Staff Working Paper/Document de travail du personnel—2024-11

Last updated: April 11, 2024

Unintended Consequences of the Home Affordable Refinance Program

by Phoebe Tian[1] and Chen Zheng[2]

[1] Financial Stability Department
Bank of Canada
xtian@bankofcanada.ca

[2] The Brattle Group
chen.zheng@brattle.com

DOI: https://doi.org/10.34989/swp-2024-11 | ISSN 1701-9397

Acknowledgements

We thank Jean-François Houde for his invaluable support and guidance. This work also benefited from suggestions and comments of Jason Allen, Robert Clark, Ken Hendricks, Katya Kartashova, Shaoteng Li, Lorenzo Magnolfi, Nuno Marques da Paixao, Alan Sorensen, Yu Zhu, participants at the 21st Annual International Industrial Organization Conference, and many seminar participants at the Bank of Canada and the University of Wisconsin–Madison. We also thank Maxim Ralchenko for technical assistance. Any remaining errors are our own. The views expressed are our own and do not necessarily represent those of any institution.

Abstract

We study the unintended effects of the Home Affordable Refinance Program (HARP) on mortgage borrowers. Originally designed to help financially distressed borrowers refinance after the 2008–09 global financial crisis, HARP inadvertently amplified the market power of incumbent lenders by introducing a cost differential between incumbents and their competitors. To assess the welfare implications of this cost advantage, we develop and estimate a structural model of dynamic refinancing decisions with lenders' offers arising from a search and negotiation process. Our findings reveal that although the cost asymmetry was rectified by a 2013 policy, it still resulted in a welfare loss exceeding the impact of search frictions.

Topic: Financial institutions
JEL codes: G21, G51, L51

Résumé

Nous étudions les effets non désirés du Home Affordable Refinance Program sur les emprunteurs hypothécaires. Conçu initialement pour aider les emprunteurs aux prises avec des difficultés financières à obtenir un refinancement après la crise financière mondiale de 2008-2009, ce programme a eu pour effet d'augmenter involontairement le pouvoir de marché des prêteurs initiaux en introduisant un écart de coût entre ces prêteurs et leurs compétiteurs. Pour évaluer les incidences qu'a sur le bien-être cet avantage de coût, nous élaborons et estimons un modèle structurel qui simule des décisions de refinancement en tenant compte de considérations dynamiques, où les offres des prêteurs découlent d'un processus de recherche et de négociation. Nous constatons que même si l'asymétrie des coûts a été corrigée par une politique adoptée en 2013, il en a découlé des pertes de bien-être supérieures à l'impact des frictions de recherche.

Sujet : Institutions financières
Codes JEL : G21, G51, L51

1 Introduction

Economic crises are often associated with a rise in household stress. To help financially distressed households, governments have regularly initiated large-scale mortgage relief programs, since mortgages are the most prominent source of household debt. The implementation of such programs often relies on financial intermediaries. For example, in the United States, while the CARES Act guaranteed that individuals with federally backed mortgages had the right to pause their mortgage payments in response to COVID-19-induced distress, it did not automatically place mortgages in forbearance. Borrowers had to contact their lenders to benefit from the program.

The indirect implementation of such programs leaves room for distortions due to market imperfections, which could, unintentionally, lead to an incomplete pass-through of benefits. This study examines the welfare implications of the 2009 refinance relief program called the Home Affordable Refinance Program (HARP). The Global Financial Crisis (GFC) wiped out home equity for many borrowers, precluding them from refinancing despite declining interest rates. HARP was launched to open up the refinance channel for underwater borrowers in order to boost household consumption. However, a substantial proportion of eligible borrowers did not take up the program (Agarwal et al., 2023).[1] In this paper we focus on the role of lender market power. The program was implemented through mortgage lenders. The initial design of the program, however, unintentionally gave incumbent lenders a cost advantage, amplifying any market power they might already enjoy. In this paper, we ask the following questions: How does the incumbent advantage granted by the program distort market outcomes? What are the implications for consumer welfare?

This paper contributes to the literature by developing a dynamic refinancing model with search frictions and price negotiation. By estimating the model using data from HARP, we quantify the extent to which the program-granted advantage led to an incomplete pass-through of benefits by exacerbating incumbent market power.

Incumbent lenders have an intrinsic first-mover advantage in the mortgage market characterized by search frictions. It is costly for borrowers to obtain additional quotes from lenders other than their incumbent lender, which allows incumbent lenders to price discriminate

[1] The Treasury Department and Federal Housing Finance Agency (FHFA) estimated that 8 million borrowers could have been eligible for the program. The number of borrowers that refinanced during the first five years of the program is approximately 3 million, according to Agarwal et al. (2023).

based on borrowers' outside options and/or search costs (Allen and Li, 2020).

The first four years of HARP exacerbated this advantage. HARP was launched during a period when waves of mortgages were audited for potential violations of underwriting requirements and mortgage lenders were forced to buy back the problematic loans from the government-sponsored enterprises (GSEs). Such forced repurchase is also known as a mortgage *put-back*, which is costly for lenders since they have to bear the credit loss of the mortgage, typically already in a state of severe delinquency. High-loan-to-value (LTV) mortgages are especially prone to mortgage put-back due to their high-risk profile. To encourage the participation of mortgage lenders in high-LTV refinances, HARP lessened legal burdens on incumbent lenders but did not initially extend this treatment to competing lenders. As a result, incumbent lenders have less exposure to forced mortgage put-backs from GSEs.

Using the performance data on HARP loans, we find that a HARP mortgage refinanced through the incumbent lender is half as likely to be put-back than through a competing lender. Therefore, the incumbent lender is faced with lower expected put-back costs when refinancing through HARP. This created an asymmetry in the cost of refinancing between the incumbent lenders and their would-be competitors under HARP (Amromin and Kearns, 2014). We document that incumbent lenders were able to retain 84% of the HARP borrowers in the first phase of HARP. We also find that borrowers who refinanced with their incumbent lenders paid 12 basis points (bps) higher on average than those who switched to competing lenders, conditional on observed characteristics.

Such incumbent markup disappeared after a sharp policy change at the beginning of 2013, four years after the launch of HARP. We find that the new policy had two direct effects on put-back risk. First, it homogenized the risk exposure between incumbent and competing lenders, removing the aforementioned incumbent advantage in terms of put-back risk. Second, the new policy significantly lowered the overall put-back risk for all lenders. The effect of this change was immediate and significant. The market share of incumbent lenders declined and the interest rate on HARP refinances dropped by 41 bps.

The decrease in HARP interest rates reflects both the elimination of the incumbent's cost advantage and the general decline in put-back risk, and their contributions to the rate reduction are unlikely to be proportional to their marginal effects on put-back probabilities. This is because the removal of the cost asymmetry lessens the competitive frictions, leading to a decrease in markup in addition to the cost-driven price decline. Thus, to separate the two effects calls for a model of the incumbent's pricing decision with imperfect competition.

Even if we can separately identify the rate decrease in the absence of incumbency advantage in HARP, the quantification of the welfare loss of such advantage requires more than applying the rate decrease to the affected borrowers. This is because refinancing is an inherently dynamic problem. The timing of refinance is an endogenous outcome, changing with the market condition. Borrowers who refinanced after the removal of the cost asymmetry would do so earlier if the cost asymmetry was absent from the beginning. An earlier refinancing decision shortens the total amortization of their mortgage, lowering the total mortgage payments. On the extensive margin, eliminating the cost advantage in the beginning could encourage more refinancing activities from those who would never refinance otherwise.

Therefore, quantifying the welfare loss from the four years of incumbency advantage in HARP requires accounting for the dynamic consideration in borrower's refinancing decisions as well as the lender's pricing problem. We develop a model of the borrower's dynamic refinancing problem in which interest rate offers come from a two-stage bargaining process with lenders. The model starts with a mortgage origination. The underlying house value is subject to idiosyncratic shocks in addition to market-level price changes in each subsequent period, updating the borrower's LTV ratio. The new LTV determines the borrower's eligibility for HARP. At the beginning of each period, the borrower makes a discrete choice of whether to refinance if not hit by a default shock. This decision involves a dynamic tradeoff between the expected reduction in future mortgage payments and the lump-sum cost of refinancing in the current period. The lump-sum cost includes fixed transaction costs and potential search costs.[2] If the borrower decides to refinance, she first interacts with her incumbent lender, who offers an initial quote from a quasi-monopoly position. The borrower also observes her search cost at this point, which remains her private information. If the borrower rejects the initial offer, a search process begins, with the borrower paying the search cost. In the search stage, the borrower organizes an English auction among incumbent and competing lenders and takes the best offer.

In the model, the borrower's refinancing decision depends on the competitiveness of the lending market as well as the idiosyncratic housing shock. When the incumbent has a

[2]Transaction costs in mortgage refinancing include application fees, property appraisal fees, title search fees, etc.

lower expected cost than competing lenders, it charges a higher markup in the initial offer instead of passing through its low cost. This is because the lower cost gives the incumbent a competitive advantage and thus higher expected profits in the search stage; therefore, the threat of searching is not as effective in the first stage. As the borrower anticipates unattractive rate offers, she is inhibited from refinancing in the beginning.

We tailor the model in to fit the actual roll-out of HARP during 2009–2018. First, we introduce the put-back risk into a lender's expected cost of refinancing in the model. Until the end of 2012, the incumbent has a lower expected put-back risk than competing lenders, which translates into a cost advantage. Since 2013, all lenders' risk exposure is symmetric and generally lower, mimicking the policy change documented in the empirical analysis. Second, HARP notably underwent major modifications to the program rules, with the most prominent ones in 2012 collectively known as HARP 2.0. To reflect these modifications, we allow the fixed cost, LTV ceiling, and accessibility of the program to change after 2011. In the model, the fixed cost of the program affects a borrower's decision of whether to refinance, the LTV ceiling determines a borrower's eligibility for HARP, while the accessibility of the program governs how likely a HARP-eligible borrower is to choose HARP instead of regular refinancing. Accommodating for these changes also allows us to evaluate the welfare implications of such modifications, which serves as an interesting comparison with our main quantification goal, the welfare implication of the cost advantage in HARP.

These changes also provide over-time variations to help identify model primitives. For example, the link between the observed drop in prices and in put-back probabilities after 2012 helps in determining the cost parameter of mortgage put-back. We estimate the model using maximum likelihood, which is an easier incorporation of observed borrower characteristics in our estimation.

The results can be summarized as follows. First, the expected put-back cost is economically meaningful in the pre-2013 period. With an average of \$7.90 per \$100 of the mortgage, it makes up 18.8% of the total cost for a competing lender. Second, the asymmetric risk exposure leads to a significant cost advantage for the incumbent lenders. The incumbent's expected put-back cost is less than \$3 per \$100 mortgage, or 8.3% of its total cost. This leads to a cost advantage of \$4.90 in terms of put-back risk, while the cost differential from all other sources is less than a dollar. The cost advantage implies that the incumbent lender's winning probability is double that of competing lenders in a typical market. In the post-change period, the expected put-back cost is almost negligible because of the extremely low put-back probability.

We use the model estimates to evaluate the implication of the cost advantage on market outcomes and borrower welfare, and how it interacts with the pre-existing search friction in the market. In our first counterfactual exercise, we shut down the asymmetric risk exposure since the beginning of HARP. The welfare improvement mainly comes from the intensive margin, with a relatively small impact on the overall refinancing rate (0.3%). Among the HARP borrowers who refinanced post-change in the baseline model, 3.8% of them chose to refinance earlier in the counterfactual, with an average four years difference in the timing. Their eagerness was driven by their higher-than-average initial LTV and loan balance. Overall, the interest savings for those who refinanced increased by 19 bps on average, which translates into a 1.8% reduction in annual mortgage payments. Accounting for the shorter amortization due to an earlier refinance timing, this leads to a 2% reduction in lifetime mortgage payments. The average borrower welfare across all sampled borrowers (i.e., all Freddie homeowners with a loan originated between 2003 to 2006) increased by $2,100 compared to the baseline level.

Our second counterfactual experiment not only shut down the asymmetric risk exposure, but also introduces the general decline in put-back risk from the beginning of 2009. We find that the general decline leads to another 1.6% reduction in lifetime mortgage payments with an additional total welfare effect of $1,900. Although the general risk reduction has an almost three times larger marginal effect on put-back probabilities than the removal of the cost asymmetry, it does not lead to higher welfare implications. This implies that the cost asymmetry not only has a direct effect of a higher expected cost, but also exacerbates competitive frictions.

