust standard errors in parentheses. Interaction terms use equally-sized quintiles of the corresponding variable (higher bins correspond to lower credit risk). ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A23: Alternate Credit Constraint Probit Model with Denial

	Model 1	Model 2	Model 3	Model 4
Income	-3.24E-7***	-3.23E-7***	-3.24E-7***	-3.25E-7***
	(8.20E-9)	(8.19E-9)	(8.20E-9)	(8.22E-9)
Orig. Balance	-1.73E-7***	-1.80E-7***	-1.77E-7***	-1.72E-7***
	(5.40E-9)	(5.41E-9)	(5.40E-9)	(5.40E-9)
Owner-Occ.	-3.03E-2***	-3.03E-2***	-3.03E-2***	-3.09E-2***
	(1.32E-3)	(1.32E-3)	(1.32E-3)	(1.32E-3)
FICO	-1.66E-5	4.08E-4***	-2.38E-5	-1.70E-4***
	(2.98E-5)	(3.69E-5)	(2.98E-5)	(2.99E-5)
LTV	1.39E-3***	1.32E-3***	1.36E-3***	-6.98E-4***
	(5.32E-5)	(5.35E-5)	(5.33E-5)	(8.28E-5)
DTI	-0.0035***	-0.0036***	-0.0062***	-0.0034***
	(1.20E-4)	(1.21E-4)	(1.84E-4)	(1.20E-4)
Mtg. Spread	0.0029***	0.0024**	0.0029***	0.0032***
	(1.00E-3)	(1.00E-3)	(1.00E-3)	(1.00E-3)
G-Fee	-0.0002	-0.0002	-0.0001	0.0000
	(1.66E-4)	(1.66E-4)	(1.66E-4)	(1.66E-4)
$\widehat{\textbf{Unemp. Rate}}$	0.0070***	-	-	-
	(5.68E-4)	-	-	-
$\widehat{\textbf{Unemp.}}$ **x Bin 1**	-	0.0501***	0.0109***	0.0078***
	-	(3.00E-3)	(6.11E-4)	(5.70E-4)
$\widehat{\textbf{Unemp.}}$ **x Bin 2**	-	0.0054***	0.0078***	0.0036***
	-	(8.20E-4)	(5.88E-4)	(5.82E-4)
$\widehat{\textbf{Unemp.}}$ **x Bin 3**	-	0.0150***	0.0090***	0.0014**
	-	(6.08E-4)	(5.77E-4)	(5.92E-4)
$\widehat{\textbf{Unemp.}}$ **x Bin 4**	-	0.0074***	0.0071***	-0.0008
	-	(5.69E-4)	(5.73E-4)	(6.05E-4)
$\widehat{\textbf{Unemp.}}$ **x Bin 5**	-	0.0052***	0.0052***	-0.0037***
	-	(1.06E-3)	(5.78E-4)	(6.46E-4)
Year FE?	No	No	No	No
County FE?	Yes	Yes	Yes	Yes
Bin Type	None	FICO	DTI	LTV
N	16,613,070	16,613,070	16,613,070	16,613,070
Pseudo-R^2	0.0231	0.0232	0.0231	0.0232

Notes: Dependent variable is indicator for application denial, sample is all loans in HMDA dataset. Robust standard errors in parentheses. Hats denote instrumented variables. Interaction terms use equally-sized quintiles of the corresponding variable (higher bins correspond to lower credit risk). ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A24: Baseline Refinancing Probit Model

	Model 1	Model 2	Model 3	Model 4
Initial LTV	2.53E-3***	2.45E-3***	2.50E-3***	8.59E-3***
	(2.72E-4)	(2.75E-4)	(2.73E-4)	(1.81E-4)
FICO	1.24E-4***	-1.13E-3***	1.25E-4***	1.29E-4***
	(2.53E-5)	(5.16E-5)	(2.52E-5)	(2.57E-5)
DTI	-5.56E-4***	-5.12E-4***	-5.95E-4***	-8.42E-4***
	(1.01E-4)	(1.01E-4)	(1.01E-4)	(1.01E-4)
Loan Age	1.44E-2***	1.45E-2***	1.26E-2***	1.25E-2***
	(2.97E-4)	(2.98E-4)	(3.01E-4)	(2.91E-4)
$\widehat{\textbf{Equity}}$	7.66E-8**	7.67E-8**	7.64E-8**	6.61E-8**
	(3.05E-8)	(3.07E-8)	(3.05E-8)	(2.58E-8)
Δ Mtg. Rt.	-0.1448***	-0.1445***	-0.1394***	-0.1519***
	(2.14E-3)	(2.14E-3)	(2.13E-3)	(2.11E-3)
Investor	-0.3177***	-0.3196***	-0.3182***	-0.2835***
	(6.28E-3)	(6.31E-3)	(6.30E-3)	(6.32E-3)
TPO	0.0678***	0.0678***	0.0682***	0.0600***
	(2.48E-3)	(2.48E-3)	(2.48E-3)	(2.47E-3)
Purchase	0.0539***	0.0546***	0.0492***	0.0868***
	(4.28E-3)	(4.29E-3)	(4.28E-3)	(4.34E-3)
Cash-Out Refi.	0.0258***	0.0259***	0.0269***	0.0274***
	(2.62E-3)	(2.63E-3)	(2.62E-3)	(2.65E-3)
$\widehat{\textbf{Unemp. Rate}}$	-0.0422***	-	-	-
	(1.52E-3)	-	-	-
$\widehat{\textbf{Unemp.}}$ **x Bin 1**	-	-0.0672***	-0.0388***	-0.0482***
	-	(2.08E-3)	(1.56E-3)	(1.59E-3)
$\widehat{\textbf{Unemp.}}$ **x Bin 2**	-	-0.0571***	-0.0346***	-0.0195***
	-	(1.67E-3)	(1.57E-3)	(1.59E-3)
$\widehat{\textbf{Unemp.}}$ **x Bin 3**	-	-0.0492***	-0.0492***	-0.0007
	-	(1.55E-3)	(1.64E-3)	(1.61E-3)
$\widehat{\textbf{Unemp.}}$ **x Bin 4**	-	-0.0359***	-0.0183***	0.0077***
	-	(1.53E-3)	(1.68E-3)	(1.66E-3)
$\widehat{\textbf{Unemp.}}$ **x Bin 5**	-	-0.0308***	-0.0208***	0.0154***
	-	(1.61E-3)	(1.92E-3)	(1.87E-3)
Year FE?	Yes	Yes	Yes	Yes
County FE?	Yes	Yes	Yes	Yes
Bin Type	None	FICO	HP App.	Equity
N	6,911,395	6,911,395	6,911,395	6,911,395
Pseudo-R^2	0.0609	0.0617	0.0632	0.0689

Notes: Dependent variable is indicator for refinancing, sample is all loans in matched dataset. Standard errors clustered at individual-level in parentheses. Hats denote instrumented variables. Interaction terms use equally-sized quintiles of the corresponding variable (higher bins correspond to lower credit risk). ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A25: Alternate Refinance Probit Model Without Instruments

	Model 1	Model 2	Model 3	Model 4
Initial LTV	2.54E-3***	2.44E-3***	2.29E-3***	7.97E-3***
	(2.94E-4)	(2.96E-4)	(2.78E-4)	(1.73E-4)
FICO	1.21E-4***	-1.27E-3***	1.38E-4***	1.24E-4***
	(2.52E-5)	(5.08E-5)	(2.56E-5)	(2.58E-5)
DTI	-5.55E-4***	-5.23E-4***	-4.43E-4***	-7.48E-4***
	(1.01E-4)	(1.01E-4)	(1.02E-4)	(1.01E-4)
Loan Age	1.20E-2***	1.20E-2***	9.91E-3***	1.13E-2***
	(2.94E-4)	(2.95E-4)	(2.94E-4)	(2.86E-4)
Equity	7.73E-8**	7.70E-8**	7.25E-8**	6.00E-8**
	(3.25E-8)	(3.26E-8)	(3.03E-8)	(2.38E-8)
Δ Mtg. Rt.	-0.1448***	-0.1449***	-0.1294***	-0.1491***
	(2.14E-3)	(2.14E-3)	(2.14E-3)	(2.11E-3)
Investor	-0.3176***	-0.3198***	-0.3204***	-0.2860***
	(6.32E-3)	(6.35E-3)	(6.35E-3)	(6.31E-3)
TPO	0.0673***	0.0672***	0.0624***	0.0606***
	(2.47E-3)	(2.48E-3)	(2.50E-3)	(2.47E-3)
Purchase	0.0554***	0.0560***	0.0412***	0.0856***
	(4.29E-3)	(4.30E-3)	(4.35E-3)	(4.37E-3)
Cash-Out Refi.	0.0262***	0.0261***	0.0216***	0.0269***
	(2.62E-3)	(2.63E-3)	(2.66E-3)	(2.64E-3)
Unemp. Rate	-0.0619***	-	-	-
	(1.47E-3)	-	-	-
Unemp. x Bin 1	-	-0.0902***	-0.0523***	-0.0636***
	-	(2.03E-3)	(1.47E-3)	(1.43E-3)
Unemp. x Bin 2	-	-0.0777***	-0.0496***	-0.0369***
	-	(1.61E-3)	(1.49E-3)	(1.44E-3)
Unemp. x Bin 3	-	-0.0684***	-0.0571***	-0.0179***
	-	(1.50E-3)	(1.53E-3)	(1.45E-3)
Unemp. x Bin 4	-	-0.0540***	-0.0295***	-0.0117***
	-	(1.48E-3)	(1.54E-3)	(1.48E-3)
Unemp. x Bin 5	-	-0.0494***	-0.0099***	-0.0023
	-	(1.59E-3)	(1.75E-3)	(1.59E-3)
Year FE?	Yes	Yes	Yes	Yes
Orig. Year FE?	Yes	Yes	Yes	Yes
County FE?	Yes	Yes	Yes	Yes
Bin Type	None	FICO	HP App.	Equity
N	6,911,395	6,911,395	6,911,395	6,911,395
Pseudo-R^2	0.0622	0.0631	0.0655	0.0699

Notes: Dependent variable is indicator for refinancing, sample is all loans in matched dataset. Standard errors clustered at individual-level in parentheses. Interaction terms use equally-sized deciles of the corresponding variable (higher bins correspond to lower credit risk). ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A26: Alternate Refinance Probit Model with Decile Bins

	Model 1	Model 2	Model 3	Model 4
$\widehat{\text{Unemp. Rate}}$	-0.0422***	-	-	-
	(1.52E-3)	-	-	-
$\widehat{\text{Unemp.}}$ x Bin 1	-	-0.0874***	-0.0482***	-0.0535***
	-	(2.28E-3)	(1.63E-3)	(1.63E-3)
$\widehat{\text{Unemp.}}$ x Bin 2	-	-0.0805***	-0.0270***	-0.0364***
	-	(2.08E-3)	(1.60E-3)	(1.64E-3)
$\widehat{\text{Unemp.}}$ x Bin 3	-	-0.0682***	-0.0254***	-0.0218***
	-	(1.90E-3)	(1.60E-3)	(1.64E-3)
$\widehat{\text{Unemp.}}$ x Bin 4	-	-0.0625***	-0.0346***	-0.0083***
	-	(1.76E-3)	(1.66E-3)	(1.63E-3)
$\widehat{\text{Unemp.}}$ x Bin 5	-	-0.0580***	-0.0644***	0.0018
	-	(1.67E-3)	(1.79E-3)	(1.64E-3)
$\widehat{\text{Unemp.}}$ x Bin 6	-	-0.0529***	-0.0288***	0.0087***
	-	(1.62E-3)	(1.70E-3)	(1.67E-3)
$\widehat{\text{Unemp.}}$ x Bin 7	-	-0.0490***	-0.0185***	0.0138***
	-	(1.58E-3)	(1.75E-3)	(1.69E-3)
$\widehat{\text{Unemp.}}$ x Bin 8	-	-0.0409***	0.0036**	0.0165***
	-	(1.56E-3)	(1.83E-3)	(1.74E-3)
$\widehat{\text{Unemp.}}$ x Bin 9	-	-0.0332***	0.0001	0.0212***
	-	(1.55E-3)	(1.99E-3)	(1.81E-3)
$\widehat{\text{Unemp.}}$ x Bin 10	-	-0.0255***	-0.0105***	0.0299***
	-	(1.56E-3)	(2.35E-3)	(2.16E-3)
Year FE?	Yes	Yes	Yes	Yes
Orig. Year FE?	Yes	Yes	Yes	Yes
County FE?	Yes	Yes	Yes	Yes
Bin Type	None	FICO	HP App.	Equity
N	6,911,395	6,911,395	6,911,395	6,911,395
Pseudo-R^2	0.0609	0.0667	0.0700	0.0623