To further explore the interaction between the cost asymmetry and market power, we examine the welfare gain from removing the cost asymmetry in an environment free of search friction. When the incumbent no longer has the first-mover advantage, does the cost advantage have larger or smaller welfare implications? We find that the cost asymmetry is associated with an average decrease of $1,900 in borrower welfare if the incumbent does not have first-mover advantage, which is smaller than the welfare effects in the baseline case. In other words, the pre-existing competitive friction has an amplifying effect on the policy-related market inefficiencies.

Another finding from our counterfactual experiments is that the inefficiencies from search friction are not as high as that from the cost advantage. To make further comparisons, we evaluate the effects of HARP 2.0 modifications in comparison to the cost asymmetry. We find that among all three main targeted areas of HARP 2.0 modifications, the changes related to fixed costs (e.g., waiving appraisal fees) had the largest overall impact on borrower welfare, while improving program accessibility and awareness contributed most to the utilization rate of the program. In total, HARP 2.0 contributed to an average $ 1,200 increase in borrower welfare. Although smaller than the effects of the cost asymmetry, this is still an economically significant enhancement to the program. Furthermore, our counterfactual experiments also show a marginal decline in default outcome related to HARP 2.0 and the change in put-back policy.

In sum, our results show that HARP in general helps borrowers to refinance, save on interest costs, and even avoid default, especially with the subsequent modifications in HARP 2.0. However, we also find that the benefit of the program would be much higher if the program did not start with an unintentional cost advantage for the incumbent lenders. This exacerbated the pre-exisitng incumbency advantage and eroded the pass-through of the program benefits. Our analysis suggests that, for programs implemented indirectly through intermediaries to reach the targeted agents, details that have market power implications can have has much impact as the high-level design of the program.

Related Literature Our study relates to several strands of literature. First, it is related to the literature that examines the importance of institutional frictions and financial intermediaries in the effective implementation of stabilization programs, particularly in housing markets. Piskorski et al. (2010) and Agarwal et al. (2017) study the Home Affordable Modification Program (HAMP), another federal program that tried to help households modify their mortgages during the GFC. They find that lender-specific factors, such as servicing capacity and cost structure, negatively impacted the effectiveness of the policy intervention. Abel and Fuster (2019) study the effects of HARP on household debt and spending. Amromin and Kearns (2014) and Agarwal et al. (2023) point out the design flaw in HARP that gave rise to competitive frictions. Agarwal et al. (2023) estimate that the competitive frictions reduced the take-up rate and annual savings among those who refinanced by 10–20%. We add to this literature by quantifying the welfare loss associated with the incumbency advantage induced by HARP and put it into perspective by comparing it to the welfare effects of search frictions and high-level program design. We also examine the interaction between the policy-related incumbency advantage with pre-existing market conditions.

This study also contributes to the literature on refinance decisions in the U.S. mortgage market (Keys et al., 2016; Agarwal et al., 2017; DeFusco and Mondragon, 2020). It is well documented that many U.S. mortgage borrowers do not refinance, even in the presence of seemingly large financial gains from doing so. To explain the lack of refinancing, this literature focuses on borrower-specific factors, such as inattention or liquidity. While such borrower-specific factors can partially account for the muted response to HARP, our paper emphasizes the search friction and imperfect competition of financial intermediaries in explaining part of this shortfall. Closely related, Ambokar and Samaee (2019) explore the role of search costs in explaining such inaction by developing and estimating a dynamic discrete choice model of refinancing with search friction. Our model differs in two ways: First, we highlight the special position of the incumbent lender. We give the incumbent lender a first-mover advantage by permitting an initial quote from the incumbent to preempt the borrower's search efforts. Second, we allow for price negotiation to be embedded in the search process. The incumbent could revise its offer in response to a competing offer.

This paper fits into the literature that examines market power in household finance. Previous studies (Woodward and Hall, 2012; Honka, 2014; Scharfstein and Sunderam, 2016; Allen et al., 2019; Allen and Li, 2020; Agarwal et al., 2020) have documented various sources of market power. We add to this literature by studying the U.S. refinance market and highlighting the role of stimulus policies on market power. Our paper is most closely related to Allen et al. (2019), who propose a search and negotiation framework to quantify the magnitude of incumbent advantage in the Canadian mortgage market. The search and negotiation process in our model is somewhat simpler because we assume no recall of the initial offer in the auction stage, and this simplification gives us a closed-form solution of the initial offer and distribution of competitive offers. This allows us to embed this search and negotiation process into the borrower's dynamic refinancing problem and check its implications for the timing of refinance decisions and dynamic selection. Another related paper, Allen and Li (2020), examines dynamic competition in the market with search and price negotiation. They focus on a setting in which the timing of refinance is fixed and does not involve the borrower's dynamic decision. In our paper, we focus on the dynamic problem of a borrower's refinance decision and treat the lender's problem as static. Finally, this study is related to papers on price dispersion in the U.S. mortgage market. There is substantial dispersion in lenders' offered rates (Alexandrov and Koulayev, 2018; McManus et al., 2018) as well as in transacted rates (Gurun et al., 2016; Agarwal et al., 2020; Bhutta et al., 2020).

The remainder of this paper is organized as follows. Section 2 provides the institutional background on HARP. Section 3 describes in detail the data source used for the analysis and then uses the data to document some key patterns of the program's features. The model is presented in Section 4. Section 5 discusses the empirical specification of the model, identification arguments, and estimation procedures. Section 6 presents estimated parameters and assesses model fit. Section 7 presents results from counterfactual simulations. Section 8 concludes. Additional technical details, tables, and figures can be found in the Appendix.

2 Program and Background

2.1 U.S. Mortgage Market

The U.S. mortgage market is organized into two segments, primary and secondary. The primary market is where borrowers and lenders meet and negotiate lending terms to create a mortgage transaction, while the secondary market trades mortgage loans and mortgage-backed securities (MBS). The primary buyers in the secondary mortgage market are government-sponsored enterprises (GSEs), Fannie Mae and Freddie Mac. After acquiring the mortgages, they bundle them into the MBS, which is later sold to investors. GSEs guarantee full payment of interest and principal to investors on behalf of lenders. In return, they charge the lenders an upfront guarantee fee. Mortgage lenders that securitized loans through a GSE typically retain mortgage servicing rights, which is the main source of cash flow for mortgage lenders.[3]

The majority of mortgage contracts in the U.S. are 15- or 30-year fixed rate mortgages. In the U.S., most borrowers can repay their mortgage at any point in time without penalties. This is usually done by refinancing their mortgage backed by the same property with either

[3]The role of a servicer includes collecting payments, advancing them to the MBS trustee, and engaging in various loss-mitigating actions on delinquent loans. The terms "servicer" and "lender" are used interchangeably.

the incumbent lender or a new lender. Therefore, mortgage borrowers can take advantage of a decrease in interest rates by refinancing their loans. However, since any new mortgage needs to be underwritten, its availability depends on the borrower's creditworthiness and whether the borrower has enough equity in their home. Traditionally, lenders require an LTV ratio of no more than 80% for refinance transactions, although the maximum they are willing to accept is 95% if the borrower is willing to pay an upfront mortgage insurance premium.

2.2 HARP Program

We focus on the period immediately following the collapse of U.S. housing prices during the GFC. As house prices fell, many borrowers had near zero or negative equity in their home (they were "underwater"). The Federal Reserve cut interest rates, meaning many households could refinance at interest rates lower than their current rate. Despite lower rates, however, underwater households were unable to refinance because their LTV ratios made them ineligible. In response, the federal government, working with Fannie Mae and Freddie Mac, developed HARP in 2009 to expand the set of borrowers who could refinance their loans. The goal was to help underwater borrowers regain access to the refinance market, which could lower their mortgage payments and thus reduce mortgage default rates. The program allowed borrowers with LTV ratios higher than 80% to refinance their mortgages by extending federal credit guarantees on those loans. Other qualification requirements included no delinquency record in the previous 12 months and that the original mortgage was owned by a GSE. Crucially, the program allowed each borrower only one chance to take advantage of HARP. Borrowers with previous HARP refinancing would be ineligible to do so again.

The program went through a series of changes after its initial launch, with the most notable modifications known as HARP 2.0.[4] This enhancement of HARP targeted three main areas. First, it lowered certain costs and fees related to the refinancing process. Specifically, it eliminated the manual property appraisal requirement, which helped streamline the refinancing process and reduced borrower costs. Second, it removed the 125% LTV ceiling, which expanded the HARP-eligible population to include seriously underwater borrowers. Last but not least, it implemented a nationwide public relations campaign to educate bor-

[4]Details of the modifications can be found at https://www.fhfaoig.gov/Content/Files/EVL-2013-006.pdf.

rowers and increase their awareness about HARP. Previously, borrowers may not have been aware that they were permitted to refinance their mortgages with any participating lender.

The program was initially set to expire at the end of 2012, with a series of subsequent announcements extending that deadline. The program officially ended at the end of 2018.

2.3 Representations and Warranties and Mortgage Put-Back

HARP was created as part of the post-GFC stimulus efforts, but its progress was impeded by lenders' concerns over loan put-back risk stemming from the GSE representation and warranty framework. When selling mortgage loans to the GSE, lenders must assure that the loan selling and servicing processes comply with the guidelines outlined by the GSE. These are formally known as representations and warranties ("reps and warrants") contracts. Reps and warrants relate to factors such as mortgage underwriting, borrower eligibility, the mortgage product, the property, and the project in which the property is located. With this assurance, the GSE does not need to conduct a thorough evaluation on each individual loan when purchasing from the lender, which streamlines the loan delivery process and facilitates the growth of the U.S. mortgage market. The GSE can, however, conduct reviews on any loan after it is delivered to the GSE. If the GSE determines that the loan violates any of the reps and warrants, the GSE is entitled to require the lender that delivers the defective loan to buy it back. This forced repurchase is referred to as a loan *put-back*. A loan put-back is often extremely costly for the lender because the GSE typically conducts reviews only when a loan is in the state of default. By buying back the loan, the lender has to bear all the credit loss. There is also significant uncertainty for lenders since GSEs have discretion in what constitutes a trigger for a loan put-back.

During the GFC, with large waves of mortgage defaults, the GSEs became more aggressive in terms of auditing the delinquent loans for any defects and breaches of reps and warrants contracts. In Figure 1(a), we plot the distribution of loan put-backs by year, calculated as the number of put-backs in each year divided by the total number of put-backs. Loan put-backs were uncommon prior to the crisis; however, they peaked following the crisis, when the rate of put-back tripled compared to pre-crisis levels. Given the uncertainty around the reps and warrants framework, the put-back risk perceived by lenders during that time was likely to be higher than the ex post rates.

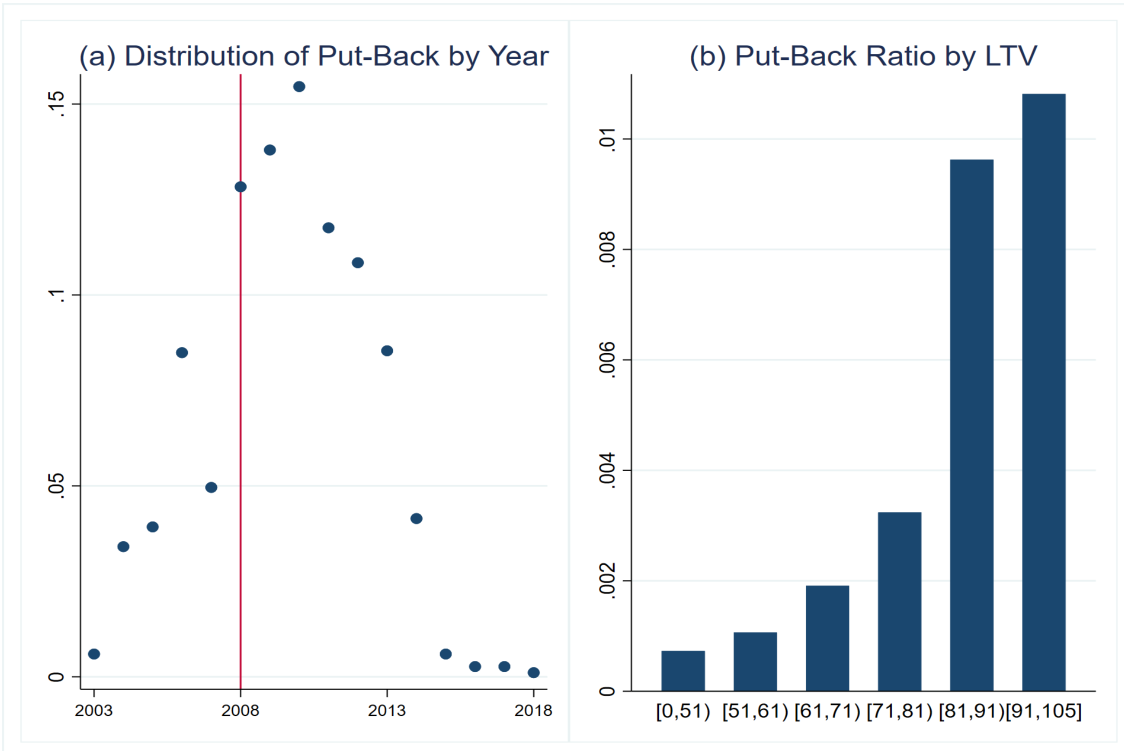

Figure 1: Tightening Underwriting Standard during HARP

Figure 1(a) shows the fraction of loans that are put-back in a given year among all put-back loans in the sample. Figure 1(b) shows the fraction of loans that are put-back among all loans in each LTV bucket.

It is against this backdrop of heightened put-back risk that HARP was launched. From a lender's perspective, put-back risk was one of the most prominent concerns for HARP in its early years. This is because HARP loans are high LTV in nature. This makes them more likely to be audited and reviewed by the GSE and thus more likely to be put-back. As shown in Figure 1(b), only 0.2% of borrowers with LTV $\leq$ 80 resulted in a mortgage put-back, whereas 1% of high-LTV borrowers (i.e., LTV over 80%) have their mortgages put-back from GSEs. In other words, a lender's exposure to put-back risk when originating a high-LTV mortgage is five times as high as a regular mortgage. Therefore, put-back risk remains a great concern for lenders to accept high-LTV mortgages during the post-crisis period.

Moreover, lenders' exposure to put-back risk was not symmetric during the first few years of HARP because of ambiguity in the reps and warrants framework. The incumbent—that is, the lender that is currently servicing the original mortgage—retains all the documents and payment history of the original loan. Although this is not directly related to the new loan after refinance, these documents serve as additional evidence to prove that the lender has done its due diligence. In an environment where lenders are uncertain about the enforcement rules of the reps and warrants, any supplemental material could be of help. Thus, the residual risk of loan put-back fell unevenly on lenders, depending on whether they had a previous relationship with the borrower or not.

2.4 The 2013 Policy Change on Mortgage Put-Back

In response to this issue, the GSEs and FHFA launched a new framework for the enforcement of the reps and warrants, effective January 1, 2013. The goal of the new framework was to reduce uncertainty for lenders by clarifying lenders' exposure and liability to loan put-back. First, the new framework established a unifying criterion for a loan to be relieved from the reps and warrants liabilities and thus a future put-back. For HARP loans, lenders were spared from loan put-back if the HARP borrower made on-time payments during the first 12 months after the acquisition of the loan by the GSEs. Second, the new framework directed the GSEs to evaluate loan files and identify potential defects earlier in the loan process, rather than when a loan defaults. Early quality control reviews like this avoid the worst outcome, namely to put-back a defaulted loan and let the lender bear all the credit loss. Third, the review process is conducted on a more consistent and systematic basis, rather than using discretion and relying on isolated instances of misstatements and misrepresentations. Overall, the new framework made it clear for the lenders what is needed to minimize a loan put-back, and it does not involve documents or history from the previous relationship. With transparent enforcement rules, lenders other than the incumbent lender could now manage the put-back risk to the same degree as the incumbent.

In the next section we perform two descriptive analyses using the HARP sample to illustrate that the data pattern is consistent with the discussion in this section.

3 Data

Our data come from three sources. The first is the single-family loan-level dataset from Freddie Mac. The second is HARP, which are made public by Freddie to promote the transparency of the program. The third is the House Price Index from the FHFA, which measures the price movement of single-family houses. We then use the data to document the key features of the program.

3.1 Data Sources

Single-Family Loan-Level Dataset Freddie Mac started publishing single-family loan-level data to support risk sharing and transparency. The dataset starts in 2000 and are updated quarterly. It comprises two parts: acquisition and performance. The acquisition file provides the characteristics of loans acquired by Freddie at the loan origination level. The loan characteristics that we observe include credit score (FICO), LTV ratio, debt-to-income (DTI) ratio, loan amount, loan purpose (e.g., home purchase, cash-out refinance, or no cash-out refinance), quarter of origination, property ZIP code (three-digit), and the name of the lending institution. The performance file is a panel that provides monthly credit performance, which includes the monthly loan balance and delinquency status. The loan exits the performance file if it is terminated by the borrower via a prepay/refinance or foreclosure.

HARP Data The HARP data, which is a subset of the U.S. single-family loan-level dataset, uniquely allows us to link every HARP refinance to its previous mortgage information. This allows us to identify the households that were refinanced under the program, as well as constructing the key variables of the analysis, such as whether they refinanced with their incumbent lender and the interest rate reduction they received from HARP.

House Price Index We also use the FHFA House Price Index (FHFA HPI©). FHFA uses data on mortgage transactions from Fannie Mac and Fannie Mae to calculate the index using a modified version of the weighted-repeated sales methodology. This quarterly index measures changes in single-family home values at the national, census division, state, metro area, county, ZIP code, and census tract levels. We match borrowers in our main sample with the HPI at the three-digit ZIP code level, which is the finest geographic location disclosed by the single-family loan-level data. The HPI aids in the estimation of borrowers' home values after loan origination. The estimation procedure is discussed in detail in Section 5.1.2.

National Survey of Mortgage Originations We use the National Survey of Mortgage Originations (NSMO) as a complementary data source for search behavior. The NSMO survey represents a random sample of about 6,000 mortgages drawn quarterly from loans newly reported to one of the three national credit bureaus. It is a nationally representative sample of newly originated, closed-end, first-lien residential mortgages in the U.S. We use this external dataset to construct an auxiliary moment on search behaviors in our estimation.

3.2 Sample and Summary Statistics

3.2.1 Sample Construction and Variable Definition

Our sample consists of homeowners with a 30-year fixed-rate mortgage owned by Freddie Mac that originated between 2003Q1 and 2006Q4 with the purpose of purchasing a property. We refer to the origination year as the cohort of a borrower. These "purchase mortgages" are referred to as the original mortgage. We obtain the loan and borrower characteristics on the original mortgages from the origination data file in the single-family loan-level dataset. This is then merged with the HARP origination data file, which contains the loan identification number of the corresponding legacy mortgage. This allows use to identify loans in the main dataset that are refinanced through HARP following their origination. Furthermore, the merged data contains the information on their subsequent HARP refinances, such as the interest rate, LTV ratio, and lender information. We consider a borrower to switch lenders if the lender on the HARP refinance is not the same as the lender on the original mortgage.[5] Since the single-family loan-level dataset does not provide specific lender names when a lender's market share is too small, we discard observations when both the previous lender

[5] Technically speaking, we define a switching behavior when the servicer of the original mortgage is not the same as the seller of the new mortgage. This is because a mortgage's servicing right is often sold by the originator of the mortgage to other financial intermediaries after its origination. From the borrower's perspective, the servicer is the one with whom they directly interact and build familiarity at the time of refinancing. On the other hand, the seller of the new mortgage is more likely to be the one that interacts with the borrower during the refinancing process. We also used other ways to define switching and found similar results.

and new lender's names are missing. We also do not include HARP mortgages that are not of standard term length.[6] These account for 21,542 observations, or 1% of the whole sample.

For each original mortgage, we construct the loan outcome variable from the monthly performance data file, which contains information on the repayment status of each loan up until June 2018, the end of the sample period. We classify each loan into four outcomes: default, HARP refinanced, other prepaid, and no action. Default includes two scenarios. First, the loan's balance is reduced to zero for reasons other than voluntary payoff. Second, the zero balance is due to voluntary payoff, but the loan is at least 90 days in delinquency in the last period before being paid off. We treat the second scenario as a voluntary default, likely caused by the owners selling their home voluntarily to avoid foreclosure. Other voluntary payoffs that don't appear in the HARP dataset are considered "other prepaid." A loan is considered "no action" if it is still active by the end of the sample period.

3.2.2 Summary Statistics

Table 1 reports the summary statistics for a number of variables of interest. Panel A is the main data. This sample contains people who purchased a house before the crisis during 2003–2006. These are all purchase loans with 30-year fixed interest rates. Their FICO score on average is 729, and the LTV is on average 78%, or 22% down payment. The mean of initial interest rate is 600 bps, with an average loan size of $172,000.

Panel B then reports the HARP program takers among those from Panel A. We report the characteristics for HARP takers separately before and after the mid-HARP policy change on underwriting standards between incumbent and competing lenders that took place in 2013. First of all, between 2009–2012, FICO scores for HARP takers actually increased from 729 to 750, presumably because HARP has a requirement that borrowers cannot have a missing mortgage payment in 12 consecutive months. However, LTV for those borrowers increased from 0.78 to 1.04, suggesting a loss of home equity for those households as a consequence of the 2008 housing crisis. The (refinance) interest rate that households obtained from the program was 452 bps between 2009–2012, compared to 412 bps between 2013–2018, a period when the market interest rate (i.e., the cost of credit) also decreased.

During the first half of HARP, the switching rate is only 16% , which implies a market

[6] We keep the three most predominant term lengths for HARP loans, which are 180 months (18%), 240 months (14%) and 360 months (64%).

Table 1: Loan-Level Summary Statistics

Panel A: GSE Single-Family Data, 2003–2006

	Mean	S.D.
Loan Characteristics		
FICO Score	729	54
LTV	0.78	0.14
Interest Rate (bps)	600	46
Loan Size (1,000$)	172	84
Cohort Distribution		
2003	0.25	0.43
2004	0.24	0.43
2005	0.27	0.44
2006	0.24	0.43
Loan Outcome		
Default	0.060	0.24
Other Prepaid	0.78	0.42
HARP Refinanced	0.088	0.28
No Action	0.075	0.26
# of Observations	2,124,685	

Panel B: HARP Refinance, 2009–2018

	2009–2012		2013–2018	
	Mean	S.D.	Mean	S.D.
FICO Score	750	59	732	74
LTV	1.04	0.26	1.06	0.27
Interest Rate (bps)	452	63	412	55
Loan Size (1,000$)	197	77	163	69
Switching Rate	0.16	0.36	0.27	0.44
Put-back	0.002	0.046	0.001	0.023
# of Observations	130,329		56,204	

This table presents descriptive statistics for the data source used in this paper. Panel A shows the statistics for the parent data, which is the main GSE data that contains purchase loans from 2003–2006. Panel B presents the HARP takers among those in Panel A. This is separated by those who participated in HARP before the mid-HARP policy change on underwriting standards between the incumbent and competing lenders.

share of 84% for incumbent lenders, compared to a regular refinance market where the incumbent market share is 28% to 33% across different years (Agarwal et al. (2023)). The switching rate increases to 27% after the policy change in 2013. Figure 2 plots the switch rate over time. It is low during the first half of the program due to the asymmetric put-back probabilities between incumbent and competing lenders. It started to gradually increase after the policy change in 2013. Finally, the probability of put-back also decreased by twofold following the policy change.