Notes: Dependent variable is indicator for refinancing, sample is all loans in matched dataset. Standard errors clustered at individual-level in parentheses. Additional controls include LTV, FICO, DTI, loan age, instrumented home equity, change in mortgage rates, and indicators for investor-owned properties, third-party origination, purchase loans, and cash-out refinancing loans. Hats denote instrumented variables. Interaction terms use equally-sized deciles of the corresponding variable (higher bins correspond to lower credit risk). ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A27: Refinance Behavior of HARP-Eligible High-LTV Borrowers

	Model 1	Model 2	Model 3	Model 4
FICO	2.04E-5***	2.09E-5***	1.27E-5***	1.42E-5***
	(1.72E-6)	(1.70E-6)	(1.93E-6)	(1.93E-6)
DTI	-2.63E-4***	-2.57E-4***	-3.38E-4***	-3.33E-4***
	(6.13E-6)	(6.11E-6)	(7.05E-6)	(7.05E-6)
Initial LTV	6.23E-3***	6.25E-3***	6.87E-3***	8.85E-3***
	(1.19E-5)	(1.18E-5)	(1.35E-5)	(3.85E-5)
Int. Rt. Spread	2.54E-2***	2.58E-2***	2.65E-2***	2.72E-2***
	(1.87E-4)	(1.85E-4)	(2.17E-4)	(2.24E-4)
Loan Age	-5.41E-4***	-6.60E-4***	-7.94E-4***	2.59E-4*
	(1.09E-5)	(1.06E-5)	(1.49E-5)	(1.47E-4)
Δ Mtg. Rt.	-1.05E-2***	-1.07E-2***	-1.00E-2***	-
	(1.87E-4)	(1.86E-4)	(2.03E-4)	-
Equity	2.87E-6***	2.87E-6***	2.65E-6***	2.66E-6***
	(2.50E-9)	(2.49E-9)	(2.63E-9)	(2.62E-9)
Unemp. Rate	1.23E-2***	-	1.11E-2***	1.16E-2***
	(7.72E-5)	-	(1.65E-4)	(1.64E-4)
Dec. 2 Unemp.	-	5.75E-3***	-	-
	-	(1.69E-4)	-	-
Dec. 3 Unemp.	-	1.34E-2***	-	-
	-	(2.04E-4)	-	-
Dec. 4 Unemp.	-	2.11E-2***	-	-
	-	(2.29E-4)	-	-
Dec. 5 Unemp.	-	2.33E-2***	-	-
	-	(2.50E-4)	-	-
Dec. 6 Unemp.	-	2.86E-2***	-	-
	-	(2.63E-4)	-	-
Dec. 7 Unemp.	-	3.14E-2***	-	-
	-	(2.84E-4)	-	-
Dec. 8 Unemp.	-	3.82E-2***	-	-
	-	(3.19E-4)	-	-
Dec. 9 Unemp.	-	5.50E-2***	-	-
	-	(3.94E-4)	-	-
Dec. 10 Unemp.	-	1.20E-1***	-	-
	-	(4.67E-4)	-	-
Lender FE?	Yes	Yes	Yes	Yes
State FE?	Yes	Yes	Yes	Yes
Instruments?	None	None	Unemp./Equity	Unemp./Equity/App.
N	12,191,940	12,191,940	12,191,940	12,191,940
R^2	0.5981	0.5978	0.5982	0.5992

Notes: Dependent variable is indicator for prepayment, sample is all loans in GSE dataset originated in 12 months prior to April 2009 with LTV ratios above 75% observed from January 2012 to December 2013. Standard errors clustered at individual-level in parentheses. Additional controls include include initial loan interest rate and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. Models 2 and 4 use instrumented values for unemployment rates and home equity, and Model 4 also includes controls for combined LTV, age squared, instrumented home price appreciation, and quartiles of mortgage-rate changes as indicators and interacted with changes in mortgage rates. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A28: Refinance Behavior of HARP-Eligible Borrowers Post-2013

	Model 1	Model 2	Model 3	Model 4
FICO	9.44E-5***	9.42E-5***	9.35E-5***	9.34E-5***
	(2.92E-6)	(2.94E-6)	(2.92E-6)	(3.13E-6)
DTI	-2.31E-4***	-2.30E-4***	-2.49E-4***	-2.38E-4***
	(1.88E-5)	(1.86E-5)	(1.99E-5)	(1.85E-5)
Initial LTV	2.62E-3***	2.62E-3***	2.71E-3***	3.68E-3***
	(3.11E-4)	(3.12E-4)	(3.17E-4)	(4.92E-4)
Int. Rt. Spread	5.68E-4***	5.66E-4***	1.41E-4	8.95E-4***
	(1.94E-4)	(1.93E-4)	(1.89E-4)	(1.90E-4)
Loan Age	1.15E-3***	1.15E-3***	1.10E-3***	6.61E-3***
	(3.32E-5)	(3.46E-5)	(6.26E-5)	(2.57E-4)
Δ Mtg. Rt.	-7.13E-3***	-7.13E-3***	-5.90E-3***	-
	(2.26E-4)	(2.25E-4)	(2.51E-4)	-
Equity	4.60E-7***	4.61E-7***	4.82E-7***	4.85E-7***
	(5.81E-8)	(5.84E-8)	(5.99E-8)	(6.08E-8)
Unemp. Rate	4.44E-3***	-	9.00E-3***	1.19E-2***
	(7.42E-4)	-	(7.48E-4)	(7.39E-4)
Dec. 2 Unemp.	-	1.83E-3***	-	-
	-	(4.90E-4)	-	-
Dec. 3 Unemp.	-	4.83E-3***	-	-
	-	(1.07E-3)	-	-
Dec. 4 Unemp.	-	7.59E-3***	-	-
	-	(1.50E-3)	-	-
Dec. 5 Unemp.	-	1.31E-2***	-	-
	-	(2.13E-3)	-	-
Dec. 6 Unemp.	-	1.20E-2***	-	-
	-	(2.08E-3)	-	-
Dec. 7 Unemp.	-	1.60E-2***	-	-
	-	(2.79E-3)	-	-
Dec. 8 Unemp.	-	1.85E-2***	-	-
	-	(2.98E-3)	-	-
Dec. 9 Unemp.	-	3.24E-2***	-	-
	-	(5.43E-3)	-	-
Dec. 10 Unemp.	-	3.94E-2***	-	-
	-	(6.73E-3)	-	-
Lender FE?	Yes	Yes	Yes	Yes
State FE?	Yes	Yes	Yes	Yes
Instruments?	None	None	Unemp./Equity	Unemp./Equity/App.
N	6,283,568	6,283,568	6,283,568	6,283,568
R^2	0.1064	0.1066	0.1111	0.1118

Notes: Dependent variable is indicator for prepayment, sample is all loans in GSE dataset originated in 12 months prior to April 2009 with LTV ratios above 75% observed from January 2013 to December 2013. Standard errors clustered at individual-level in parentheses. Additional controls include include initial loan interest rate and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. Models 2 and 4 use instrumented values for unemployment rates and home equity, and Model 4 also includes controls for combined LTV, age squared, instrumented home price appreciation, and quartiles of mortgage-rate changes as indicators and interacted with changes in mortgage rates. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Appendix G: Alternate Difference-in-Differences Model Specifications

Table A29: Differences in Credit Risk Characteristics by Eligibility

	Ineligible	Eligible	t-Stat.
FICO	761.1	758.1	63.9
LTV	67.4	68.1	38.7
DTI	33.2	33.5	19.6
Int. Rt.	5.02	5.18	280.0
FTHB	10.3%	7.4%	91.0
SF	91.4%	92.2%	23.5
Own. Occ.	90.1%	91.5%	41.9
Purch.	31.0%	22.7%	165.0
Retail	52.8%	49.8%	52.6
# Obs.	1,569,630	1,493,292	-

Notes: Summary statistics for subsets of mortgage loans in HARP dataset at origination. "Eligible" and "Ineligible" refer to loans originated in 6-month window before and after April 2009. "t-Stat" denotes the t-statistics for the difference between the eligible and ineligible subsets.

Table A30: Refinancing Probability After X Years by Origination Date

Orig. Month	Refi. 2-Year	Refi. 3-Year	Refi. 4-Year
2008-10	36.66%	50.79%	65.17%
2008-11	35.99%	50.15%	66.36%
2008-12	26.27%	43.21%	64.83%
2009-01	19.13%	38.39%	62.97%
2009-02	17.05%	37.38%	62.12%
2009-03	14.73%	35.45%	59.55%
2009-04	12.26%	31.78%	54.61%
2009-05	12.19%	29.67%	49.97%
2009-06	15.51%	33.65%	51.81%
2009-07	18.70%	37.72%	53.63%
2009-08	20.80%	40.33%	54.79%
2009-09	20.88%	40.88%	53.95%

Notes: Probability that a loan originated in a given month has terminated via prepayment within specified number of years of origination.

Table A31: Unconditional Difference-in-Differences Using 1-Month HARP Eligibility Cutoff

Q1 Unemp.	Pre-HARP	Post-HARP	Difference	Diff.-in-Diff.
Ineligible	14.86%	27.17%	12.31%	-
Eligible	18.55%	34.57%	16.02%	3.71%
Q2 Unemp.	**Pre-HARP**	**Post-HARP**	**Difference**	**Diff.-in-Diff.**
Ineligible	11.69%	25.54%	13.84%	-
Eligible	15.37%	33.69%	18.32%	4.48%
Q3 Unemp.	**Pre-HARP**	**Post-HARP**	**Difference**	**Diff.-in-Diff.**
Ineligible	12.05%	23.73%	11.68%	-
Eligible	14.96%	31.52%	16.56%	4.88%
Q4 Unemp.	**Pre-HARP**	**Post-HARP**	**Difference**	**Diff.-in-Diff.**
Ineligible	10.78%	23.62%	12.84%	-
Eligible	12.68%	31.45%	18.77%	5.93%

Notes: Average CPR for subsets of mortgage loans in HARP dataset. “Eligible” and “Ineligible” refer to loans originated in 6-month window before and after June 2009, “Pre-HARP” and “Post-HARP” refer to monthly observations before and after January 2012. “Difference” denotes difference in CPR between loans observed during pre- and post-HARP period, “Diff.-in-Diff.” denotes difference in differences between eligible and ineligible loans. “Q1 Unemp.” through “Q2 Unemp.” refer to quartiles of the unemployment rate distribution, in order of increasing unemployment rates.

Table A32: Program Effect Estimates Using 1-Month HARP Eligibility Cutoff

	3-Month Window	6-Month Window	1-Year Window
Eligible	0.0069***	0.0053***	0.0035***
	(3.67E-5)	(6.45E-5)	(1.52E-4)
Post-HARP	0.0070***	0.0090***	0.0119***
	(4.74E-5)	(6.04E-5)	(5.06E-5)
Post-HARP 2.0	-0.0011***	-0.0024***	0.0011***
	(7.30E-5)	(9.65E-5)	(1.38E-4)
Post-HARP 2.0 x Eligible	0.0051***	0.0066***	0.0041***
	(5.37E-5)	(6.15E-5)	(7.73E-5)
FICO	5.06E-5***	5.38E-5***	5.34E-5***
	(6.82E-7)	(5.25E-7)	(5.61E-7)
DTI	-1.45E-4***	-1.63E-4***	-1.61E-4***
	(4.56E-6)	(5.11E-6)	(6.04E-6)
LTV	9.96E-4***	1.22E-3***	1.30E-3***
	(6.67E-5)	(7.51E-5)	(9.44E-5)
Int. Rt. Spread	5.26E-3***	7.09E-3***	8.81E-3***
	(1.01E-4)	(9.63E-5)	(1.11E-4)
Unemp.	-2.50E-4***	-1.81E-4***	-1.71E-4***
	(1.13E-5)	(8.85E-6)	(9.08E-6)
WALA	1.67E-4***	1.89E-4***	2.17E-4***
	(5.19E-6)	(5.02E-6)	(4.95E-6)
Δ Mtg. Rt.	-1.48E-2***	-1.64E-2***	-1.21E-2***
	(3.52E-5)	(6.73E-5)	(1.23E-4)
Equity	1.47E-7***	1.83E-7***	2.08E-7***
	(1.00E-8)	(1.06E-8)	(1.29E-8)
Lender FE?	Yes	Yes	Yes
State FE?	No	No	No
N	70,340,519	124,678,689	177,172,126
R^2	0.0624	0.0589	0.0488

Notes: Dependent variable is indicator for prepayment, sample is all loans in HARP dataset. Standard errors clustered at individual-level in parentheses. Additional controls include loan interest rate and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A33: Program Effect Estimates with Instruments