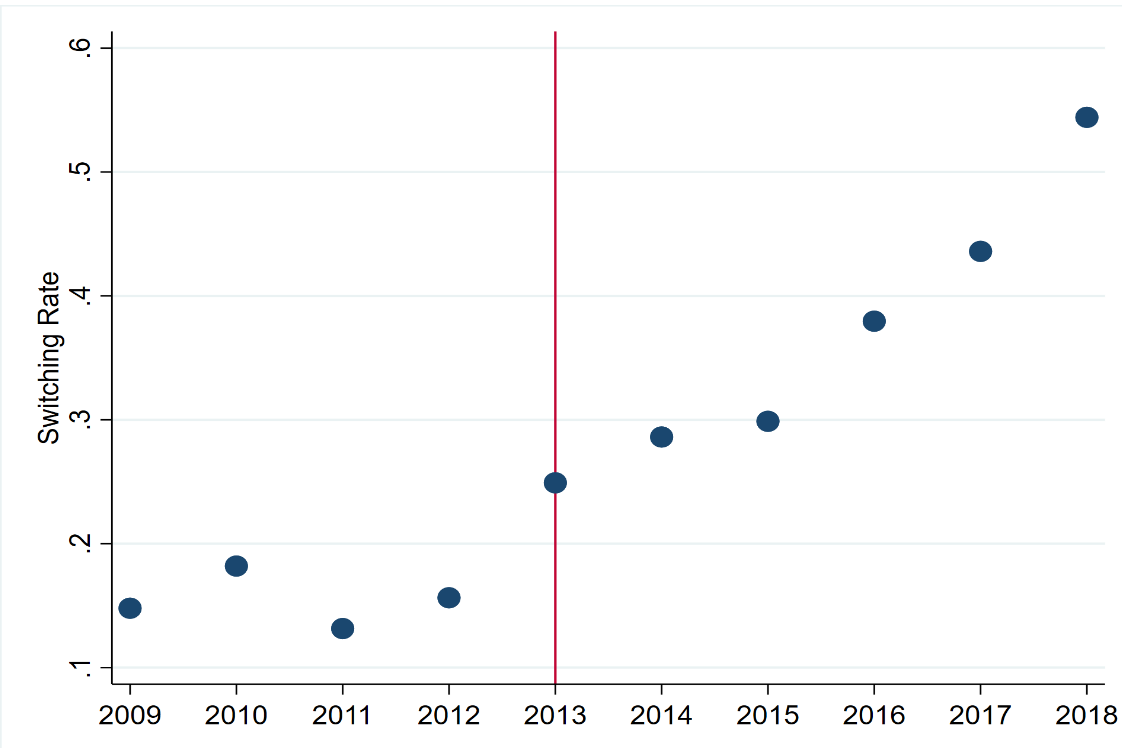

Figure 2: HARP Switch Rate Over Time

This figure shows the average switching rate among HARP borrowers in each year from 2009 to 2018. A borrower is considered to switch lenders if the lender on the new mortgage is not the same as the lender on the original mortgage.

3.3 Descriptive Analysis

In this section, we perform two descriptive analyses using the HARP sample to illustrate the effect of the mid-program policy change on put-back probability and prices.

3.3.1 The Effect of the 2013 Policy Change on Put-Back Probability

Since we can only observe switching behavior for HARP refinances, we use the put-back outcomes on HARP refinances to access the effect of the policy change. We estimate the put-back probability p^{PB} estimated via a logistic regression model:

$$p_{ij}^{PB} = \frac{\exp\left(X_{ij}'\hat{\delta}\right)}{1 + \exp\left(X_{ij}'\hat{\delta}\right)} \tag{1}$$

The dependent variable X_{ij} includes $Incumbent \times Pre$, $Post$, $Incumbent$, FICO, income, LTV and loan amount of the new mortgage, interest rate on the original mortgage, market-cohort fixed effects, and fixed effects for HARP origination year and HARP loan term. The dummy variable $Incumbent = 1$ indicates same-lender refinance, $Pre = 1$ indicates the period before 2013, and $Post = 1 - Pre$ indicates the post-change period.

Column (1) of Table 2 presents the estimates for δ. $Incumbent \times Pre$ has a negative effect, suggesting that during the first half of the program borrowers who refinanced through HARP with incumbent lenders were less likely to have their loans put-back than those who switched to competing lenders. The marginal effect is on average 0.23%, while the average put-back rate in the pre-2013 period is 0.45%, which makes same-lender refinances half as likely to be put-back than other refinances. The coefficient for $Incumbent$ is not statistically significant, suggesting that the asymmetry no longer holds in the second half of the program, which eliminates the difference between incumbent and competing lenders regarding put-back probabilities. Moreover, the policy change is also associated with a general reduction in put-back probabilities in the second half of the program, as is evident from the coefficient on $Post$. This is consistent with the policy background, which is intended to clarify lender's risk exposure, lower uncertainty, and create a level playing field for every lender. The marginal effect of the $Post$ is large, with an average of 0.53%. This suggests that the put-back risk is no longer a significant concern after the new policy.

Table 2: Descriptive Analysis

	(1) Put-back	(2) HARP Rate
Incumbent X Pre	-1.242^{**}	0.145^{***}
	(0.452)	(0.00748)
Post	-2.877^{**}	-0.411^{***}
	(1.054)	(0.0384)
Incumbent	0.158	-0.0248^{**}
	(0.433)	(0.00836)
LTV	0.00651^{*}	0.00185^{***}
	(0.00286)	(0.000108)
FICO	-0.00560^{***}	-0.000272^{***}
	(0.000943)	(0.0000547)
log(Income)	-0.712^{***}	0.00705^{**}
	(0.201)	(0.00228)
log(Balance)	0.647^{*}	-0.0729^{***}
	(0.255)	(0.00460)
Previous Rate	0.532^{**}	0.203^{***}
	(0.176)	(0.00536)
R-squared	0.116	0.702
HARP Orig Year FE	Yes	Yes
Seller FE	No	Yes
Cohort X Market FE	Yes	Yes
Observations	183,331	186,533

This table reports the results of the descriptive analysis using the sample of HARP loans (i.e., data from Panel B of Table 1). Column (1) reports the estimates for δ from a logit model in Equation 1 where the dependent variable is put-back probability. The figures in parentheses are standard errors with 1, 2, and 3 asterisks indicating statistical significance at 10%, 5%, and 1%, respectively. R-squared is the pseudo R-squared from the logit model. Column (2) reports the coefficient estimates from a regression model in Equation 2 where the dependent variable is the HARP refinance interest rate. The figures in parentheses are cohort-market clustered standard errors, with 1, 2, and 3 asterisks indicating statistical significance at 10%, 5%, and 1%, respectively.

3.3.2 The Effect of the 2013 Policy Change on Prices

We now examine the change in interest rate on HARP refinances after the new policy through a regression design shown in Equation 2. We keep the same set of independent variables except for additional fixed effects of the incumbent lender's identity.

$$r_{ijtm} = \beta_0 + \beta_1 Incumbent_{ij} + \beta_2 \text{Post}_{it} + \beta_3 Incumbent_{ij} \times \text{Pre}_{it} + Z'_{ijtm}\gamma + \epsilon_{ijtm}, \tag{2}$$

Column (2) of Table 2 presents the results. During the first half of the program, borrowers with incumbent lenders on average pay 12 bps higher than those who switched to competing lenders.[7] After the policy change, interest rates dropped significantly for both stayers (41 bps) and switchers (56 bps), which is consistent with the presence of an incumbent advantage and the cost-reduction effect of the mid-program policy.[8] Note that the price differential between stayers and switchers is no longer positive after the policy change; it is negative but economically insignificant (2.5 bps), which could be a result of other cost differences between the incumbent and competing lenders.[9]

In sum, this section shows some key patterns from the program and how they change before and after the policy change. It suggests that the mid-program change had implications for the incumbent advantage. However, for quantification purposes, the regression results cannot be directly extrapolated. This is because of the presence of dynamic selection, which implies that the early HARP takers are unobservably different from the late HARP takers. Accounting for this calls for a structural model of borrowers' refinancing choices that endogenizes the timing of refinance given the market structure and program design. This model is described in the next section.

4 Model

Our model is finite horizon with discrete time periods. In this section, we focus on a single borrower's dynamic refinancing problem and omit the borrower index i in the subscript. We also omit all time-invariant variables, including observed borrower characteristics (income,

[7] 0.145 - 0.0248 = 0.1202.

[8] -0.411 - 0.145 = -0.556.

[9] We do not find any significant difference in default risk for stayers and switchers.

FICO, market) and characteristics of the original mortgage (origination year, interest rate, principal, LTV).

4.1 Timing and information

In our model, a borrower starts with an existing fixed-rate mortgage ($t = 0$). From the next year ($t = 1$), the borrower's dynamic refinancing problem begins. At the beginning of each period t the borrower is faced with a probability of default $(1-p_t^C)$. In the case of non-default, she checks her updated house value, h_t, and the current cost of funds in the market, c_t^m, to make a refinance decision. The superscript m specifies the market in which the borrower is situated. We assume that the change in her house value relative to the original house value at $t = 0$, denoted as $\Delta h_t = h_t/h_0$, is a known function of the market-level change in house value, $\Delta h_t^m = h_t^m/h_0^m$, and an individual-specific temporary shock, q_t. The transition of the market-level variables $z_t = (h_t^m, c_t^m)$ is assumed as a Markov process. The idiosyncratic housing shock q_t is an i.i.d. draw from $N(0, \sigma_q)$ in each period. It is unobserved by the econometrician, but known by agents in the model.

The updated house value determines the borrower's current LTV, and therefore her eligibility for HARP refinancing. If the LTV is less than 80%, or it exceeds the ceiling imposed by HARP, the borrower is not qualified for HARP refinancing. Otherwise, HARP refinancing remains an option in addition to regular refinancing. The two types of refinancing differ in the fixed costs borne by the borrower, represented by $\phi^k(z, q)$, where $k = H$ (HARP refinancing) or $k = R$ (regular refinancing). This cost encompasses property appraisal fees, mortgage insurance, and various transaction fees. It is a function of LTV, with borrowers at a higher LTV typically incurring greater fixed costs, particularly when the LTV exceeds 80%. We assume that the fixed cost is generally lower for HARP refinancing, $\phi^H(z, q) < \phi^R(z, q)$ for $\forall(z, q)$, due to the waived appraisal and streamlined underwriting process.

If the borrower decides to refinance and qualifies for HARP, we assume that she chooses HARP refinancing with probability ξ. If the borrower is well informed about the program, ξ should equal 1 for because it dominates regular refinancing for qualified borrowers. A value of $\xi < 1$ indicates limited awareness of HARP. In our empirical specification, we allow the fixed cost of HARP refinancing, the probability of offering HARP, and the LTV ceiling for HARP qualification to vary with different phases of HARP's roll-out.

Given the type of refinancing, the borrower then negotiates a price with mortgage lenders through a two-stage process. In the first stage, the borrower contacts the incumbent lender, namely the one servicing her current mortgage. The incumbent lender makes an initial offer r^I. At this point, the borrower privately observes her search cost κ, which is a uniformly distributed random variable with mean $\bar{\kappa}$. Without loss of generality, we can write $\kappa = \bar{\kappa} \cdot \epsilon$, with $\epsilon \sim U[1-e, 1+e]$ and $e \in [0, 1]$.

Then, the borrower decides whether to take the initial offer or to reject it and search for a competitive offer by paying the search cost κ. If the initial offer is rejected, the borrower organizes an English auction among all lenders in the market and takes the lowest offer, thus ending the dynamic refinancing problem. If she chooses not to refinance, she will still have the option of refinancing in the next period and the process continues.

Figure 3 summarizes the timing of events. The tradeoff faced by the borrower is that refinancing in the current period might lower the borrower's interest rate and future monthly payments, but it involves a lump-sum fixed cost to refinance—transaction cost and potential search cost. The dynamic refinancing decision also depends on the expected LTV trajectory, because a lower LTV generally leads to better interest rates on refinance.

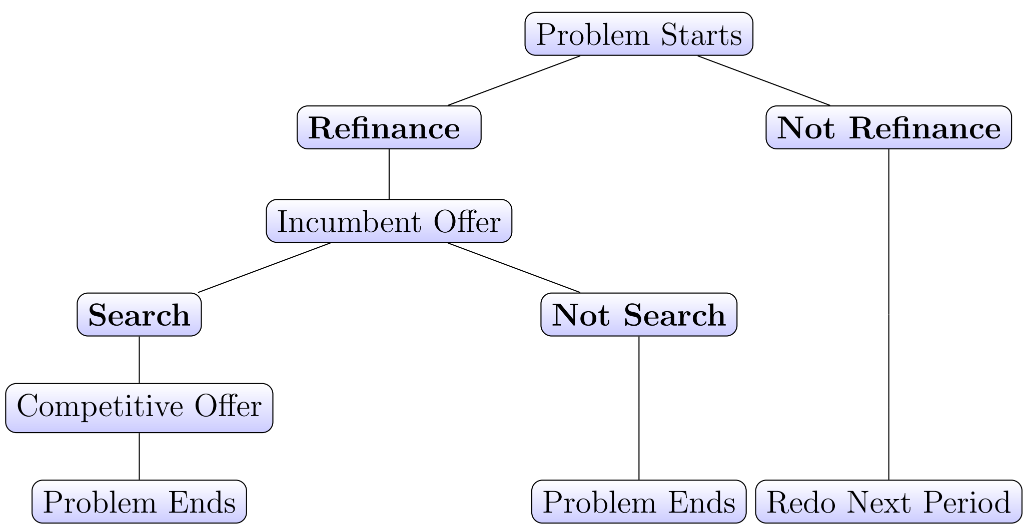

Figure 3: Timing of Borrower Decision

This figure shows the timing of borrowers' decisions. Borrowers first decide whether to refinance. If so, they decide whether to accept their incumbent lender's offer after receiving a free quote or pay a search cost to gain additional quotes. They must accept a competitive offer once they decide to search. If they decide not to refinance, they make the decision again in the next period.

Before solving the model, two remarks are in order. First, HARP imposes a one-time-only requirement, allowing each borrower only one chance to take advantage of the program. Therefore we assume that in the model a borrower has only one opportunity to refinance. Second, we assume that once a borrower decides to refinance at the beginning of a period, she commits to the refinance decision. In other words, she either takes the incumbent's initial offer or the competing offer by the end of the period. This assumption greatly simplifies the game by ruling out non-refinance as an outside option in the price-setting game, thus making the model tractable.

4.2 Price Negotiation

In this dynamic discrete choice problem, the decision to refinance is terminal, meaning that no further dynamic decisions will be required in subsequent periods after the refinancing choice has been made in the current period. A price negotiation process ensues, which we model as a two-stage game. This section focuses on the price negotiation process contingent on the time period of refinancing t, the type of refinancing k, and state variables (z_t, q_t). They are omitted in the following discussion for notational simplicity.

4.2.1 Utility and Profit Functions

We start the discussion by specifying the expected utility and the profit of refinancing as functions of price r. Here, r refers to the *amortized* interest rate, defined as $r =$ *Monthly Payment/Loan Size*.[10]

Utility of Mortgage Refinance After refinancing the mortgage at price r, a borrower makes repayment $m(r)$ on the new mortgage for T periods (unless hit by the default shock) and becomes a mortgage-free homeowner afterwards. We define $U_t(r)$ as the lifetime utility of refinancing at price r, which can be calculated recursively by

$$U_\tau(r) = u(y - m(r)) + \beta \left[p_\tau^C U_{\tau+1}(r) + \left(1 - p_\tau^C\right) \underline{U}_\tau \right], \ \tau = t, ..., t+T. \tag{3}$$

[10]The relationship between amortized interest rate and annualized percentage rate $\tilde{r}$ is given by

$$r = \frac{r\%/12}{1 - (1 + r\%/12)^{-12T}}$$

$u(\cdot)$ is the flow utility on consumption, which is income y net of mortgage payment $m(r)$.[11] β is the discount factor. p_t^C is the probability of non-default until the next period, $t+1$, conditional on non-default until period t. In other words, it is the probability that the borrower can continue to make the $t+1$th payment conditional on having made t payments. It is an exogenous function of borrower and loan characteristics as well as loan age.[12] $\underline{U}_t$ is the lifetime utility after default, given by $\sum_{\tau=t+1}^{\bar{T}} \beta^{\bar{T}-\tau} u(y) - h_0$, where $\bar{T}$ is the last period of the borrower's life ($\bar{T} > T$). The terminal value of the recursive calculation is $U_{t+T+1} = \sum_{\tau=t+T+1}^{\bar{T}} \beta^{\bar{T}-\tau} u(y)$, which is the discounted sum of utility flows of a mortgage-free homeowner.

To facilitate the analytical derivation of price offers, we specify a linear utility function: $u(y - m(r)) = y - m(r)$. It follows that $U_t(r)$ is also linear in r. Let α_t denote the slope of $U_t(r)$ with respect to r, which varies with the time of refinancing, t, and the loan balance at the time of refinancing. The constant term in $U_t(r)$ is represented by $\bar{U}_t$. Therefore, we use the following representation of $U_t(r)$ in the subsequent discussion:

$$U_t(r) = \bar{U}_t - \alpha_t r \tag{4}$$

Profit Function Mortgage lenders in the market are indexed by j. We reserve $j = 0$ for the incumbent lender and $j = 1, ...J$ for outside (competing) lenders. For each dollar of the loan amount, r is the lender's incoming monthly cash flow, and the outgoing cash flow has two components: the guarantee fee paid to GSEs, denoted as g, and the cost of funds c. Similarly, we convert both g and c to amortized rates, so that $r - g - c$ is the net cash flow per dollar of the mortgage in each month. Following Fuster et al. (2013), we multiply the net cash flow with a predetermined multiplier, M, to obtain the expected revenue per dollar throughout the span of the mortgage: $M(r - g - c)$.

This leads to the following baseline profit function, $\pi(r)$:

$$\pi(r) = M(r - g - c) - P_0, \tag{5}$$

where P_0 represents the expected cost of a potential put-back for the incumbent lender.

[11] We abstract way from the borrower's saving choice and other non-mortgage borrowings.

[12] We treat default as an event triggered by exogenous shocks rather than modeling it as a choice. In our empirical specification, we estimate p_t^C by using a survival function.

For all other lenders, the expected put-back cost is $P_0 + \Delta_P$. Here, $\Delta_P > 0$ prior to the implementation of the put-back policy change, and $\Delta_P = 0$ following the policy change. The variable Δ_F encapsulates the cost differential from other sources that competing lenders may have relative to the incumbent lender. The total cost advantage of the incumbent, Δ, is thus the sum of Δ_P and Δ_F. In addition, there is an idiosyncratic shock to lender j's lending cost, ω_j. Therefore, the expected per-dollar profit of refinancing at price r is:[13]

$$\pi(r) - \Delta\, \mathbb{1}\{j \neq 0\} - \omega_j, \tag{6}$$

where the distribution of ω_j is given by a (minimum) Gumbel distribution with mean zero and scale parameter σ_ω.

4.2.2 Competitive Stage

We now describe the solution of the negotiation by backward induction, starting with the competition stage. If the borrower rejects the initial offer and starts to search ($S = 1$), the incumbent lender enters into an English auction, competing with other lenders in the market. The competition stage commences with each lender observing an idiosyncratic shock to his lending cost for the borrower, ω_j. We define the effective cost shock $\tilde{\omega}_j = \omega_j + \Delta \mathbb{1}\{j \neq 0\}$ to capture the cost asymmetry, and rewrite Equation (6) as $\pi(r) - \tilde{\omega}_j$.

The winning lender, denoted as j^*, is the one with the lowest effective cost shock: $j^* = \arg\min_j\{\tilde{\omega}_j\}$. The probability that lender j wins the auction is given by

$$p^W_{j^*} = \begin{cases} \frac{1}{J \cdot \exp(-\Delta/\sigma_\omega)+1}, \text{ if } j^* = 0 \\ \frac{\exp(-\Delta/\sigma_\omega)}{J \cdot \exp(-\Delta/\sigma)+1}, \text{ if } j^* \neq 0. \end{cases} \tag{7}$$

When $\Delta = 0$, the incumbent wins with the same chance as other lenders, $p^W_0 = 1/(J + 1)$. When $\Delta > 0$, that is, the incumbent has a cost advantage, the incumbent wins with a higher probability than any competing lender, $p^W_0 > 1/(J + 1)$, and the incumbent's chance of winning increases with the extent of the advantage Δ.

The winner in the auction charges an interest rate r^C that makes the closest runner-up

[13] We assume that M and g are known functions of borrower characteristics.

just break even:

$$r^C = \pi^{-1}\big(\tilde{\omega}_{(2)}\big) \tag{8}$$

where $\tilde{\omega}_{(2)}$ is the runner-up lender's effective cost shock, or the second lowest of all. The distribution of $\tilde{\omega}_{(2)}$ conditional on the winner j^*, denoted as $F_{\tilde{\omega}_{(2)}|j^*}$, has an analytical form with the following conditional expectation (Brannman and Froeb, 2000).The $\tilde{\omega}_{(2)}$ given j^* is:

$$E\big[\tilde{\omega}_{(2)} \,\big|\, j^*\big] = -\sigma_\omega \log\left(J \exp\left(-\Delta/\sigma_\omega\right) + 1\right) - \frac{\sigma_\omega \log(1 - p^W_{j^*})}{p^W_{j^*}}. \tag{9}$$

It follows that the expected competitive offer, denoted as $\bar{r}^C$, is

$$\bar{r}^C = -\frac{\sigma_\omega}{M}\,\pi^{-1}\left(\log(J \exp\left(-\Delta/\sigma_\omega\right) + 1) + \sum_{j^*=0}^{J} \log(1 - p^W_{j^*})\right) \tag{10}$$

Finally, the incumbent's expected profit in the competitive stage is given by:

$$\bar{\pi}^S_0 \equiv E[\pi_0 \,|\, S = 1] = -\sigma_\omega \log(1 - p^W_0), \tag{11}$$

which increases with the incumbent's cost advantage Δ.

4.2.3 Initial Stage and Search Decision

The incumbent lender solves the profit-maximization problem upon receiving an inquiry from a borrower. Given any incumbent's quote r, the borrower chooses to search if the net gain from searching, $\Delta U(r) = U(\bar{r}^C) - U(r) = \alpha(r - \bar{r}^C)$, is greater than the search cost, κ.

$$S = \mathbb{1}\{\kappa < \Delta U(r)\} \tag{12}$$

Letting H denote the distribution function of the search cost, it then follows that the rejection probability is $H\left(\Delta U(r)\right)$. Thus, the incumbent's initial offer comes from the following problem:[14]

$$\max_r (1 - H(\Delta U(r)))\,\pi_0(r) + H(\Delta U(r))\bar{\pi}^S_0. \tag{13}$$

[14]Note that we normalize the incumbent's cost shock in this stage as zero.

The specification of linear utility function and uniform distribution of search cost transforms the incumbent's problem in Equation (13) into a quadratic optimization problem, which has a closed-form solution. Specifically, the initial offer, r^I, is a piecewise linear function:

$$r^I = \begin{cases} \bar{r}^C + \frac{\bar{\kappa}(1-e)}{\alpha}, & \text{if } \hat{r} - \bar{r}^C \leq \frac{\bar{\kappa}(1-3e)}{\alpha}, \\ \frac{1}{2}\left[\hat{r} + \bar{r}^C + \frac{\bar{\kappa}(1+e)}{\alpha}\right], & \text{if } \frac{\bar{\kappa}(1-3e)}{\alpha} < \hat{r} - \bar{r}^C \leq \frac{\bar{\kappa}(1+e)}{\alpha}, \\ \hat{r}, & \text{if } \hat{r} - \bar{r}^C > \frac{\bar{\kappa}(1+e)}{\alpha}. \end{cases} \tag{14}$$

The associated search probability is

$$\Pr(S = 1) = \begin{cases} 0, & \text{if } \hat{r} \leq \bar{r}^C + \frac{\bar{\kappa}(1-3e)}{\alpha}, \\ \frac{\alpha\left(\hat{r} - \bar{r}^C\right)}{4\bar{\kappa}e} - \frac{1-3e}{4e}, & \text{if } \bar{r}^C + \frac{\bar{\kappa}(1-3e)}{\alpha} < \hat{r} < \bar{r}^C + \frac{\bar{\kappa}(1+e)}{\alpha}, \\ 1, & \text{if } \hat{r} \geq \bar{r}^C + \frac{\bar{\kappa}(1+e)}{\alpha}, \end{cases} \tag{15}$$

where $\hat{r} = \pi^{-1}(\bar{\pi}_0^S)$ is the incumbent's reservation price. It is the interest rate at which the incumbent lender is indifferent whether the offer is accepted or not, because the expected profits are the same. A borrower's reservation price is $\bar{r}^C + \kappa/\alpha$, which is additive in $\bar{r}^C$. Thus, the term $\hat{r} - \bar{r}^C$ governs the difference in reservation price between the incumbent and the borrower. A lower value of $\hat{r} - \bar{r}^C$ increases the likelihood that the borrower's reservation price exceeds that of the incumbent, thereby increasing the probability of the initial offer being accepted. In essence, $\hat{r} - \bar{r}^C$ serves as a measure of the incumbent's pricing advantage during the initial stage. Depending on the value of $\hat{r} - \bar{r}^C$, the pricing function has different slopes.