	3-Month Window	**6-Month Window**	**1-Year Window**
Eligible	0.0022***	0.0025***	0.0028***
	(4.36E-5)	(4.11E-5)	(4.17E-5)
Post-HARP	0.0048***	0.0093***	0.0126***
	(1.06E-4)	(1.01E-4)	(1.08E-4)
Post-HARP 2.0	-0.0010***	-0.0005***	0.0021***
	(2.01E-4)	(1.49E-4)	(1.15E-4)
Post-HARP 2.0 x Eligible	0.0077***	0.0065***	0.0038***
	(8.01E-5)	(6.82E-5)	(5.80E-5)
FICO	5.32E-5***	5.41E-5***	5.21E-5***
	(5.35E-7)	(5.65E-7)	(6.36E-7)
DTI	-1.62E-4***	-1.64E-4***	-1.59E-4***
	(5.98E-6)	(4.78E-6)	(3.95E-6)
LTV	1.16E-3***	1.26E-3***	1.31E-3***
	(9.48E-5)	(7.60E-5)	(6.42E-5)
Int. Rt. Spread	4.24E-4***	5.91E-4***	1.90E-3***
	(4.31E-5)	(3.85E-5)	(4.29E-5)
$\widehat{\textbf{Unemp.}}$	-1.12E-3***	-1.27E-3***	-1.40E-3***
	(1.15E-4)	(9.34E-5)	(7.85E-5)
WALA	3.81E-4***	4.03E-4***	2.89E-4***
	(5.22E-6)	(5.37E-6)	(3.40E-6)
Δ Mtg. Rt.	-1.10E-2***	-8.99E-3***	-8.99E-3***
	(6.22E-5)	(4.82E-5)	(3.48E-5)
$\widehat{\textbf{Equity}}$	1.73E-7***	1.92E-7***	2.12E-7***
	(1.34E-8)	(1.11E-8)	(9.85E-9)
Lender FE?	Yes	Yes	Yes
State FE?	No	No	No
N	82,998,950	120,265,979	175,338,982
$\mathbf{R^2}$	0.0562	0.0595	0.0626

Notes: Dependent variable is indicator for prepayment, sample is all loans in HARP dataset. Standard errors clustered at individual-level in parentheses. Additional controls include loan interest rate and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. Hats denote instrumented variables. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A34: Heterogeneity in Program Effect by Unemployment Rate Using 1-Month HARP Eligibility Cutoff

	3-Month Window	6-Month Window	1-Year Window
Eligible	0.0068***	0.0049***	0.0022***
	(2.78E-4)	(2.03E-4)	(1.72E-4)
Post-HARP	0.0070***	0.0090***	0.0118***
	(4.34E-5)	(5.75E-5)	(4.96E-5)
Post-HARP 2.0	-0.0004	-0.0022***	-0.0006
	(3.89E-4)	(4.24E-4)	(5.20E-4)
Post-HARP 2.0 x Eligible	0.0025***	0.0034***	0.0026***
	(2.49E-4)	(2.66E-4)	(3.31E-4)
Eligible x Unemp.	1.805E-5	4.566E-5*	1.390E-4***
	(2.99E-5)	(2.55E-5)	(2.25E-5)
Post-HARP 2.0 x Unemp.	-9.320E-5*	-3.665E-5	1.900E-4***
	(3.84E-5)	(4.11E-5)	(4.94E-5)
Post-HARP 2.0 x Eligible x Unemp.	3.407E-4***	4.199E-4***	2.214E-4***
	(2.74E-5)	(3.14E-5)	(3.97E-5)
FICO	5.06E-5***	5.39E-5***	5.35E-5***
	(6.80E-7)	(5.22E-7)	(5.59E-7)
DTI	-1.45E-4***	-1.63E-4***	-1.61E-4***
	(4.55E-6)	(5.11E-6)	(6.04E-6)
LTV	9.96E-4***	1.22E-3***	1.30E-3***
	(6.67E-5)	(7.52E-5)	(9.44E-5)
Int. Rt. Spread	5.64E-3***	4.46E-3***	-8.50E-4***
	(1.29E-4)	(3.15E-4)	(4.70E-4)
Unemp.	-2.84E-4***	-2.58E-4***	-3.28E-4***
	(1.74E-5)	(2.02E-5)	(2.42E-5)
WALA	1.67E-4***	1.90E-4***	2.20E-4***
	(5.43E-6)	(5.32E-6)	(5.35E-6)
Δ Mtg. Rt.	-1.48E-2***	-1.63E-2***	-1.21E-2***
	(3.34E-5)	(6.07E-5)	(1.12E-4)
Equity	1.47E-7***	1.83E-7***	2.08E-7***
	(1.00E-8)	(1.06E-8)	(1.29E-8)
Lender FE?	Yes	Yes	Yes
State FE?	No	No	No
N	70,340,519	124,678,689	177,172,126
R^2	0.0624	0.0589	0.0488

Notes: Dependent variable is indicator for prepayment, sample is all loans in HARP dataset. Standard errors clustered at individual-level in parentheses. Additional controls include loan interest rate and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A35: Heterogeneity in Program Effect by Unemployment Rate with Instruments

	3-Month Window	6-Month Window	1-Year Window
Eligible	0.0000	-0.0012***	-0.0010**
	(3.26E-4)	(3.59E-4)	(4.13E-4)
Post-HARP	0.0047***	0.0092***	0.0125***
	(1.00E-4)	(9.45E-5)	(9.80E-5)
Post-HARP 2.0	-0.0083***	-0.0075***	-0.0027***
	(6.88E-4)	(5.53E-4)	(4.45E-4)
Post-HARP 2.0 x Eligible	0.0051***	0.0054***	0.0025***
	(4.38E-4)	(3.75E-4)	(3.25E-4)
Eligible x $\widehat{\text{Unemp.}}$	2.37E-4***	3.96E-4***	4.04E-4***
	(3.56E-5)	(3.87E-5)	(4.39E-5)
Post-HARP 2.0 x $\widehat{\text{Unemp.}}$	8.60E-4***	8.08E-4***	5.54E-4***
	(6.12E-5)	(4.98E-5)	(3.97E-5)
Post-HARP 2.0 x Eligible x $\widehat{\text{Unemp.}}$	3.73E-4***	2.08E-4***	2.31E-4***
	(5.38E-5)	(4.57E-5)	(3.85E-5)
FICO	5.32E-5***	5.42E-5***	5.21E-5***
	(5.35E-7)	(5.63E-7)	(6.35E-7)
DTI	-1.62E-4***	-1.64E-4***	-1.59E-4***
	(5.99E-6)	(4.78E-6)	(3.94E-6)
LTV	1.16E-3***	1.26E-3***	1.31E-3***
	(9.49E-5)	(7.61E-5)	(6.43E-5)
Int. Rt. Spread	4.12E-4***	5.78E-4***	1.89E-3***
	(4.29E-5)	(3.83E-5)	(4.24E-5)
$\widehat{\text{Unemp.}}$	-1.48E-3***	-1.67E-3***	-1.73E-3***
	(1.40E-4)	(1.19E-4)	(1.04E-4)
WALA	3.92E-4***	4.11E-4***	2.95E-4***
	(5.73E-6)	(5.74E-6)	(3.65E-6)
Δ Mtg. Rt.	-1.08E-2***	-8.85E-3***	-8.89E-3***
	(5.88E-5)	(5.14E-5)	(3.53E-5)
$\widehat{\text{Equity}}$	1.73E-7***	1.92E-7***	2.12E-7***
	(1.34E-8)	(1.11E-8)	(9.86E-9)
Lender FE?	Yes	Yes	Yes
State FE?	No	No	No
N	82,998,950	120,265,979	175,338,982
R^2	0.0562	0.0595	0.0626

Notes: Dependent variable is indicator for prepayment, sample is all loans in HARP dataset. Standard errors clustered at individual-level in parentheses. Additional controls include loan interest rate and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. Hats denote instrumented variables. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A36: Heterogeneity in Program Effect by Quartiles of Unemployment Rate

Panel A: Baseline			
	3-Month Window	**6-Month Window**	**1-Year Window**
Post-HARP 2.0 x Eligible	0.0060***	0.0051***	0.0025***
	(9.22E-5)	(1.08E-4)	(1.26E-4)
Post-HARP 2.0 x Eligible x Q2 Unemp.	2.027E-3***	2.051E-3***	2.148E-3***
	(1.36E-4)	(1.61E-4)	(1.92E-4)
Post-HARP 2.0 x Eligible x Q3 Unemp.	2.571E-3***	2.319E-3***	2.242E-3***
	(1.58E-4)	(1.87E-4)	(2.16E-4)
Post-HARP 2.0 x Eligible x Q4 Unemp.	2.926E-3***	2.028E-3***	2.218E-3***
	(2.08E-4)	(2.51E-4)	(3.03E-4)
Panel B: Using 1-Month HARP Eligibility Cutoff			
	3-Month Window	**6-Month Window**	**1-Year Window**
Post-HARP 2.0 x Eligible	0.0038***	0.0049***	0.0025***
	(8.51E-5)	(1.04E-4)	(1.38E-4)
Post-HARP 2.0 x Eligible x Q2 Unemp.	1.886E-3***	2.196E-3***	2.136E-3***
	(1.21E-4)	(1.43E-4)	(1.79E-4)
Post-HARP 2.0 x Eligible x Q3 Unemp.	2.022E-3***	2.823E-3***	2.499E-3***
	(1.57E-4)	(1.83E-4)	(2.32E-4)
Post-HARP 2.0 x Eligible x Q4 Unemp.	2.386E-3***	3.175E-3***	2.823E-3***
	(2.03E-4)	(2.43E-4)	(3.03E-4)
Panel C: With Instruments			
	3-Month Window	**6-Month Window**	**1-Year Window**
Post-HARP 2.0 x Eligible	0.0037***	0.0068***	0.0076***
	(9.47E-5)	(1.10E-4)	(1.29E-4)
Post-HARP 2.0 x Eligible x Q2 Unemp.	5.91E-4***	3.69E-4**	7.49E-4***
	(1.40E-4)	(1.66E-4)	(1.94E-4)
Post-HARP 2.0 x Eligible x Q3 Unemp.	1.09E-3***	8.69E-4***	9.87E-4***
	(1.51E-4)	(1.77E-4)	(2.08E-4)
Post-HARP 2.0 x Eligible x Q4 Unemp.	5.47E-4**	-1.37E-4	8.57E-4***
	(2.26E-4)	(2.69E-4)	(3.20E-4)

Notes: Dependent variable is indicator for prepayment, sample is all loans in HARP dataset. Standard errors clustered at individual-level in parentheses. Additional controls include FICO, DTI, LTV, unemployment rate, WALA, change in mortgage rates, equity, loan interest rate and interest rate spread, and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A37: Heterogeneity in Program Effect by Unemployment with LTV Controls Using Additional HARP Controls

	3-Month Window	6-Month Window	1-Year Window
Eligible x High LTV	0.0119***	0.0090***	0.0072***
	(9.36E-5)	(1.25E-4)	(1.72E-4)
Post-HARP	0.0148***	0.0202***	0.0146***
	(3.80E-4)	(4.40E-4)	(4.37E-4)
Post-HARP 2.0	-0.0033***	-0.0035***	-0.0019***
	(4.05E-4)	(4.99E-4)	(6.38E-4)
Post-HARP x Eligible x High LTV	-0.0117***	-0.0129***	-0.0131***
	(6.66E-4)	(6.98E-4)	(8.00E-4)
Post-HARP 2.0 x Eligible x High LTV	1.145E-2***	1.146E-2***	8.626E-3***
	(3.67E-4)	(4.31E-4)	(5.22E-4)
Eligible x High LTV x Unemp.	-3.361E-4***	-6.050E-4***	-6.724E-4***
	(3.42E-5)	(3.29E-5)	(3.48E-5)
Post-HARP x Unemp.	-6.610E-4***	-9.052E-4***	-2.583E-4***
	(4.64E-5)	(4.85E-5)	(4.79E-5)
Post-HARP 2.0 x Unemp.	4.379E-4***	5.143E-4***	4.209E-4***
	(3.94E-5)	(4.89E-5)	(6.10E-5)
Post-HARP x Eligible x High LTV x Unemp.	8.365E-4***	1.112E-3***	1.291E-3***
	(8.38E-5)	(8.29E-5)	(9.42E-5)
Post-HARP 2.0 x Eligible x High LTV x Unemp.	-7.872E-4***	-8.342E-4***	-6.925E-4***
	(4.23E-5)	(5.10E-5)	(6.29E-5)
FICO	5.41E-5***	5.51E-5***	5.49E-5***
	(6.40E-7)	(5.86E-7)	(5.26E-7)
DTI	-1.61E-4***	-1.62E-4***	-1.56E-4***
	(4.41E-6)	(5.36E-6)	(6.76E-6)
LTV	1.22E-3***	1.33E-3***	1.37E-3***
	(6.92E-5)	(8.26E-5)	(1.02E-4)
Int. Rt. Spread	7.19E-3***	1.09E-2***	9.52E-3***
	(1.18E-4)	(1.54E-4)	(1.43E-4)
Unemp.	3.92E-4***	6.07E-4***	-4.86E-5
	(3.71E-5)	(3.76E-5)	(3.46E-5)
WALA	1.82E-4***	2.64E-4***	2.24E-4***
	(5.26E-6)	(6.88E-6)	(6.56E-6)
Δ Mtg. Rt.	-1.82E-2***	-1.41E-2***	-1.32E-2***
	(5.16E-5)	(5.48E-5)	(1.22E-4)
Equity	1.70E-7***	1.89E-7***	2.08E-7***
	(1.01E-8)	(1.14E-8)	(1.37E-8)
Lender FE?	Yes	Yes	Yes
State FE?	No	No	No
N	83,824,948	121,448,185	177,172,126
R^2	0.0627	0.0596	0.0563