When $\hat{r} - \bar{r}^C$ is smaller than the first cutoff, $\frac{\bar{\kappa}(1-3e)}{\alpha}$, the initial offer is flat at $\bar{r}^C + \frac{\bar{\kappa}(1-e)}{\alpha}$. This is the price at which the borrower with the lowest search cost $\bar{\kappa}(1-e)$ is indifferent between searching and not searching. Thus, the offer is accepted with probability 1. Any price lower than it cannot further increase the acceptance probability, and thus it serves as a floor on the initial offer. When $\hat{r} - \bar{r}^C$ falls between the two cutoffs, $\hat{r} - \bar{r}^C \in \left(\frac{\bar{\kappa}(1-3e)}{\alpha}, \frac{\bar{\kappa}(1+e)}{\alpha}\right)$, the initial offer increases with $\hat{r}$ at a slope of 1/2. In this interval, the monopolistic incumbent faces the classic tradeoff between price and demand, and the price is determined by the interior solution to the first-order condition of the profit-maximization problem. The slope of r^I with respect to $\hat{r}$ is analogous to the pass-through rate of marginal cost, which is 1/2 due to the incumbent's monopoly position in this case. The probability of searching changes linearly in $\hat{r}$ from 0 to 1. In the last scenario, $\hat{r} - \bar{r}^C$ is larger than the second cutoff, $\frac{\bar{\kappa}(1+e)}{\alpha}$. At this point, even the borrower with the highest search cost would search, since the net gain from searching outweighs the search cost. Thus, an initial offer in this interval is rejected with probability 1.

An interesting observation is that the incumbent pricing advantage is higher with a smaller number of competing lenders or a smaller dispersion of cost shock. In other words, $\hat{r} - \bar{r}^C$ increases with J and σ_ω, while the effect of Δ is ambiguous.[15] This is because although Δ inflates the expected profit in the competitive stage (from Equation (11)), it also drives up the expected competitive offer $\bar{r}^C$, so the net effect depends on the comparison of the two opposing forces.

An analysis of the incumbent's market power calls for an examination of the markup. Let $r^b = P_0/M + g + c$ denote the break-even price. Then the incumbent's markup in the initial stage can be decomposed into two parts:

$$r^I - r^b = \underbrace{r^I - \hat{r}}_{\text{Markup from search friction}} + \underbrace{\hat{r} - r^b}_{\text{Markup from cost advantage}}.$$

The first part, $r^I - \hat{r}$, measures how much more the incumbent charges above its reservation price in the first stage. The second part, $\hat{r} - r^b$, measures the difference between its reservation price and the break-even price. Since $\hat{r}$ increases with $\bar{\pi}_0^S$, it therefore increases with the cost advantage Δ. In other words, the incumbent's cost advantage in the competition stage drives up the reservation price it is willing to offer in the initial stage.

To illustrate how the first part of the markup, $r^I - \hat{r}$, arises from search friction, we plot r^I as a function of $\hat{r}$ in Figure 4. The dashed line is at 45 degrees, so the distance between the solid line and the dashed line represents $r^I - \hat{r}$. This markup term is higher with lower values of $\hat{r}$, which means higher pricing advantage.

When $\bar{\kappa}$ increases, the pricing function shifts up to the red line, suggesting that the initial quote increases with $\bar{\kappa}$. Intuitively, $\bar{\kappa}$ is the average search cost, and higher average

[15] This can be seen from

$$\hat{r} - \bar{r}^C = \frac{\sigma_\omega}{M}\, \pi^{-1}\big(\log(J \exp(-\Delta/\sigma_\omega) + 1\ ) + J \log\big(1 - (1 - P_0^W)/J\big)\big).$$

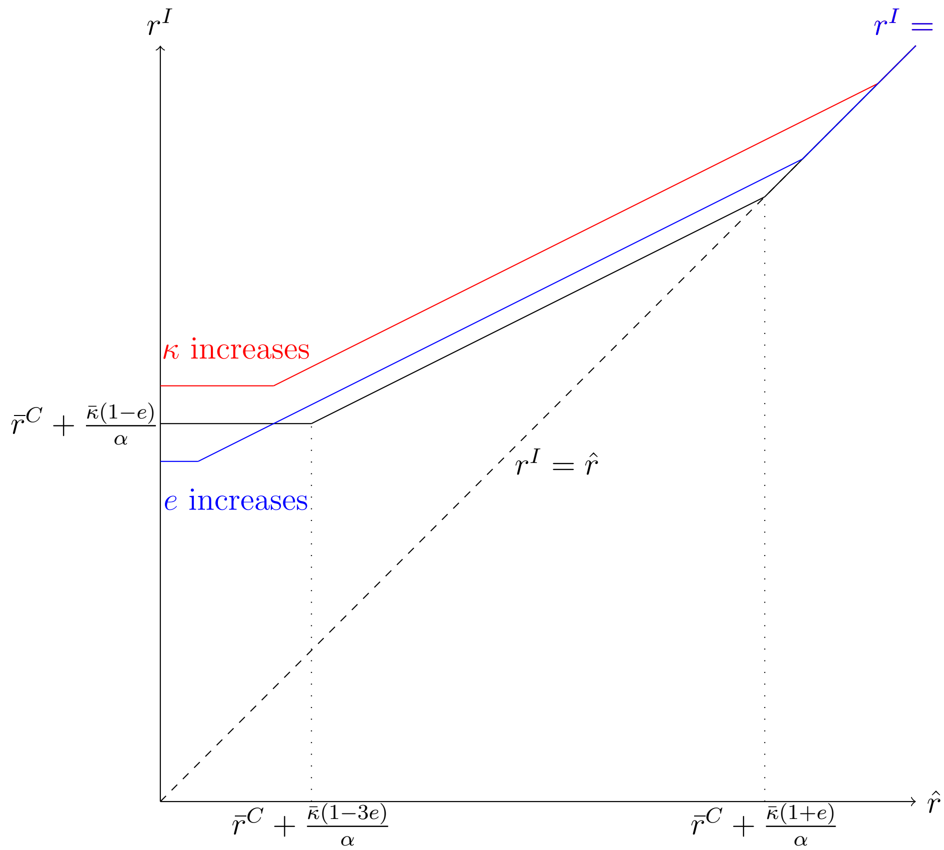

Figure 4: Pricing Function in the Initial Stage

This figure illustrates the initial offer r^I as a function of the reservation price $\hat{r}$, as shown in Equation (14). The black solid line is the pricing function at the baseline level. The red solid line shows the pricing function after an increase in κ, and the blue solid line represents the pricing function after an increase in e.

search cost gives the incumbent higher market power in the first stage, thus extracting more surplus.

Interestingly, the effect of e on the initial quote is ambiguous, depending on the value of $\hat{r} - \bar{r}^C$. In the case of a small $\hat{r} - \bar{r}^C$, the incumbent has a high pricing advantage and is able to preempt searching efforts by offering an attractive initial offer. If e goes up in this case, the marginal borrower's search cost, $\bar{\kappa}(1-e)$, goes up. Therefore, the initial offer required to preempt searching must be lower. On the other hand, this also implies that it is more costly for the incumbent to preempt searching, so the incumbent is more likely to seek an interior solution instead (the middle piece of the pricing function). In this classic scenario of the monopoly pricing problem, a higher e results in a less elastic demand curve, because a 1 percent change in price now leads to less change in demand. Consequently, the optimal price is higher in response to the increasingly inelastic demand.[16]

4.3 Refinancing Decision

In this subsection, we lay out the borrower's value function and the associated policy function for the refinancing decision. We add back the state variables (z, q) and time subscript t while suppressing refinancing type k.

The ex-ante value of refinance, viewed at the beginning of period t, is the maximum between the expected value of accepting the incumbent's offer and the expected value of searching for competing lenders' offers, net of the fixed cost of refinancing:

$$V_t^{refi}(z,q) = E\big[\max\big\{U_t(r^I),\ U_t(r^C) - \kappa\big\}\big] - \phi(z,q), \tag{16}$$

If the borrower does not refinance in period t, she retains the chance to refinance in the future. The value of waiting is thus the sum of flow utility and the discounted expectation of the continuation value:

$$V_t^{wait}(z,q) = u(y - m^0) + \beta\left[p_t^C EV_{t+1}(z',q') + \left(1 - p_t^C\right)\underline{U}_t\right], \tag{17}$$

where m^0 is the mortgage payment on the original mortgage.

In addition to the factors accounted for V^{wait} and V^{refi}, there are other unobserved

[16]Similar results hold for the incumbent's expected profit in this stage.

determinants on a borrower's refinancing decision. On the positive side, refinancing can present opportunities such as cashing out for home renovations or debt consolidation, shortening the loan term, or even improving credit scores for borrowers who struggle to make payments. Conversely, refinancing activity involves a significant time commitment for borrowers in addition to monetary costs. We use a scalar, μ, to reflect the time-invariant utility effect of refinancing, with $\mu > 0$ indicating a net benefit and $\mu < 0$ suggesting a net cost. Other idiosyncratic unobserved determinants are summarized in a pair of utility shocks, $(\epsilon_t^0, \epsilon_t^1)$, which are i.i.d. random variables with zero means. The expected value of having a refinancing opportunity in period t is thus given by:[17]

$$V_t(z,q) = \max\left\{V_t^{wait}(z,q) + \epsilon_t^0,\ V_t^{refi}(z,q) + \mu + \epsilon_t^1\right\},\ t = 1, ..., T. \tag{18}$$

5 Estimation and Identification

We now discuss our method to estimate model primitives. We begin by discussing specifications we make to the model in order to fit the empirical settings. Section 5.1.1 specifies borrowers' beliefs in different time periods, in line with the actual timeline of HARP's roll-out and the policy change regarding mortgage put-back. In Section 5.1.2, we describe methods employed for the parameterization of a variety of functions and distributions, as well as off-model estimations of some of the functions. We discuss the sources of identification of model parameters in Section 5.2, and then derive the likelihood function and describe the estimation procedure in Section 5.3.

5.1 Empirical Specifications

5.1.1 Timeline

As mentioned in Section 2, HARP was launched in 2009 followed by several changes to the program rules and the related policy on mortgage put-back. In our empirical model, the launch of HARP and the subsequent modifications to the program, as well as the change in put-back policy, are not foreseeable by a borrower. In other words, a borrower's belief on

[17]The terminal value is given by $V_{T+1} = \sum_{\tau=T+1}^{\bar{T}} \beta^{\bar{T}-\tau} u(y,h)$, i.e., the discounted sum of the utility flow of a mortgage-free homeowner until the end of life.

HARP and the put-back policy changes with different phases of the program. We identify four phases with different beliefs and solve the dynamic refinancing problem corresponding to each belief. We then keep the implied refinancing decision within each corresponding phase. Details on the four phases are as follows:

1. From the year of mortgage origination, Y_0, to 2008, a borrower's refinancing decisions are derived under the belief that HARP does not exist and the put-back policy remains the same. The borrower's belief on HARP will go through a series changes after the launch of HARP and its subsequent modifications in the next four years, but the borrower's belief on put-back policy will stay the same until the beginning of 2013.

2. From 2009 to 2011, HARP becomes available (HARP 1.0). During this period, we assume that a borrower's refinancing decisions are made under the belief that HARP ends at the end of 2012. We also specify an LTV ceiling of 125% for HARP qualification and assume that the fixed cost associated with HARP refinancing during this period is $\phi^H = \phi_0 > 0$. These assumptions aim to capture the initial specifications of the program when it was first announced by FHFA.

3. In 2012, a modified version of HARP becomes available, also known as HARP 2.0. We obtain a borrower's refinancing decision in 2012 under the belief that HARP 2.0 ends at the end of 2013. HARP 2.0 differs from HARP 1.0 in that the LTV ceiling is removed as is the fixed cost $\phi_H = 0$. This assumption reflects the program enhancements made by FHFA at the beginning of 2012, addressing several key aspects of HARP including removing the 125% LTV ceiling, extending the end date of HARP to the end of 2013, eliminating the need for a new property appraisal, and streamlining the underwriting process. This largely reduced the fixed cost of HARP refinancing, which we normalize to zero in the model.

4. Since the beginning of 2013, the put-back policy is changed, and the end date of HARP is postponed to the end of 2018. Under this belief, we calculate a borrower's refinance decisions from 2013 to the end of the refinancing window. The post-change put-back policy eliminates the difference in expected put-back costs faced by the incumbent and the competing lenders, and lowers the expected put-back costs in general.

Table 3 summarizes borrowers' beliefs during the four phases in the empirical model.

Table 3: Timeline of the Empirical Model

Phase	**Borrower's Belief**	
	HARP	Put-Back Policy
$Y0$–2008	No HARP	No change
2009–2011	HARP 1.0 from 2009–2012	No change
2012	HARP 2.0 from 2012–2013	No change
2013–End	HARP 2.0 from 2013–2018	Post-change

This table summarizes the timeline of the empirical model. The first column shows the time periods for each of the four phases, where $Y0$ refers to the start year of a borrower's original mortgage. The second column describes the borrower's belief about HARP in each phase, and the third column shows the borrower's belief about the put-back policy in each phase.

5.1.2 Parametrization

Since we focus on 30-year fixed-rate mortgages, $T = 30$, we set a borrower's life horizon to $\bar{T} = 50$.[18] The number of competing lenders J is set as a quarter of the total number of lenders in the market, rounded to the nearest integer. We specify the guarantee fee g as a function of the borrower's FICO and LTV, based on the g-fee matrix in the annual report published by GSEs. The distribution of utility shocks to refinancing decisions, $(\epsilon_t^0, \epsilon_t^1)$, are assumed T1EV with mean zero and scale parameter σ_V.

The fixed cost of regular refinancing is

$$\phi^R = \phi_0 + 1\% \cdot LoanBalance \cdot \mathbb{1}\{LTV > 80\%\}$$

where $\phi_0 > 0$ represents the property appraisal fees and other transaction costs, and the second term reflects the extra private mortgage insurance involved in regular refinancing when LTV exceeds 80%. As mentioned in the previous subsection, the fixed cost of HARP refinancing during 2009–2011 is ϕ_0, and it reduces to zero afterwards.

Repayment Probability We use a log-logistic survival function to model the probability of non-default until period t: $\left[1 + (\lambda_i t)^{1/s}\right]^{-1}$, where λ_i is parameterized as $\exp\left(-X_i' b\right)$ and

[18] Note that borrowers cannot refinance from $t = 21$ to $t = 30$ because the remaining lifetime is shorter than the mortgage term.

X_i includes borrower characteristics (FICO and income), original interest rate and LTV, principal, cohort fixed effects, and market fixed effects. Using the monthly performance data on the original mortgages in the sample, we estimate b and s using the maximum likelihood method, and the results are presented in Section A.1. Using the model estimates, we then calculate p_t^C, the probability of non-default until period $t+1$ conditional on non-default until period t, as

$$p^C = \frac{1 + \left(t\hat{\lambda}\right)^{1/\hat{s}}}{1 + \left((t+1)\hat{\lambda}\right)^{1/\hat{s}}}. \tag{19}$$

The Market-Level House Value Index and the Cost of Funds We use the yearly average coupon rate in the MBS market as the measurement of c_t, which is on the national level. We define market m on the state level, and we use the HPI for each state as the measure of h_t^m. For each $m = 1, ..., 51$, we estimate a VAR(1) process for $(\log(h_t^m), c_t)$ and use a discrete approximation to the VAR(1) process via the method proposed by Farmer and Toda (2017).[19]

Idiosyncratic House Value Shock Given the market-level change in house value, Δh_t^m, we assume that the individual-level change in house value, Δh_{it}, is determined by

$$\log(\Delta h_{it}) = \beta_0 + \beta_1 \log(\Delta h_{it}^m) + q_{it},\ \ q_{it} \sim N(0, \sigma_q). \tag{20}$$

The conditional distribution of Δh_t given Δh_{it}^m is thus determined by β_0, β_1, and σ_q. Notice that Equation (20) can only be estimated for borrowers who refinance under HARP, because the available data include both the original home value and the new home value exclusively for the HARP takers and not for other borrowers. However, this subsample of borrowers with HARP refinances is highly selective, so a direct OLS estimation of Equation (20) using this subsample would yield biased results for the whole sample of borrowers. We tackle this problem by applying a two-step Heckman selection model. The first step involves analyzing the choice to opt for HARP refinancing. The set of variables in the first-stage regression that are excluded from the main regression model include borrower characteristics (FICO, income, whether first-time home buyer) and loan characteristics (interest rate, principal,

[19] We assume a linear trend for $\log(h_t^m)$ in the VAR(1) estimation.

LTV, insurance percentage, etc). Section A.2 provides further details of the estimation procedure and results.

Multiplier We impute the multiplier M in a lender's profit function using the predicted mortgage duration based on borrower characteristics. We estimate a log-normal survival model using the monthly performance data on the original mortgages in the sample. Covariates include borrower characteristics (FICO and income), loan characteristics (interest rate, LTV, principal), market fixed effects, and cohort fixed effects. Results are presented in Section A.1.

Expected Put-Back Cost We decompose the expected put-back cost P_j as the product of the put-back probability p^{PB} and a cost parameter P_{cost}:

$$P_{ijt} = P_{cost} \cdot p^{PB}(X_i, t; Incumbent_j, Post_t). \tag{21}$$

We use the estimated logit model of the put-back probability in Section 3.3.1 as $p^{PB}(\cdot)$, where covariates of the logit regression include borrower characteristics X_i, the duration of the original mortgage (i.e., period t), $Incumbent_j \times (1 - Post_t)$ dummy, and $Post_t$ dummy, where $Incumbent_j = \mathbb{1}\{j = 0\}$ and $Post_t = \mathbb{1}\{year_t \geq 2013\}$. The cost parameter P_{cost} is left for structural estimation. This specification implies that

$$\Delta_P = P_{cost} \cdot \big[p^{PB}(X_i, t; Incumbent_j = 0, Post_t) - p^{PB}(X_i, t; Incumbent_j = 1, Post_t)\big].$$

Note that $\Delta_P = 0$ if $Post_t = 1$ since the interaction term is only present in the pre-2013 period.

5.2 Identification

The set of model parameters to estimate, θ, include: (i) parameters in the search cost distribution: $\bar{\kappa}$, e, (ii) supply-side parameters: P_{cost}, σ_ω, Δ_F, (iii) the fixed-cost parameter, ϕ_0, (iv) the fixed effect of refinancing, μ, (v) the scale parameter in the distribution of utility shocks to refinancing decisions, σ_V, and (vi) the probability of choosing HARP for qualified borrowers during HARP 1.0 and 2.0, respectively: ξ_1, ξ_2.

We start with supply-side parameters. First, notice that the incumbents' market share among HARP borrowers is the sum of two components: the share of non-searchers and the share of searchers multiplied by the incumbent's winning probability in the competition stage, p_0^W. Thus the level of p_0^W is implied from the observed fraction of searchers and the incumbents' market share. From Equation (7), p_0^W is determined by Δ/σ_ω given the observable J. In the post-2013 period, $\Delta = \Delta_F$. Therefore, we can pin down Δ_F/σ_ω using the post-2013 market share of the incumbent and the fraction of searchers. Intuitively, Δ_F/σ_ω measures the post-2013 cost advantage of the incumbent lender relative to the dispersion of idiosyncratic cost shocks, and a larger advantage grants the incumbent higher market share in the competition stage.

The price of HARP refinancing from competing lenders experiences a change after 2013 because of the general reduction of put-back probability. The extent of the price change, together with the price level post-2013, helps determine the supply-side parameters. To see this, we calculate the difference between the mean HARP price offered by competing lenders post-2013 and the corresponding pre-2013 price (by Equations (8) and (9)):

$$E[r \mid j^* \neq 0, \text{post}] - E[r \mid j^* \neq 0, \text{pre}] = -\frac{\sigma_\omega}{M} \log\left(\frac{J + \exp(\Delta_F/\sigma_\omega)}{J\exp(-\Delta_P/\sigma_\omega) + \exp(\Delta_F/\sigma_\omega)}\right) + \frac{P_{0|post} - P_{0|pre}}{M} \tag{22}$$

where $P_{0|post} = P_{cost} \cdot p^{PB}(x, t; j = 0, Post = 1)$ and $P_{0|pre} = P_{cost} \cdot p^{PB}(x, t; j = 0, Post = 0)$ are the expected put-back costs for the incumbent post-2013 and pre-2013, respectively. Note that $\Delta_P = P_{cost} \cdot \left[p^{PB}(x, t; j \neq 0, Post = 0) - p^{PB}(x, t; j = 0, Post = 0)\right]$, and thus the only unknown part in $P_{0|post}$, $P_{0|pre}$, and Δ_P is P_{cost}. Therefore, given Δ_F/σ_ω, this price change is determined by two parameters: σ_ω and P_{cost}. The two parameters have opposing effects on the equation: P_{cost} drives up the price change while σ_ω mitigates it. In terms of magnitude, we expect P_{cost} to have a more pronounced effect on the pre-post price change, while the influence of σ_ω is more nuanced, given that it is divided by the multiplier, M. The two parameters also jointly determine the competing lender's average HARP price post-2013:

$$E[r \mid j^* \neq 0, \text{post}] = -\frac{\sigma_\omega}{M}\left[\log(J\exp(-\Delta_F/\sigma_\omega) + 1) + \frac{J\log\left(1 - (1 - p_0^W)/J\right)}{1 - p_0^W}\right] + \frac{P_{0|post}}{M} + g + c \tag{23}$$

Thus, P_{cost} and σ_ω are simultaneously determined by Equations (22) and (23), and therefore,

Δ_F.

Parameters in the search cost distribution, $\bar{\kappa}$ and e, are then determined through the incumbent's prices and the fraction of searchers. Let j^o denote the observed lender, with $j^o = 0$ indicating a refinancing with the incumbent. The expected price for HARP refinancing with the incumbent lender is a linear combination of the initial quote and the conditional expectation of competitive offer, weighted by the search probability:

$$E[r \mid j^o = 0] = \Pr(S = 0)r^I + \Pr(S = 1)E[r \mid j^* = 0], \tag{24}$$

where the expected competitive offer $E[r \mid j^* = 0]$ is pinned down by supply-side parameters. As mentioned in Section 4.2.3, $\bar{\kappa}$ and e govern both the initial quote r^I and search probability. Therefore, the system of the two equations determining $E[r \mid j^o = 0]$ and $P(S = 1)$ (Equations (24) and (15)) pins down the two unknown parameters ($\bar{\kappa}$ and e).

The fixed cost parameter ϕ_0 governs the increase in refinance activity in response to HARP 2.0. The fixed cost of HARP 1.0 is ϕ_0, but it reduces to zero during HARP 2.0, which induces more refinancing activity in the era of HARP 2.0. The magnitude of such increase helps to identify ϕ_0. On the other hand, the unobserved utility effect of refinancing, μ, is constant over time, which can be pinned down by the overall rate of refinancing. For example, μ tends to be positive if the predicted refinancing rate based on the calculated monetary value functions is lower than the observed level, suggesting the presence of unobserved utility gain from refinancing.

The scale parameter of utility shocks to refinancing cost, σ_V, is identified by the cross-sectional variation in refinancing decisions across different markets. σ_V controls the sensitivity of refinancing decisions with respect to the value of refinancing relative to the value of waiting, which is lower if the current LTV is high but it is expected to decline as house prices in the market gradually recover from the crisis. During 2009–2011, the recovery of house prices took different trajectories in different states. If σ_V is small, the timing of refinancing decisions would exhibit significant variation across different states. Specifically, states with a faster recovery of house prices would have more refinancing activities later in that period, compared to states with a slower recovery path. Conversely, if σ_V is high, refinancing decisions are not sensitive to the calculations of future LTV changes, and there would be less variation in terms of refinancing decisions across different states. This suggests that σ_V plays a crucial role in the heterogeneity of refinancing decisions across states.

Finally, under the assumption that the decision of refinancing type is made after the refinancing decision, the relative share of HARP refinancing compared to regular refinancing during HARP 1.0 and HARP 2.0 identifies ξ_1 and ξ_2, respectively.