Notes: Dependent variable is indicator for prepayment, sample is all loans in HARP dataset. Standard errors clustered at individual-level in parentheses. Additional controls include loan interest rate and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A38: Heterogeneity in Program Effect by Unemployment with LTV Controls Using 1-Month HARP Eligibility Cutoff

	3-Month Window	6-Month Window	1-Year Window
Eligible x High LTV	0.0114***	0.0103***	0.0077***
	(9.08E-5)	(1.41E-4)	(2.46E-4)
Post-HARP	0.0078***	0.0097***	0.0124***
	(5.76E-5)	(6.87E-5)	(5.68E-5)
Post-HARP 2.0	0.0014***	0.0002	0.0017***
	(2.35E-4)	(3.14E-4)	(4.32E-4)
Post-HARP 2.0 x Eligible x High LTV	-0.0031***	-0.0025***	-0.0039***
	(3.88E-4)	(4.15E-4)	(5.14E-4)
Eligible x High LTV x Unemp.	-6.031E-4***	-6.962E-4***	-6.653E-4***
	(1.44E-5)	(1.66E-5)	(2.16E-5)
Post-HARP 2.0 x Unemp.	6.417E-5	8.133E-5***	1.383E-4***
	(2.40E-5)	(3.15E-5)	(4.14E-5)
Post-HARP 2.0 x Eligible x High LTV x Unemp.	2.027E-4***	3.734E-4***	4.518E-4***
	(4.10E-5)	(4.34E-5)	(5.45E-5)
FICO	5.05E-5***	5.39E-5***	5.42E-5***
	(6.66E-7)	(5.21E-7)	(5.61E-7)
DTI	-1.43E-4***	-1.59E-4***	-1.57E-4***
	(4.49E-6)	(5.03E-6)	(5.98E-6)
LTV	1.07E-3***	1.31E-3***	1.38E-3***
	(6.87E-5)	(7.76E-5)	(9.74E-5)
Int. Rt. Spread	5.12E-3***	6.86E-3***	8.63E-3***
	(9.79E-5)	(9.62E-5)	(1.11E-4)
Unemp.	-6.77E-5***	4.62E-5***	5.24E-5***
	(1.12E-5)	(8.49E-6)	(1.07E-5)
WALA	1.64E-4***	1.96E-4***	2.31E-4***
	(5.43E-6)	(5.22E-6)	(5.18E-6)
Δ Mtg. Rt.	-1.52E-2***	-1.68E-2***	-1.24E-2***
	(3.68E-5)	(6.87E-5)	(1.22E-4)
Equity	1.48E-7***	1.84E-7***	2.08E-7***
	(1.01E-8)	(1.06E-8)	(1.29E-8)
Lender FE?	Yes	Yes	Yes
State FE?	No	No	No
N	70,340,519	124,678,689	177,172,126
R^2	0.0626	0.0591	0.0490

Notes: Dependent variable is indicator for prepayment, sample is all loans in HARP dataset. Standard errors clustered at individual-level in parentheses. Additional controls include loan interest rate and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A39: Heterogeneity in Program Effect by Unemployment with LTV Controls with Instruments

	3-Month Window	6-Month Window	1-Year Window
Eligible x High LTV	0.0077***	0.0082***	0.0079***
	(1.70E-4)	(1.47E-4)	(1.23E-4)
Post-HARP	0.0056***	0.0100***	0.0131***
	(1.19E-4)	(1.13E-4)	(1.19E-4)
Post-HARP 2.0	-0.0062***	-0.0047***	-0.0002
	(4.88E-4)	(3.66E-4)	(2.55E-4)
Post-HARP 2.0 x Eligible x High LTV	-0.0016	-0.0041***	-0.0089***
	(1.21E-3)	(1.03E-3)	(9.67E-4)
Eligible x High LTV x $\widehat{\text{Unemp.}}$	-9.03E-4***	-9.94E-4***	-9.66E-4***
	(3.36E-5)	(2.85E-5)	(2.43E-5)
Post-HARP 2.0 x $\widehat{\text{Unemp.}}$	9.60E-4***	7.42E-4***	3.78E-4***
	(4.28E-5)	(3.27E-5)	(2.25E-5)
Post-HARP 2.0 x Eligible x High LTV x $\widehat{\text{Unemp.}}$	4.90E-4***	7.05E-4***	1.07E-3***
	(1.31E-4)	(1.13E-4)	(1.07E-4)
FICO	5.36E-5***	5.49E-5***	5.37E-5***
	(5.34E-7)	(5.48E-7)	(6.03E-7)
DTI	-1.59E-4***	-1.60E-4***	-1.54E-4***
	(5.88E-6)	(4.68E-6)	(3.85E-6)
LTV	1.25E-3***	1.36E-3***	1.40E-3***
	(9.93E-5)	(7.97E-5)	(6.70E-5)
Int. Rt. Spread	4.86E-4***	5.86E-4***	1.80E-3***
	(4.28E-5)	(3.79E-5)	(4.11E-5)
$\widehat{\text{Unemp.}}$	-1.11E-3***	-1.18E-3***	-1.23E-3***
	(1.18E-4)	(9.45E-5)	(7.80E-5)
WALA	4.03E-4***	4.24E-4***	3.11E-4***
	(5.90E-6)	(5.91E-6)	(3.92E-6)
Δ Mtg. Rt.	-1.09E-2***	-9.08E-3***	-9.18E-3***
	(6.10E-5)	(4.73E-5)	(3.53E-5)
$\widehat{\text{Equity}}$	1.74E-7***	1.93E-7***	2.13E-7***
	(1.35E-8)	(1.12E-8)	(9.93E-9)
Lender FE?	Yes	Yes	Yes
State FE?	No	No	No
N	82,998,950	120,265,979	175,338,982
R^2	0.0564	0.0598	0.0629

Notes: Dependent variable is indicator for prepayment, sample is all loans in HARP dataset. Standard errors clustered at individual-level in parentheses. Additional controls include loan interest rate and indicators for first-time home buyers, property type, occupancy status, loan purpose, Fannie Mae loans, and origination channel. Hats denote instrumented variables. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Appendix H: Details on Structural Model Estimation

First Stage

The first stage of the model is estimated with a combination of maximum likelihood and a selection-corrected regression. I observe four separate discrete outcomes in the data: loan applications that are denied, approved loan applications where the offer is rejected, and originated loans that are either held or securitized. Denote these outcomes, respectively, as $a_{i,t} \in \{1, 2, 3, 4\}$. I also observe the continuous interest rate $r_{i,t}$, but only for a sample of loans that were securitized ($a_{i,t} = 4$). Suppressing for clarity any time- and agent subscripts as well as any dependence on state $(X_{i,t}, Z_{i,t}, S_{i,j,t})$ or structural parameters $(\phi_0, \phi_1, \tau_L, \beta)$, the likelihood function is given by:

$$
\begin{aligned}
\text{Application Denied:} \quad & \mathcal{L}(a = 1) = 1 - P_{FUND}(c) \\
\text{Offer Rejected:} \quad & \mathcal{L}(a = 2) = P_{FUND}(c) \int_{e=-\infty}^{\infty} P_{REJ}(c, r^*(c) + e) \, d\Phi(e) \\
\text{Originated, Held:} \quad & \mathcal{L}(a = 3) = P_{FUND}(c) \int_{e=-\infty}^{\infty} [1 - P_{REJ}(c, r^*(c) + e)][1 - P_{SEC}(c, r^*(c) + e)] \, d\Phi(e) \\
\text{Originated, Securitized:} \quad & \mathcal{L}(a = 4) = P_{FUND}(c) \int_{e=-\infty}^{\infty} [1 - P_{REJ}(c, r^*(c) + e)] P_{SEC}(c, r^*(c) + e) \, d\Phi(e)
\end{aligned}
$$

where the probabilities P_{FUND}, P_{REJ} and P_{SEC} are as defined in Section 4.3. I use the assumption that e is distributed Standard Normal to write the numerically-approximated log likelihood as:

$$\begin{aligned}
\log \mathcal{L}(a=1) &= \tau_L - \log(e^{\tau_L} + e^{\Sigma_m w_m[log(e^{\pi^0(r^*+n_m)}+e^{\pi^1(r^*+n_m)})+\gamma][1-P_{REJ}(r^*+n_m)]}) \\
\log \mathcal{L}(a=2) &= \Sigma_m w_m[log(e^{\pi^0(r^*+n_m)} + e^{\pi^1(r^*+n_m)}) + \gamma][1 - P_{REJ}(r^* + n_m)] \\
&- \log(e^{\tau_L} + e^{\Sigma_m w_m[log(e^{\pi^0(r^*+n_m)}+e^{\pi^1(r^*+n_m)})+\gamma][1-P_{REJ}(r^*+n_m)]}) \\
&+ \log(\Sigma_m w_m P_{REJ}(r^* + n_m)) \\
\log \mathcal{L}(a=3) &= \Sigma_m w_m[log(e^{\pi^0(r^*+n_m)} + e^{\pi^1(r^*+n_m)}) + \gamma][1 - P_{REJ}(r^* + n_m)] \\
&- \log(e^{\tau_L} + e^{\Sigma_m w_m[log(e^{\pi^0(r^*+n_m)}+e^{\pi^1(r^*+n_m)})+\gamma][1-P_{REJ}(r^*+n_m)]}) \\
&+ \log(\Sigma_m w_m[1 - P_{REJ}(r^* + n_m)](1 - P_{SEC}(r^* + n_m))) \\
\log \mathcal{L}(a=4) &= \Sigma_m w_m[log(e^{\pi^0(r^*+n_m)} + e^{\pi^1(r^*+n_m)}) + \gamma][1 - P_{REJ}(r^* + n_m)] \\
&- \log(e^{\tau_L} + e^{\Sigma_m w_m[log(e^{\pi^0(r^*+n_m)}+e^{\pi^1(r^*+n_m)})+\gamma][1-P_{REJ}(r^*+n_m)]}) \\
&+ \log(\Sigma_m w_m[1 - P_{REJ}(r^* + n_m)]P_{SEC}(r^* + n_m))
\end{aligned}$$

where (w_m, n_m) are a set of m Gauss-Hermite weights and nodes.

If the optimal interest rate r^* were observed, solving for the likelihood and gradient would be straightforward. Because r^* is not observed, I infer it as the solution to the lender's choice problem from Section 4.3. The first-order condition for $r^*(c)$ is given by:

$$\begin{aligned}
0 &= \int_{e=-\infty}^{\infty} [1 - P_{REJ}(c, r^*(c) + e)][\phi_r^0(1 - P_{SEC}(c, r^*(c) + e)) + \phi_r^1 P_{SEC}(c, r^*(c) + e) \\
&+ \beta_r P_{REJ}(c, r^*(c) + e)[\log(e^{\pi^0(c,r^*(c)+e)} + e^{\pi^1(c,r^*(c)+e)}) + \gamma]]\, d\Phi(e)
\end{aligned}$$

where β_r, ϕ_r^0 and ϕ_r^1 are the elasticities with respect to the interest rate of, respectively, the borrower's offer acceptance value and the lender's profit from holding and securitizing the loan.

At each guess of the structural parameters $(\phi_0, \phi_1, \tau_L, \beta)$, I first solve for r^* for each observation by iterating on a linearized first-order condition. I then project the price of the originated mortgage using the recovered latent interest rate and MBS pricing model from Table A40 of Appendix I, estimated in advance on MBS data. I then also calculate the gradient of optimal r^* with respect to the structural parameters under the assumption that the interest rate offer is linear[52] in the neighborhood of the optimum given the current-iterate values of the structural parameters. With these objects in hand I construct the likelihood and likelihood gradient[53], approximating any integrals using Gauss-Hermite quadrature with 8 quadrature nodes. I iterate to convergence using KNITRO, a nonlinear optimization software tool, with the overall convergence parameters set at 10^{-14} in parameter space and 10^{-8} for the objective function. I set the tolerance for solving the interest rate offer at 10^{-5} and cap the number of iterations on the linearized first-order condition at 25. Standard errors are calculated with a finite-differences approximation of the Hessian, using the analytical gradient, again approximated with Gauss-Hermite quadrature, and with a step size of 10^{-8} relative to the maximum likelihood parameter estimates.

Having thus recovered estimates for $(\hat{\phi}_0, \hat{\phi}_1, \hat{\tau}_L, \hat{\beta})$, the final task is to estimate α, the parameters governing the lender's interest rate offer. I assume that observed interest rate $r_{i,t} = r^*_{i,t} + e_{i,t}$, as in the model and seek to estimate a model $r^*_{i,t} = \Gamma(c_{i,t}, X_{i,t}, S_{i,j,t})\alpha + u_{i,t}$ for the optimal mean interest rate offer. Since I observe a sample only of loans for which $a_{i,t} = 4$, I face the sample-selection problem that in the regression:

$$\mathbb{E}[r_{i,t}|a_{i,t} = 4, c_{i,t}, X_{i,t}, S_{i,j,t}] \quad = \quad \Gamma(c_{i,t}, X_{i,t}, S_{i,j,t})\alpha + \mathbb{E}[e_{i,t}|a_{i,t} = 4, c_{i,t}, X_{i,t}, S_{i,j,t}]$$

the error term on the right is neither mean zero nor uncorrelated with $(X_{i,t}, S_{i,j,t})$. However, by applying Bayes' rule and the model-implied probability that a loan application results in securitization, I numerically approximate the conditional

[52]Literally, that the linear approximation to the interest rate first-order condition is true.
[53]While not shown, given the model's assumptions the gradient can be expressed analytically.

distribution of the error, as in Heckman (1979),[111] by:

$$
\begin{aligned}
f(e_{i,t}|a_{i,t}=4, c_{i,t}, X_{i,t}, S_{i,j,t}) &= \frac{f(e_{i,t})P(a_{i,t}=4|e_{i,t})}{P(a_{i,t}=4)} \\
&= \frac{\frac{1}{\sqrt{2\pi}} e^{\frac{-e_{i,t}^2}{2}} P_{FUND}(r_{i,t}^*)[1-P_{REJ}(r_{i,t}^*+e_{i,t})]P_{SEC}(r_{i,t}^*+e_{i,t})}{P_{FUND}(r_{i,t}^*)\Sigma_m w_m[1-P_{REJ}(r_{i,t}^*+n_m)]P_{SEC}(r_{i,t}^*+n_m)}
\end{aligned}
$$

and hence construct the following sample-selection correction term:

$$
\mathbb{E}[e_{i,t}|a_{i,t}=4, c_{i,t}, X_{i,t}, S_{i,j,t}] = \Sigma_m w_m \frac{P_{FUND}(r_{i,t}^*)[1-P_{REJ}(r_{i,t}^*+n_m)]P_{SEC}(r_{i,t}^*+n_m)}{P_{FUND}(r_{i,t}^*)\Sigma_{m'} w_{m'}[1-P_{REJ}(r_{i,t}^*+n_{m'})]P_{SEC}(r_i^*+n_{m'})} n_m
$$

I construct this term for each observed originated refinance mortgage at the estimated values of $(\hat{\phi}_0, \hat{\phi}_1, \hat{\tau}_L, \hat{\beta})$. I then regress the spread of the interest rate on those loans over the ten year swap rate and current guarantee fees, assuming that guarantee fees pass through to estimates one-for-one, as I estimate in Chapter 2, less the selection correction term, on a variety of borrower credit risk characteristics to estimate $\hat{\alpha}$.

Second Stage

The second stage of the model is estimated with filtered maximum likelihood. For each borrower-month in the data, I observe three discrete outcomes, either default, continuation, or refinance, denoted $a_{i,t} \in \{0,1,2\}$ respectively. Whenever $a_{i,t}=2$, I also observe the continuous LTV choice $c_{i,t} \in [0, \bar{c}]$. Three key pieces of information are hidden. First, I do not know whether an observed continuation or default followed from a failed refinance application or not, and as such must estimate these choice probabilities as a mixture. Second and relatedly, in cases where a borrower first applied for credit but subsequently continued or defaulted, I do not observe their LTV choice or offered interest rate, and hence must project these latent choices. Finally, I do not observe the borrowers state $Z_{i,t}$, only the average probability that $Z_{i,t}=1$ (the unemployment rate) in the borrower's area, and hence must filter out an estimate for $P(Z_{i,t}=1)$.

I specify that the realized-action-specific value function, $V(a_{i,t}, c_{i,t}, X_{i,t}, Z_{i,t}$, is a parametric function of the observed state and continuous choice with parameters θ that depend on the discrete action and latent state:

$$V(a_{i,t}, c_{i,t}, X_{i,t}, Z_{i,t}, \theta) = v(c_{i,t}, X_{i,t})\theta_{a,Z}$$

and that the latent-state transition probabilities are a logistic function of observed state $X_{i,t}$ and parameters ρ:

$$\begin{aligned} P(Z_{i,t} = 0 | Z_{i,t-1} = 0, \rho) &= \frac{e^{X_{i,t}\rho_0}}{1 + e^{X_{i,t}\rho_0}} \\ P(Z_{i,t} = 1 | Z_{i,t-1} = 1, \rho) &= \frac{e^{X_{i,t}\rho_1}}{1 + e^{X_{i,t}\rho_1}} \end{aligned}$$

This specification implies that unemployment is Markovian and that the probability of entering or exiting the unemployment state is time-varying and depends on, among other things, the current unemployment rate and average unemployment duration. My specification for the value function uses different set of parameters $\theta_{a,Z}$ for each of four (a, Z) combinations, normalizes $V(0, 0, X_{i,t}, Z_{i,t})$ to zero for all $(X_{i,t}, Z_{i,t})$, and allows the value function to depend on a cubic spline in four important state variables (principal balance, home price, income, and monthly payment) with ten gridpoints, plus other macroeconomic variables (such as interest rates) and home-value proxies (such as square feet) and measures of total cash extracted on refinance.

Under these assumptions, I can write the likelihood of each potential outcome as follows:

$$\text{Default: } \mathcal{L}(a=0) = \frac{1}{1+e^{v_{1,Z}}+e^{EV_{2,Z}-\tau_B}} + \frac{1}{1+e^{v_{1,Z}}} \frac{e^{EV_{2,Z}-\tau_B}}{1+e^{v_{1,Z}}+e^{EV_{2,Z}-\tau_B}}$$
$$\times \quad [1 - \int_{h=-\infty}^{\infty} P_{FUND}(c^*(e^h)) \int_{e=-\infty}^{\infty} [1 - P_{REJ}(c^*(e^h), r^*(c^*(e^h)) + e)] \, d\Phi(e) \, d\Phi(h)]$$
$$\text{Continue: } \mathcal{L}(a=1) = \frac{e^{v_{1,Z}}}{1+e^{v_{1,Z}}+e^{EV_{2,Z}-\tau_B}} + \frac{e^{v_{1,Z}}}{1+e^{v_{1,Z}}} \frac{e^{EV_{2,Z}-\tau_B}}{1+e^{v_{1,Z}}+e^{EV_{2,Z}-\tau_B}}$$
$$\times \quad [1 - \int_{h=-\infty}^{\infty} P_{FUND}(c^*(e^h)) \int_{e=-\infty}^{\infty} [1 - P_{REJ}(c^*(e^h), r^*(c^*(e^h)) + e)] \, d\Phi(e) \, d\Phi(h)]$$
$$\text{Refi. at } c^* < \bar{c}\text{: } \mathcal{L}(a=2, c^*) = \frac{e^{EV_{2,Z}-\tau_B}}{1+e^{v_{1,Z}}+e^{EV_{2,Z}-\tau_B}}$$
$$\times \quad [P_{FUND}(c^*) \int_{e=-\infty}^{\infty} [1 - P_{REJ}(c^*, r^*(c^*) + e)] \, d\Phi(e) \phi(\log(\eta(c^*)))]$$
$$\text{Refi. at } c^* = \bar{c}\text{: } \mathcal{L}(a=2, \bar{c}) = \frac{e^{EV_{2,Z}-\tau_B}}{1+e^{v_{1,Z}}+e^{EV_{2,Z}-\tau_B}}$$
$$\times \quad [P_{FUND}(\bar{c}) \int_{e=-\infty}^{\infty} [1 - P_{REJ}(\bar{c}, r^*(\bar{c}) + e)] \, d\Phi(e)][1 - \Phi(\log(\eta(\bar{c})))]$$

where $v_{1,Z}$ denotes the value of continuation in state Z, $\eta(c)$ the value of the continuous preference shock η required to rationalize the observed choice, and $EV_{2,Z}$ the integral over η of EV_R, and all other dependence on state ($X_{i,t}$, $Z_{i,t}$, parameters θ, and time- and agent subscripts is suppressed for clarity. Note that the likelihood of default and continuation includes a term that captures the probability that those choices are optimal as well as a term capturing the probability that refinancing was optimal, multiplied by the probability that the loan application was either denied or the interest rate offer was rejected, integrated over all potential LTV choices and interest rate offers. Then using the assumption that e is distributed Standard Normal and η Standard Lognormal, I can numerically approximate the log-likelihood as:

$$
\begin{aligned}
& \log \mathcal{L}(a=0) = -\log(1+e^{v_{1,Z}}+e^{EV_{2,Z}-\tau_B}) - \log(1+e^{v_{1,Z}}) \\
+ \quad & \log[1+e^{v_{1,Z}}+e^{EV_{2,Z}-\tau_B}(1-\Sigma_m w_m P_{FUND}(c(e^{n_m}))\Sigma_{m'} w_{m'}(1-P_{REJ}(c(e^{n_m}),r^*(c(e^{n_m}))+n_{m'})))] \\
& \log \mathcal{L}(a=1) = v_{1,Z} - \log(1+e^{v_{1,Z}}+e^{EV_{2,Z}-\tau_B}) - \log(1+e^{v_{1,Z}}) \\
+ \quad & \log[1+e^{v_{1,Z}}+e^{EV_{2,Z}-\tau_B}(1-\Sigma_m w_m P_{FUND}(c(e^{n_m}))\Sigma_{m'} w_{m'}(1-P_{REJ}(c(e^{n_m}),r^*(c(e^{n_m}))+n_{m'})))] \\
& \log \mathcal{L}(a=2,c^*) = EV_{2,Z} - \tau_B - \log(1+e^{v_{1,Z}}+e^{EV_{2,Z}-\tau_B}) \\
+ \quad & \log(P_{FUND}(c^*)\Sigma_m w_m[1-P_{REJ}(c^*,r^*(c^*)+n_m)]) + \log(\phi(\log(\eta(c^*)))) \\
& \log \mathcal{L}(a=2,\bar{c}) = EV_{2,Z} - \tau_B - \log(1+e^{v_{1,Z}}+e^{EV_{2,Z}-\tau_B}) \\
+ \quad & \log(P_{FUND}(\bar{c})\Sigma_m w_m[1-P_{REJ}(\bar{c},r^*(\bar{c})+n_m)]) + \log(1-\Phi(\log(\eta(\bar{c}))))
\end{aligned}
$$

where (w_m, n_m) and $(w_{m'}, n_{m'})$ are sets of Gauss-Hermite weights and nodes.

The set of structural parameters to be estimated is (θ, τ_B, ρ), and if I could observe latent state $Z_{i,t}$, I could write the full log-likelihood as a function of these parameters as[54]:

$$
\begin{aligned}
\log \mathcal{L}(a^T, c^T Z^T | \theta, \tau_B, \rho) \quad = \quad & \Sigma_{z\in\{0,1\}} \mathbf{1}[Z_1 = z]\mathcal{L}(a_1, c_1 | Z_1 = z, \theta, \tau_B) \\
+ \quad & \Sigma_{t=2}^{T} \Big\{ \Sigma_{z\in\{0,1\}} \mathbf{1}[Z_t = z]\mathcal{L}(a_t, c_t | Z_t = z, \theta, \tau_B) \\
+ \quad & \Sigma_{z\in\{0,1\}} \left[\mathbf{1}[Z_t = z, Z_{t-1} = 0] \log(P(Z_t = z | Z_{t-1} = 0, \rho_0))\right. \\
+ \quad & \left.\mathbf{1}[Z_t = z, Z_{t-1} = 1] \log(P(Z_t = z | Z_{t-1} = 1, \rho_1))\right] \Big\}
\end{aligned}
$$

where superscripts denote "history up to" and agent subscripts are omitted for clarity. However, since I do not observe $Z_{i,t}$, I must instead form an expectation of the probability distribution of $Z_{i,t}$ and then maximize the expected log likelihood.

Following Diebold et al (1999),[66] I iteratively filter out an estimate for the

[54]Note that if one were to write the full likelihood conditional on $Z_{i,t}$ and take the log, it would not be the same as the expression below. However, both the function written below and the true log-likelihood achieve their maximum at the same value of the structural parameters, as discussed in Baum et al (1970).[27]

distribution of $Z_{i,t}$ and then update the parameters (θ, τ_B, ρ) using an EM algorithm. In the E-step, conditional on the current guess of the structural parameters, I use a forward-backward algorithm to compute $\hat{P}(Z_t = z) = \mathbb{E}[\mathbf{1}[Z_{i,t} = z]]$ and $\hat{P}(Z_t = z, Z_{t-1} = w) = \mathbb{E}[\mathbf{1}[Z_t = z, Z_{t-1} = w]]$. These estimates are then substituted into the expression for the log-likelihood above to form the expected log-likelihood. In the M-step, conditional on these estimates for the latent state probabilities, I maximize the log-likelihood with respect to the structural parameters (θ, τ_B, ρ). Note that the log-likelihood is additively separable into a block dependent on (θ, τ_B) (the first two lines) and a block dependent on ρ (the last two lines). Maximizing the ρ-block is straightforward, and I do so using KNITRO optimization software.

If the optimal LTV choice $c_{i,t}$ and offered interest rate $r_{i,t}$ were observed for each unsuccessful refinance application, maximizing the θ-block would also be straightforward. However, because these variables are unobserved, I must project them out for each observation. At each iteration of the EM algorithm, I first use the estimated parameters $(\hat{\phi}_0, \hat{\phi}_1, \hat{\tau}_L, \hat{\beta}, \hat{\alpha})$ to project the probabilities P_{FUND} and P_{REJ}. I then iterate on the first-order condition the borrower's LTV choice problem:

$$
\begin{aligned}
-[\log(1+e^{v_{1,Z}})+\gamma] \quad &= \quad P_{FUND}(c,r^*)(1-P_{FUND}(c,r^*)) \\
&\times \quad \frac{\partial \int_{e=-\infty}^{\infty} [1-P_{REJ}(c,r^*+e)][\log(e^{\pi^0(c,r^*+e)}+e^{\pi^1(c,r^*+e)})+\gamma]\, d\Phi(e)}{\partial c} \\
&\times \quad \int_{e=-\infty}^{\infty} [1-P_{REJ}(c,r^*+e)][v_{2,Z}(c,r^*+e)\eta - \log(1+e^{v_{1,Z}})-\gamma]\, d\Phi(e) \\
&+ \quad P_{FUND}(c,r^*) \int_{e=-\infty}^{\infty} \Big\{ [1-P_{REJ}(c,r^*+e)][\frac{\partial v_{2,Z}(c,r^*+e)\eta}{\partial c} + (\beta_c + \beta_r \alpha_c) \\
&\times \quad (1-P_{REJ}(c,r^*+e))[v_{2,Z}(c,r^*+e)\eta - \log(1+e^{v_{1,Z}})-\gamma] \Big\}\, d\Phi(e)
\end{aligned}
$$

until the system converges to the optimal $c^*(\eta)$, where β_c and α_c are the elasticities with respect to LTV of, respectively, the borrower's offer acceptance value and the offered interest rate.

During each θ-block of the M-step, given the current-iterate guess of the structural parameters (θ, τ_B) and the current-iterate latent state distribution $\hat{P}(Z)$, I first solve for

$c^*(\eta)$ for each non-refinance observation by iterating on the linearized first-order condition. I solve the LTV choice problem in this fashion at 8 values of $\eta_m = e^{n_m}$, where n_m are the first 8 Gauss-Hermite nodes. I assume that the borrower's optimal $c^*(\eta)$ is capped at an exogenous upper limit $\bar{c}$, and in cases where $c^*(\eta) > \bar{c}$ I set it equal to $\bar{c}$. Because by construction the borrower's LTV choice problem satisfies the monotone choice assumption from Blevins (2010),[40] $\bar{c}$ will always be the borrower's most-preferred feasible alternative whenever $c^*(\eta) > \bar{c}$. For refinance observations, I solve the first-order equation analytically for η to obtain $\eta(c)$ for the observed LTV choice. As before I assume that $c^*(\eta)$ is linear in the neighbordhood of the optimum at the current-iterate values of the structural parameters and calculate the gradient of the optimal LTV choice with respect to parameters. I then use the estimates $(\hat{\phi}_0, \hat{\phi}_1, \hat{\tau}_L, \hat{\beta}, \hat{\alpha})$ to project r^*, P_{FUND} and P_{REJ} for each observation and construct the likelihood and likelihood gradient[55], as before approximating any integrals with Gauss-Hermite quadrature. I set the convergence criterion in parameter space to 10^{-5} for both the inner maximization problem (each M-step) and the outer maximization problem (the EM algorithm) and iterate to convergence.

Third Stage

The third and final stage in estimation involves recovering the action-specific utility functions and estimating a parametric form for them. The process to recover utility involves estimating a transition function for the observable states; simulating observable states according to that function, as well as unobservable states according to the latent-state distributions estimated in Stage 2; projecting a future discounted continuation value; and then subtracting that from the projected value. To implement this procedure, I first estimate place-by-place vector autoregressions (VARs) of the key

[55]As in Stage 1, the setup of the model allows the gradient to be computed analytically, although I for sake of space I do not present it here.

exogenous observable state variables, treating the VAR estimates as transition functions. I estimate single-lagged VARs in 10 state variables: house price indices, mean unemployment duration, unemployment rates, mortgage rates, 30-year swap rates, and 10-, 7-, 5-, 3- and 1-year treasury rates.[56] I use the standard identification (Cholesky decomposition), ordering the variables in the same order presented here.[57] I then draw a series of $s = 10$ errors from the estimate error covariance matrix, assuming that the VAR error is multivariate normal, and project a set of s simulated future states accordingly. I allow the borrowers' choices in the previous period to affect the evolution of the endogenous states (e.g. equity) in the future period. I also project the future latent state probabilities using the estimated latent state distribution and transition density from Stage 2.

Given the value function projections estimated in Stage 2, I then construct the expected future logit-inclusive value at each of these s simulated future states. I set the discount factor exogenously to .995, corresponding to a roughly 6% annual rate of discount[58]. I subtract the discounted average future expected value, which owing the transition of the latent and endogenous observable states is $Z_{i,t}$- and $(a_{i,t}, c_{i,t})$-specific, from the projected values for the observed states, which are also $Z_{i,t}$- and $(a_{i,t}, c_{i,t})$-specific. This generates an estimated utility as a function of $(a_{i,t}, c_{i,t}, X_{i,t}, Z_{i,t})$ for each observed state $X_{i,t}$, and each of four $(a_{i,t}, Z_{i,t})$ combinations. For the refinancing actions $a_{i,t} = 2$ at each observed state $X_{i,t}$ and each potential $Z_{i,t}$, I calculate the optimal LTV choice $c(\eta)$ for the mean value of $\eta = 1$ and the mean interest rate offer corresponding to this optimal LTV choice choice. Using the optimal LTV choice and mean interest rate, I calculate mean equity extracted on refinancing for each state $(X_{i,t}, Z_{i,t})$, as well as the monthly payment on the refinance loan. Note that all of these

[56]Error-correction-based panel cointegration tests strongly support the presence of a cointegrating relationship between these ten variables.

[57]The order was chosen to represent a plausible order of non-propagation in shocks, e.g. this ordering imposes that shocks to the shorter end of the term structure affect the later end, but not vice-versa.

[58]Specifying the discount factor at .99 did not substantively change the utility parameter estimates.

terms are calculated even for borrowers who are observed to either continue or default. These terms are included as regressors for refinancing utility, along additional variables intended to capture the flow value of housing services and the flow value of improved credit. I then regress the action-specific utility value $u_{a,Z}$ on these covariates in order to recover borrower preference parameters.

Appendix I: Additional Structural Estimates

Table A40: Estimated MBS Pricing Model

Variable	Est.	S.E.	Mean	S.D.
Coupon Spread $\times \mathbf{1}[\textbf{Spread} < -2]$	5.0564***	0.118	-0.001	0.039
Coupon Spread $\times \mathbf{1}[-2 < \textbf{Spread} < -1]$	4.8897***	0.023	-0.108	0.351
Coupon Spread $\times \mathbf{1}[-1 < \textbf{Spread} < -0]$	4.5139***	0.023	-0.322	0.307
Coupon Spread $\times \mathbf{1}[0 < \textbf{Spread} < 1]$	2.4250***	0.047	0.056	0.150
Coupon Spread $\times \mathbf{1}[1 < \textbf{Spread} < 2]$	2.0656***	0.070	0.012	0.121
Coupon Spread $\times \mathbf{1}[\textbf{Spread} > 2]$	2.5414***	0.375	0.001	0.034
Coupon $\times$ **1-Year Treasury**	0.0765***	0.005	11.694	11.048
Coupon $\times$ **3-Year Treasury**	0.1084***	0.021	13.276	9.682
Coupon $\times$ **5-Year Treasury**	-0.2607***	0.045	15.381	8.774
Coupon $\times$ **7-Year Treasury**	0.0704***	0.026	17.069	8.070
Coupon $\times$ **10-Year Swap**	-0.3752***	0.011	20.951	8.746
Coupon $\times$ **30-Year Swap**	0.2349***	0.008	23.016	8.217
G-Fee	-0.8963***	0.039	0.536	0.151
Avg. LTV	0.0268***	0.001	74.968	8.913
FICO (620)	0.5033***	0.068	0.003	0.055
FICO (640)	0.1958**	0.080	0.003	0.050
FICO (660)	1.5300	0.053	0.007	0.086
FICO (680)	0.8600	0.031	0.023	0.149
FICO (700)	-0.6500	0.024	0.076	0.265
FICO (720)	-1.2600	0.018	0.230	0.421
FICO (740)	0.9300	0.017	0.240	0.427
Constant	104.5745***	0.084	-	-
N	40,307	-	-	-
$\mathbf{R^2}$	0.9107	-	-	-
Dependent Variable	-	-	101.474	2.937

Notes: Estimated coefficients and standard errors from MBS pricing model used as structural model Stage 1 input. Variable means and standard deviations shown in columns 3 and 4 respectively. FICO scores refer to MBS mean values. "Coupon Spread" denotes spread over 10-year swap rate. ***/**/* indicates significance at 99%/95%/90%-level respectively. Estimated guarantee fee elasticity with respect to LTV is .59 bp per 100% LTV.

Table A41: Estimated Lender Policy Parameters

ϕ_0	Est.	S.E.	Mean	S.D.
Constant	4.6699***	0.017	1.000	0.000
LTV	-1.178***	0.009	0.625	0.123
Unemp.	-2.662***	0.031	0.074	0.029
FICO (620)	-1.370***	0.044	0.000	0.016
FICO (640)	-1.129***	0.034	0.000	0.021
FICO (660)	-1.111***	0.021	0.001	0.035
FICO (680)	-0.702***	0.009	0.006	0.078
FICO (700)	0.770***	0.003	0.064	0.245
FICO (720)	0.875***	0.002	0.359	0.480
FICO (740)	0.532***	0.002	0.259	0.438
Income	1.003***	0.063	0.008	0.010
Owner Occ.	-0.071***	0.003	0.910	0.286
Int. Rt. Spread	1.367**	0.662	0.016	0.004
ϕ_1	**Est.**	**S.E.**	**Mean**	**S.D.**
Constant	1.835***	0.067	1.000	0.000
Price	1.998***	0.064	1.020	0.020
τ_L	**Est.**	**S.E.**	**Mean**	**S.D.**
Constant	2.516***	0.010	1.000	0.000

Notes: Estimated coefficients and standard errors for parameters governing lender profits from Stage 1 structural estimates. Sample is all loans in HMDA dataset. "Int. Rt. Spread" is spread of latent interest rate over prevailing national mortgage rate. "Price" is estimated secondary-market mortgage price using pricing model from Table A44. Variable means and standard deviations shown in columns 3 and 4 respectively. ***/**/* indicates significance at 99%/95%/90%-level respectively. N = 17,109,796.

Table A42: Estimated Borrower Rejection Parameters

β	Est.	S.E.	Mean	S.D.
Constant	0.223***	0.011	1.000	0.000
Income	10.602***	0.087	0.008	0.010
LTV	1.127***	0.006	0.625	0.123
FICO (620)	1.456***	0.094	0.000	0.016
FICO (640)	1.427***	0.045	0.000	0.021
FICO (660)	1.200***	0.028	0.001	0.035
FICO (680)	0.360***	0.007	0.006	0.078
FICO (700)	-0.213***	0.002	0.064	0.245
FICO (720)	-0.280***	0.002	0.359	0.480
FICO (740)	-0.155***	0.002	0.259	0.438
Unemp.	0.288***	0.039	0.074	0.029
Int. Rt.	-2.725***	0.114	0.059	0.011
Int. Rt. Spread	-3.384***	0.127	-0.002	0.005

Notes: Estimated coefficients and standard errors for parameters governing borrower offer rejection policy from Stage 1 structural estimates. Sample is all loans in HMDA dataset. Additional controls include year- and county fixed effects. "Int. Rt. Spread" is spread of latent interest rate over prevailing national mortgage rate. Variable means and standard deviations shown in columns 3 and 4 respectively. ***/**/* indicates significance at 99%/95%/90%-level respectively. N = 17,109,796.

Table A43: Estimated Interest Rate Offer Curve

α	Est.	S.E.	Mean	S.D.
LTV	0.321***	0.002	0.607	0.170
FICO (620)	0.861***	0.005	0.032	0.177
FICO (640)	0.815***	0.005	0.033	0.178
FICO (660)	0.784***	0.005	0.049	0.215
FICO (680)	0.755***	0.005	0.066	0.248
FICO (700)	0.746***	0.005	0.081	0.273
FICO (720)	0.733***	0.005	0.096	0.295
FICO (740)	0.728***	0.005	0.105	0.306
FICO (850)	0.696***	0.005	0.538	0.499
Borrower Income	-1.51E-6***	8.30E-8	4778.4	3764.2
Owner Occ.	-0.313***	0.001	0.913	0.282
Unemp. Rate	0.064***	0.000	7.539	2.741
N	3,611,896	-	-	-
R^2	0.8389	-	-	-
Dependent Variable	-	-	1.230	0.578
Interest Rate	-	-	5.707	1.027
Int. Rt. Spread	-	-	1.444	0.599
Selection Correction Term	-	-	-0.026	0.018

Notes: Estimated coefficients and standard errors for parameters governing interest rate offer curve from Stage 1 structural estimates. Dependent variable is difference between spread of interest rate over 10-year swap rate and selection correction term, sample is all California refinance loans in GSE dataset. Additional controls include county fixed effects. Variable means and standard deviations shown in columns 3 and 4 respectively. ***/**/* indicates significance at 99%/95%/90%-level respectively. N = 3,611,896.

Table A44: Estimated Latent Transition Density Parameters

ρ	Est.	S.E.		Est.	S.E.
	ρ_0			ρ_1	
Unemp. × Q4 Inc.	46.2910***	0.4462		-37.7050***	0.6625
Mean Unemp. Dur. × Q4 Inc.	24.2090***	0.3429		-33.4930***	0.4844
ΔUnemp. × Q4 Inc.	-16.5500***	1.3201		3.4904**	1.8484
ΔMean Unemp. Dur. × Q4 Inc.	-43.2200***	1.6107		10.8350***	2.1981
Unemp.^{2} × Q4 Inc.	-187.4900***	1.6912		188.9400***	2.9259
Mean Unemp. Dur.^{2} × Q4 Inc.	-87.6690***	0.9782		105.1000***	1.3897
ΔUnemp.^{2} × Q4 Inc.	2.9542	93.9240		-2.5726	137.6700
Unemp. × Q3 Inc.	45.6820***	0.4553		-37.4980***	0.6721
Mean Unemp. Dur. × Q3 Inc.	25.3220***	0.3551		-34.2090***	0.5018
ΔUnemp. × Q3 Inc.	-18.4680***	1.3721		4.3859**	2.0240
ΔMean Unemp. Dur. × Q3 Inc.	-42.5660***	1.7130		9.3402***	2.4685
Unemp.^{2} × Q3 Inc.	-177.1500***	1.5446		180.5700***	3.0458
Mean Unemp. Dur.^{2} × Q3 Inc.	-93.8720***	0.9962		107.8700***	1.4515
ΔUnemp.^{2} × Q3 Inc.	4.9343	92.7530		-3.7695	140.2500
Unemp. × Q2 Inc.	47.0990***	0.4540		-38.7360***	0.6659
Mean Unemp. Dur. × Q2 Inc.	26.4850***	0.3566		-33.6390***	0.4988
ΔUnemp. × Q2 Inc.	-18.1970***	1.4260		3.7207**	2.0696
ΔMean Unemp. Dur. × Q2 Inc.	-44.1600***	1.7466		8.3949***	2.5142
Unemp.^{2} × Q2 Inc.	-186.9000***	1.5491		178.2400***	2.9263
Mean Unemp. Dur.^{2} × Q2 Inc.	-99.5460***	0.9953		108.0200***	1.4227
ΔUnemp.^{2} × Q2 Inc.	1.4233	95.3120		-4.4616	143.8900
Unemp. × Q1 Inc.	47.8670***	0.4549		-38.1110***	0.6634
Mean Unemp. Dur. × Q1 Inc.	25.5240***	0.3523		-33.8880***	0.4947
ΔUnemp. × Q1 Inc.	-19.3680***	1.4187		4.5482**	2.0721
ΔMean Unemp. Dur. × Q1 Inc.	-46.1540***	1.7413		10.6840***	2.5005
Unemp.^{2} × Q1 Inc.	-185.6200***	1.6115		169.1300***	3.0455
Mean Unemp. Dur.^{2} × Q1 Inc.	-98.1890***	0.9822		111.0700***	1.4130
ΔUnemp.^{2} × Q1 Inc.	-2.2251	92.9310		2.0918	145.6900

Notes: Estimated coefficients and standard errors from Stage 2 latent liquidity preference transition density. "Unemp." denotes local-area unemployment rate, "Mean Unemp. Dur." denotes county mean duration of unemployment, and income quartiles are defined within-county, i.e. Q4 income is lowest-earning 25% of sample in their county. ***/**/* indicates significance at 99%/95%/90%-level respectively.

Table A45: Credit Supply Model Fit by Borrower Type, Year, and Credit Risk

Panel A: Model Fit by Borrower Type

		P(Deny)	P(Reject)	P(Hold)	P(Securitize)
Overall	**Observed**	0.224	0.217	0.325	0.234
	Simulated	0.224	0.217	0.325	0.234
	Ratio	1.000	0.999	1.002	0.999
Low DTI	**Observed**	0.207	0.196	0.304	0.293
	Simulated	0.210	0.202	0.322	0.266
	Ratio	0.987	0.973	0.944	1.099
High DTI	**Observed**	0.241	0.238	0.346	0.176
	Simulated	0.238	0.233	0.327	0.203
	Ratio	1.012	1.021	1.058	0.869
Low LTV	**Observed**	0.211	0.221	0.334	0.235
	Simulated	0.224	0.220	0.331	0.225
	Ratio	0.941	1.001	1.009	1.044
High LTV	**Observed**	0.234	0.214	0.319	0.233
	Simulated	0.224	0.215	0.320	0.241
	Ratio	1.043	0.997	0.995	0.968
Low Balance	**Observed**	0.226	0.225	0.326	0.223
	Simulated	0.232	0.226	0.327	0.215
	Ratio	0.975	0.995	0.996	1.038
High Balance	**Observed**	0.221	0.207	0.324	0.247
	Simulated	0.215	0.207	0.322	0.257
	Ratio	1.032	1.004	1.007	0.961

Notes: Observed and estimated choice probabilities by applicant type for Stage 1 outcomes. First four columns correspond to observed choice probabilities, second four to simulated probabilities at maximum-likelihood estimates, and final four columns to ratio of observed-to-simulated.

Panel B: Model Fit by Year

	Observed				Simulated				Ratio			
	P(Deny)	P(Reject)	P(Hold)	P(Securitize)	P(Deny)	P(Reject)	P(Hold)	P(Securitize)	P(Deny)	P(Reject)	P(Hold)	P(Securitize)
2000	0.339	0.237	0.336	0.089	0.322	0.264	0.266	0.149	1.053	0.898	1.262	0.599
2001	0.216	0.215	0.364	0.204	0.210	0.228	0.363	0.199	1.029	0.946	1.003	1.026
2002	0.184	0.216	0.348	0.252	0.191	0.211	0.376	0.222	0.964	1.022	0.926	1.135
2003	0.166	0.214	0.322	0.298	0.181	0.200	0.379	0.241	0.921	1.071	0.849	1.238
2004	0.226	0.249	0.361	0.164	0.233	0.239	0.333	0.196	0.969	1.044	1.085	0.838
2005	0.241	0.254	0.377	0.127	0.244	0.247	0.327	0.182	0.988	1.030	1.153	0.700
2006	0.256	0.268	0.377	0.100	0.262	0.253	0.313	0.173	0.977	1.059	1.205	0.579
2007	0.327	0.235	0.309	0.130	0.308	0.259	0.266	0.168	1.060	0.909	1.162	0.775
2008	0.348	0.236	0.200	0.216	0.340	0.245	0.219	0.196	1.023	0.964	0.914	1.102
2009	0.221	0.187	0.253	0.339	0.221	0.186	0.267	0.326	0.999	1.002	0.949	1.041
2010	0.217	0.156	0.300	0.327	0.202	0.172	0.269	0.357	1.074	0.908	1.116	0.915
2011	0.222	0.156	0.288	0.335	0.204	0.174	0.268	0.353	1.087	0.893	1.072	0.948
2012	0.192	0.141	0.220	0.447	0.179	0.155	0.294	0.372	1.076	0.907	0.748	1.201

Notes: Observed and estimated choice probabilities by year for Stage 1 outcomes. First four columns correspond to observed choice probabilities, second four to simulated probabilities at maximum-likelihood estimates, and final four columns to ratio of observed-to-simulated.

Panel C: Model Fit by FICO Score

	Observed				Simulated				Ratio			
	P(Deny)	P(Reject)	P(Hold)	P(Securitize)	P(Deny)	P(Reject)	P(Hold)	P(Securitize)	P(Deny)	P(Reject)	P(Hold)	P(Securitize)
FICO (< 620)	0.301	0.190	0.263	0.246	0.148	0.071	0.181	0.600	2.033	2.664	1.454	0.410
FICO (640)	0.171	0.199	0.273	0.357	0.146	0.068	0.179	0.607	1.171	2.928	1.525	0.589
FICO (660)	0.194	0.191	0.267	0.349	0.150	0.071	0.173	0.607	1.295	2.681	1.546	0.575
FICO (680)	0.238	0.224	0.278	0.261	0.237	0.166	0.176	0.422	1.003	1.350	1.583	0.618
FICO (700)	0.256	0.244	0.320	0.180	0.255	0.245	0.319	0.181	1.007	0.993	1.004	0.992
FICO (720)	0.237	0.237	0.354	0.172	0.233	0.243	0.351	0.174	1.016	0.978	1.009	0.991
FICO (740)	0.214	0.226	0.343	0.217	0.220	0.220	0.346	0.214	0.972	1.028	0.991	1.014
FICO (> 740)	0.211	0.180	0.279	0.330	0.210	0.181	0.282	0.326	1.001	0.993	0.990	1.012

Notes: Observed and estimated choice probabilities by applicant FICO score for Stage 1 outcomes. First four columns correspond to observed choice probabilities, second four to simulated probabilities at maximum-likelihood estimates, and final four columns to ratio of observed-to-simulated.

Table A46: Credit Demand Model Fit by Year and FICO Score

	Observed			Simulated			Ratio		
	P(Default)	P(Continue)	P(Refinance)	P(Default)	P(Continue)	P(Refinance)	P(Default)	P(Continue)	P(Refinance)
Overall	0.000	0.981	0.019	0.000	0.981	0.019	0.799	1.000	1.010
2000	0.000	0.959	0.041	0.000	0.961	0.039	0.554	1.003	0.942
2001	0.000	0.968	0.032	0.000	0.972	0.028	1.327	1.004	0.871
2002	0.000	0.976	0.024	0.000	0.978	0.022	2.433	1.002	0.908
2003	0.000	0.985	0.015	0.000	0.984	0.016	0.340	1.000	1.015
2004	0.000	0.986	0.014	0.000	0.987	0.013	1.293	1.001	0.926
2005	0.000	0.990	0.009	0.000	0.991	0.009	0.739	1.001	0.947
2006	0.000	0.989	0.010	0.000	0.986	0.013	0.436	0.997	1.291
2007	0.000	0.988	0.011	0.000	0.984	0.016	1.047	0.995	1.399
2008	0.000	0.981	0.019	0.000	0.977	0.023	3.495	0.996	1.189
2009	0.000	0.982	0.018	0.000	0.978	0.022	0.713	0.996	1.200
2010	0.000	0.979	0.021	0.000	0.979	0.021	-	1.000	1.017
2011	0.000	0.977	0.023	0.000	0.975	0.025	-	0.999	1.050
2012	0.000	0.987	0.013	0.000	0.975	0.024	-	0.989	1.843
FICO (< 620)	0.000	0.978	0.022	0.000	0.978	0.022	0.727	1.000	1.004
FICO (640)	0.000	0.979	0.021	0.000	0.979	0.021	1.087	1.000	1.000
FICO (660)	0.000	0.978	0.022	0.000	0.977	0.022	0.723	1.000	1.012
FICO (680)	0.000	0.979	0.021	0.000	0.979	0.021	1.099	1.000	1.007
FICO (700)	0.000	0.980	0.020	0.000	0.980	0.020	0.923	1.000	1.007
FICO (720)	0.000	0.981	0.019	0.000	0.981	0.019	0.487	1.000	1.011
FICO (740)	0.000	0.982	0.018	0.000	0.982	0.018	0.807	1.000	1.009
FICO (> 740)	0.000	0.981	0.019	0.000	0.981	0.019	0.598	1.000	1.013

Notes: Observed and estimated choice probabilities by borrower type for Stage 2 outcomes. First three columns correspond to observed choice probabilities, second three to simulated probabilities at maximum-likelihood estimates, and final three columns to ratio of observed-to-simulated.

APPENDIX J: DETAILS ON SIMULATION PROCEDURE

I simulate four policy scenarios outlines in Section 4.4 by iteratively stepping through months, projecting borrower's choices, and constructing key outcome variables. Starting in January 2000, I construct the choice probabilities of all active borrowers. In the case of refinance, I integrate over the shocks (η, e) to generate choice probabilities but use the mean LTV and mean interest rate offers to evaluate outcomes. I evaluate choices separately depending on the latent state, and then weight by the estimated latent state distribution to form predictions. I track nine outcome variables: the probabilities of refinance and default, the volume of cash-out and total refinancing, the mean action value and flow payoff, and the costs, revenue and market shares of the GSEs. I measure GSE costs as the default probability times the remaining loan balance, revenue as the guarantee fee times the loan balance, and market shares as the probability of securitization conditional on refinance. In transitioning from month to month, I track ten key endogenous variables: the borrower's current interest rate, principal balance, monthly payment, and time remaining, as well as their LTV, DTI, spread over prevailing mortgage rates and average mortgage rates at origination, the probability at origination that the originated loan was a GSE securitization, and the probability that the loan has not been defaulted on up to that time. I update these ten variables as an average of the level observed in the next period, weighted by the probability that borrower continues, and the current level, modified based on the terms of refinancing and weighted by the probability of refinancing. In this way, the current level approximately captures a mean value over all possible histories of the borrower's refinancing and continuation decisions[59]. I remove borrowers when they are observed to leave the dataset.

After simulating the entire history, I rescale some of the outcome variables to match observed levels. In particular, I rescale refinancing and cash-out volume to match observed agency refinancing volumes and GSE revenues to match observed guarantee fee income. The assumptions that go into this rescaling are shown in Figure A7. I take data on outstanding GSE mortgage obligations from FHLMC and the FHFA, interpolated from annual to monthly values, and scale them by the FRM share of mortgages estimated by the FHLMC Chief Economist to estimate the volume of outstanding agency FRMs. I impute quarterly refinancing volumes using the figures for total cash-out and cash-out as a percent of refinancing originations in the FHLMC Refinance Report and interpolate to monthly frequencies. These figures are shown in Figure **??**. In Table A47, I show the combined guarantee fee income and single-family

[59] Actually simulating such a history, with 2^T possibilities per borrower, is computationally infeasible.

guarantee costs of FHLMC and FNMA, drawn from their annual 10K filings. I treat this figure as their guarantee fee revenues and costs and rescale the simulation values accordingly.

The first simulation is a baseline simulation, where I use only the observed data. For the second simulation, I analyze the effects of a HARP-type policy by assuming that starting in June 2009, loan applications are never denied for HARP-eligible borrowers. I measure eligibility as the probability in the simulation that the borrower's current mortgage was originated prior to June 2009, was GSE securitized, and has a current LTV ratio of at least 75%, effectively downweighting borrowers for histories in which they refinanced after June 2009, refinanced into a low LTV ratio, or their lender held the loan in portfolio. I assume that lender's interest rate policies and borrowers rejection policies are as in the baseline, and for ineligible borrowers I assume lenders deny applications according to their estimated policy[60]. In the third simulation I restrict borrowers to 80% LTV ratios on refinance by replacing their latent LTV choices $c^*(\eta)$ with 80% whenever they would prefer an LTV greater than 80%. Because by construction the continuous choice function $c^*(\eta)$ is monotone in η this solution method is accurate. I also use these adjusted latent LTV choices in order to predict the mean LTV borrowers select upon refinancing. Finally for the fourth simulation I replace the observed guarantee fee policy shown in blue in Figure A8 with the alternative fee policy shown in red and simulate accordingly as in the baseline model. In each simulation except the fourth I treat the path of all exogenous state variables such as house prices, unemployment rates and interest rates as given. For the fourth simulation, I assume that the average mortgage interest rate, a state variable included in the borrower's value function projection, is the average of interest rates actually received by borrowers obtaining new loans in that period, weighted by the probability that they received a new loan (1 if the borrower is entering the sample or the probability of refinance otherwise). This allows the simulation to

[60] Although the HARP program also placed restrictions on cash-out refinancing, it is beyond the scope of this simulation to model this feature.

accurately capture pass-through from guarantee fees to mortgage interest rates, as documented in Chapter 2.

Table A47: Total GSE Revenue and Cost Assumptions

	FNMA Rev.	FHLMC Rev.	Total GSE Rev.	FNMA Cost.	FHLMC Cost.	Total GSE Cost.
2000	2,430	1,489	3,919	71	40	111
2001	2,591	1,639	4,230	78	45	123
2002	2,516	1,527	4,043	284	122	406
2003	3,281	1,653	4,934	365	-5	360
2004	3,715	1,382	5,097	363	143	506
2005	3,925	2,076	6,001	428	307	735
2006	4,174	2,393	6,567	783	296	1,079
2007	5,816	2,889	8,705	5,003	3,014	8,017
2008	8,390	3,729	12,119	29,725	16,657	46,382
2009	8,002	3,448	11,450	71,320	30,273	101,593
2010	7,206	3,635	10,841	26,420	18,785	45,205
2011	7,507	3,647	11,154	27,218	12,294	39,512
2012	8,151	4,389	12,540	-919	3,168	2,249

Notes: Annual FNMA and FHLMC Single-Family Guarantee Fee Income and Provision for Credit-Related Losses, in $ million.

Figure A7: Total GSE FRMs Outstanding and Refinancing Volume Assumptions

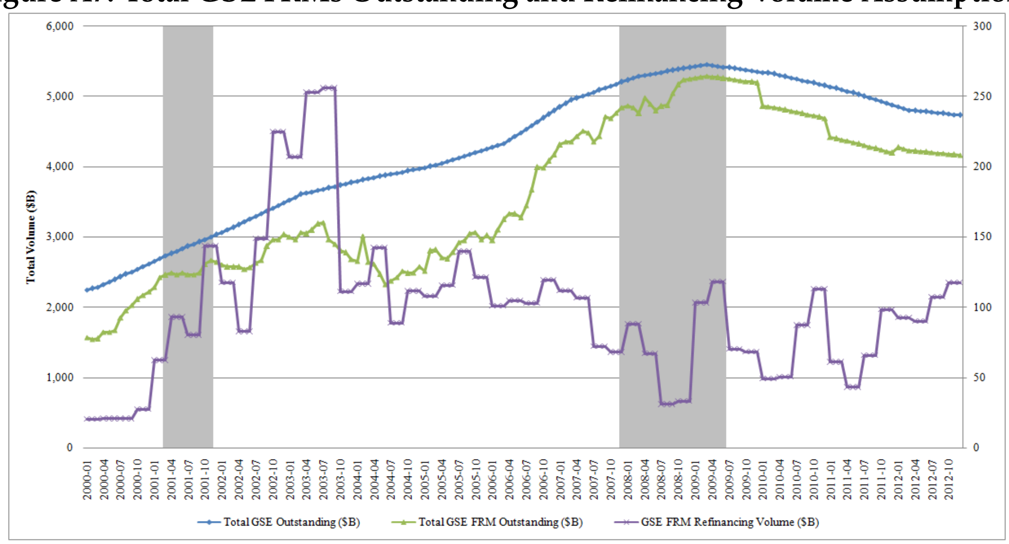

Notes: Imputed combined volume of FHLMC and FNMA FRM mortgages outstanding and refinanced, by month. Grey shading indicates NBER recession dates.

Figure A8: Observed and Alternative Guarantee Fee Policy

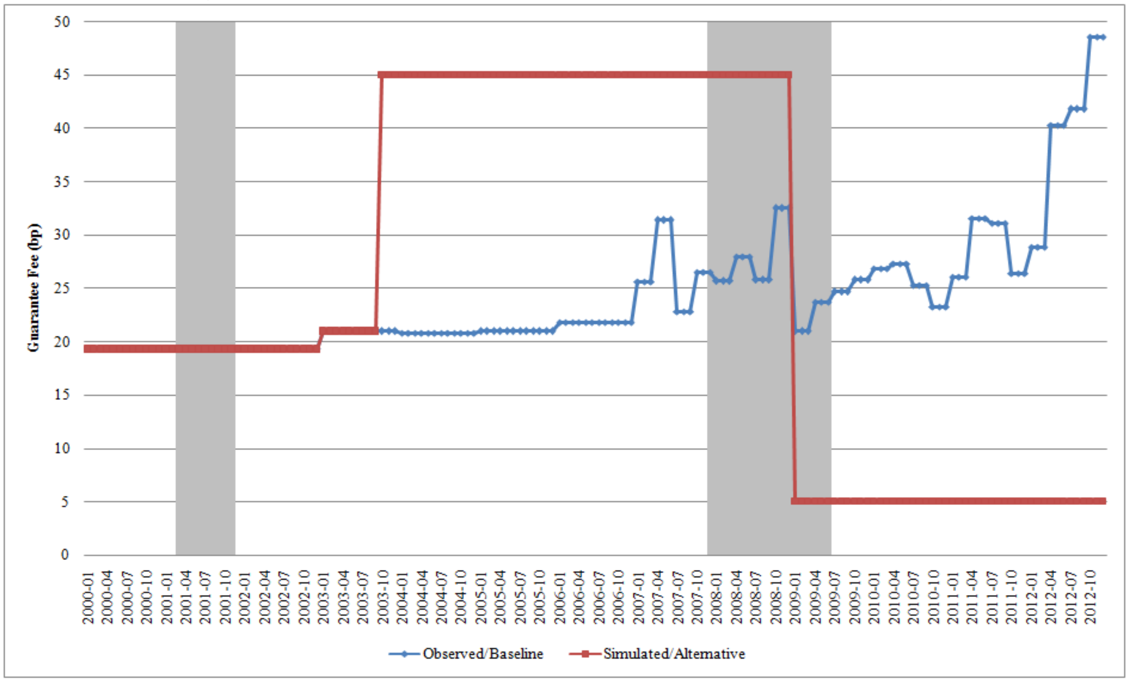

Notes: Observed and alternative average charged guarantee fees used as inputs for simulation. Grey shading indicates NBER recession dates.

Appendix K: List of Abbreviations

- ARM: Adjustable-Rate Mortgage
- BLS: Bureau of Labor Statistics
- bp: Basis Point
- CPR: Conditional Prepayment Rate
- CPS: Current Population Survey
- DTI: Debt-to-Income Ratio
- EM: Expectations Maximization
- FEDS: Finance and Economics Discussion Series
- FICO: Fair Isaac Corporation
- FHA: Federal Housing Administration
- FHFA: Federal Housing Finance Agency
- FHLMC: Freddie Mac or Federal Home Loan Mortgage Corporation

- FNMA: Fannie Mae or Federal National Mortgage Association
- FRM: Fixed-Rate Mortgage
- GSE: Government-Sponsored Enterprise
- GNMA: Ginnie Mae or Government National Mortgage Association
- HAMP: Home Affordable Modification Program
- HARP: Home Affordable Refinance Program
- HHI: Herfindahl-Hirschman Index
- HMDA: Home Mortgage Disclosure Act
- LTV: Loan-to-Value Ratio
- MBS: Mortgage-Backed Security
- MOVE: Merrill Lynch Option Volatility Index
- MSA: Metropolitan Statistical Area
- NBER: National Bureau of Economic Research
- QE: Quantitative Easing
- SCF: Sectional Center Facility or 3-Digit ZIP
- SMM: Single-Monthly Mortality
- TARP: Troubled Asset Relief Program
- TBA: To-Be-Announced Security
- VA: Department of Veterans Affairs
- VAR: Vector Autoregression
- WAC: Weighted-Average Coupon
- WRLURI: Wharton Residential Land Use Regulatory Index