5.3 Likelihood Function

Let d_t denote the refinance decision of a borrower in period t, where $d_t = 1$ stands for refinancing and $d_t = 0$ otherwise.[20] Conditional on refinancing, we observe the type of refinancing, $k = H$ or R. Conditional on HARP refinancing, we further know whether it is with the incumbent ($j^o = 0$) or another competing lender ($j^o \neq 0$). Therefore, the observed action of the borrower in period t, denoted as a_t, falls into one of the four cases:[21]

$$a_t = \begin{cases} 0, & \text{if } d_t = 0, \\ 1, & \text{if } d_t = 1, k = H, j^o = 0, \\ 2, & \text{if } d_t = 1, k = H, j^o \neq 0, \\ 3, & \text{if } d_t = 1, k = R. \end{cases} \tag{25}$$

The probability of non-refinance is given by (omitting the state variable (z, q)):

$$\Pr(a_t = 0) = \Pr(d_t = 0) = \frac{1}{1 + \exp\Big(\Big(V_t^{refi} + \mu - V_t^{wait}\Big)/\sigma_V\Big)}, \tag{26}$$

Using function $I = \mathbb{1}\{80\% < \text{LTV} < \text{cap}\}$ as an indicator for HARP eligibility, the probability of choosing HARP refinancing with the incumbent can be written as:

$$\begin{aligned} \Pr(a_t = 1) &= \Pr(d_t = 1) \Pr(k = H) \Pr(j^o = 0) \\ &= \xi I \Pr(d_t = 1) \left[\Pr(S = 0) + \Pr(S = 1) \Pr(j^* = 0 \,|\, S = 1)\right]. \end{aligned} \tag{27}$$

[20]We interpret all observed prepayment as refinancing activities, although in reality it could also include prepayment for reasons other than refinancing, such as moving.

[21]This classification is for non-default borrowers. Borrowers who end up in default are not used for likelihood estimation because their likelihood contribution is determined by parameters governing the transition of market-level variables and parameters in the survival model, which do not change with structural parameters.

Similarly, for the other two cases,

$$\Pr(a_t = 2) = \xi I \Pr(d_t = 1) \Pr(S = 1) \Pr(j^* \neq 0 \,|\, S = 1), \tag{28}$$

$$\Pr(a_t = 3) = (1 - \xi I) \Pr(d_t = 1). \tag{29}$$

Adding back the borrower index $i = 1, ..., N$, the observed outcomes for borrower i, O_i, is the collection of actions and realized macro state variables from the first year after mortgage origination ($t = 1$) to the last year that the borrower appears in the sample, T^i:

$$O_i = \left\{ a_t^{(i)}, z_t^{(i)} \right\}_{t=1}^{T^i}. \tag{30}$$

Note that T^i indicates the year of refinancing if the borrower ever refinances, otherwise it corresponds to the last year of the sample. Given the model parameters θ, the likelihood of the observed outcomes for borrower i conditional on the initial state $z_0^{(i)}$ is:

$$L\left(O_i \,\middle|\, z_0^{(i)}, \theta\right) = \prod_{t=1}^{T^i} \Pr\left(z_t^{(i)} \,\middle|\, z_{t-1}^{(i)}\right) \int \Pr\left(a_t^{(i)} \,\middle|\, z_t^{(i)}, q\right) d\Phi(q/\sigma_q) \tag{31}$$

where $\Phi(\cdot)$ is the standard normal distribution function.

The model also predicts refinancing prices, but these are only observable in the data for HARP refinancing. A natural method is to compute the likelihood of observed prices for HARP borrowers and incorporate this into the likelihood function. Consequently, the likelihood contribution of prices is solely from those opting for HARP refinancing, representing a mere 8.8% of our sample borrowers. The absence of price data for the majority of borrowers significantly constrains the role of price information in the estimation of model parameters, particularly those on the supply side. Despite attempts to use this method, it failed to yield reasonable estimates, leading us to adopt an alternative estimation procedure.

Following Allen et al. (2019), we use a quasi-likelihood estimator that incorporates a set of auxiliary moments in addition to the likelihood function. The set of moments we use, $m(\theta)$, includes four price moments and one aggregate moment on search efforts from an external source, NSMO. The four price moments come from four groups, respectively: (1) HARP refinancing with the incumbent lender prior to 2013, (2) HARP refinancing with the incumbent lender post 2013, (3) HARP refinancing with a competing lender prior to 2013, and (4) HARP refinancing with a competing lender post 2013. For each group, we calculate the expected HARP price from the model and obtain its distance from the sample average. For the aggregate moment on search effort, we use the model to calculate the average search probability for those with either HARP or regular refinancing. The analog probability from the survey is calculated as the fraction of borrowers who search more than one lender when refinancing their mortgage. The difference between the two is a mean-zero error under the null hypothesis that the model is correctly specified. Using the variance of data moments as weighting matrix $\hat{W}$, we construct the following aggregate log likelihood function:[22]

$$\max_{\theta} \sum_{i=1}^{N} \log L\left(O_i \,\middle|\, z_0^{(i)}, \theta\right) - m(\theta)^T \hat{W}^{-1} m(\theta) \tag{32}$$

In our computation of likelihood function in Equation (31), the integral over q is numerically approximated. It is important to note that directly drawing from $N(0, \sigma_q)$ is problematic in our setting because it might fail to rationalize some observed HARP refinancing decisions. Specifically, when draws of q are too centered around zero, the predicted LTV can fall below 80% for borrowers that actually choose HARP refinancing, thus being directly rejected by data. To provide enough coverage, we use Halton draws from the an auxiliary distribution (which is also a normal distribution) and use importance sampling to reweight the draws. The auxiliary distribution is chosen to rationalize all observed HARP refinances.

6 Estimation Results

6.1 Parameter Estimates

Table 4 summarizes the parameter estimates, with the standard errors enclosed in parentheses. The monetary values that directly enter the borrower's value functions, including $\bar{\kappa}$, ϕ_0, μ, and σ_V, are expressed in units of \$1,000. Supply-side parameters, including P_{cocst}, σ_ω, and Δ_F, are expressed on a per-hundred-dollar basis of the mortgage.

The search cost is on average \$5,281, ranging from $\$5,281 \times (1 - 0.496) = \$2,662$ to $\$5,281 \times (1 + 0.496) = \$7,900$. Since this is the search cost over a borrower's lifetime,

[22]See Allen et al. (2019) for more discussion on the performance of this estimation approach.

Table 4: Maximum Likelihood Estimation Results

$\bar{\kappa}$	e	ϕ_0	μ	σ_V
5.281	0.496	43.675	53.117	152.016
(0.008)	(0.001)	(0.543)	(0.565)	(0.874)

P_{cost}	σ_ω	Δ_F	ξ_1	ξ_2
969.476	3.880	-0.870	0.385	0.895
(7.789)	(0.009)	(0.003)	(0.003)	(0.006)

The first row of each table shows the estimates of the model parameters, and the second row represents the corresponding standard error for each parameter. $\bar{\kappa}$, ϕ_0, μ, and σ_V are in units of \$1,000, while P_{cocst}, σ_ω, and Δ_F are expressed on a per-hundred-dollar basis of the mortgage.

our estimate is significantly higher than the estimate of average search cost from Allen et al. (2019), where the search costs are expressed over the five-year term of the mortgage contract. Although the search cost estimates are nominally large, they represent on average only 2.67% of total interest cost over the entire horizon of the contracts. This is close to the estimate of 2.5% from Allen et al. (2019).

The fixed cost of regular refinancing is \$43,569, not including the private mortgage insurance required for loans with LTVs exceeding 80%. On average, borrowers in our sample pay about \$41,081 as the fixed cost of refinancing, including HARP 2.0 borrowers who pay zero fixed costs. However, the high fixed cost is offset by the benefit of refinancing from other sources, which is estimated to be \$53,117. Consequently, refinancing delivers a net benefit equivalent to about \$12,036. This estimate reflects all other benefits of refinancing in addition to the channels captured by the model. These include a shorter loan term, the opportunity to cash out for home renovations or debt consolidation, an improved credit score if the borrower struggles to meet payments on the exiting mortgage.

How does the net benefit of refinancing translate into refinancing probabilities? Suppose refinancing yields the same (discounted) sum of future consumption as not refinancing in a given period. Then the difference between the value of refinancing and not refinancing is given by $\mu - \phi$, with an estimated value of \$12,036. Based on the estimated value of σ_V, the probability of refinancing in this case is $1 - \Pr(a = 0) = 0.52$, according to Equation (26). In other words, a borrower is almost indifferent between refinancing and not refinancing when the two options lead to same present value of future consumption, because the net benefit of refinancing is mostly muted by the large variance of idiosyncratic utility shock to refinancing decisions. The high estimated value of σ_V might be explained by borrowers who need to move to a new house. Our data show whether an existing mortgage is prepaid, but not the information on the new mortgage if prepaid, unless the reason of prepayment is due to HARP refinancing. Thus, we generally cannot differentiate those who prepay due to moving and those who prepay due to refinancing the same house. Since moving is mainly driven by idiosyncratic factors like labor market matching frictions, its presence adds significant noise to the observed refinancing decision, thus leading to a high estimate of σ_V.

On the supply side, we find that a mortgage put-back is highly costly for a mortgage lender, according to the estimate of P_{cost}. Based on this estimate, we calculate the incumbent lender's expected put-back cost, $P_0 = P_{cost} \cdot p^{PB}$. For every \$100 of the mortgage, the pre-2013 expected put-back cost for the incumbent, $P_{0|pre}$, is \$2.976 on average. A competing lender, on the other hand, has an expected put-back cost that is \$4.928 higher than the incumbent, marking a 160% difference. The share of the expected put-back cost in a competing lender's total cost is about 18.8%, compared to 8.3% for the incumbent lenders. The asymmetry in put-back cost before 2013 dwarfs other cost differentials between the competing and incumbent lenders, Δ_F, which is less than \$1 per \$100 mortgage. Therefore, the differential exposure to put-back risk is substantial, and it constitutes the main source of the cost advantage prior to 2013. The policy change in 2013 leads to a dramatic decrease in the expected put-back cost, to an average of \$0.105. It qualitatively changed the role of put-back risk in a lender's profit function, contributing to a lower price observed in the data.

Our estimate of σ_ω implies a standard deviation of \$4.977 ($= 3.88\pi/\sqrt{6}$) for the idiosyncratic cost shock in the competition stage. This has important implications for our understanding of the importance of cost advantage in this market. In the absence of any systematic cost difference (i.e., $\Delta = 0$), our estimate of σ_ω implies that the average difference between $\omega_{(2)}$ and $\omega_{(1)}$ is \$5.379 in a duopoly market and \$4.72 with three lenders. With a systematic cost difference $\Delta = \Delta_P + \Delta_F = 4.928 - 0.87 = 4.058$, the incumbent lender's winning probability in the competition stage is 0.74, compared to 0.26 for the competing lender in a duopoly market. In a market with three lenders, it is 0.59 for the incumbent and 0.21 for the two competing lenders. This suggests that the systematic cost advantage between the incumbent and competing lenders is a more important source of market power than the idiosyncratic cost differences.

Finally, the estimated ξ_1 suggests that during the first phase of HARP, an eligible borrower takes up HARP with a 38.5% chance. This reflects poor borrower knowledge and understanding of HARP, as reported in a mid-program assessment by FHFA. The assessment pointed out three potential reasons. First, many borrowers were not aware of the program due to a lack of advertising and information campaigns. Second, borrowers may have heard of the program but confused the program with other government housing programs initiated during that time. Third, many eligible borrowers were under the mistaken impression that they were ineligible for HARP because of a lack of clarity and transparency around the program rules. An important factor contributing to the lack of borrower awareness of the program was borrower outreach. During HARP 1.0, lenders were prohibited from directly soliciting borrowers with HARP refinancing (HARP mid-program assessment). As a result, eligible borrowers may have missed the opportunity to learn about the program through lenders.

During the second phase of HARP, the take-up rate witnessed a significant increase to 89.5%. This increase is closely linked to the implementation of a nationwide public education campaign to improve borrower knowledge of the program. The solicitation guidelines for HARP loans were also revised to increase borrower outreach. Our result suggests that these measures during HARP 2.0 were effective at boosting the take-up rate of HARP.

6.2 Model Fit

This section provides a comparison between the model prediction and the observed data to assess the goodness of fit of the baseline model. We start by simulating the model $Ns = 100$ times for each borrower in the data ($i = 1, ..., 21247$). For each borrower i, we solve the model to find the refinancing probability and eligibility for HARP in each state from $t = 1, ..., T$. Then in each simulation of the borrower, we simulate the default outcome and the path of state variables (z, q) for $t = 1, ..., T$. Based on the simulated path of state variables, we then simulate the refinancing decision. If refinancing occurs, we then draw the refinancing type based on the eligibility for HARP. Next we draw the search type and find the search decision. If searching, we then draw the winner of the competition stage and find the expected price. We also re-calculate the probability of default after refinancing and simulate the default

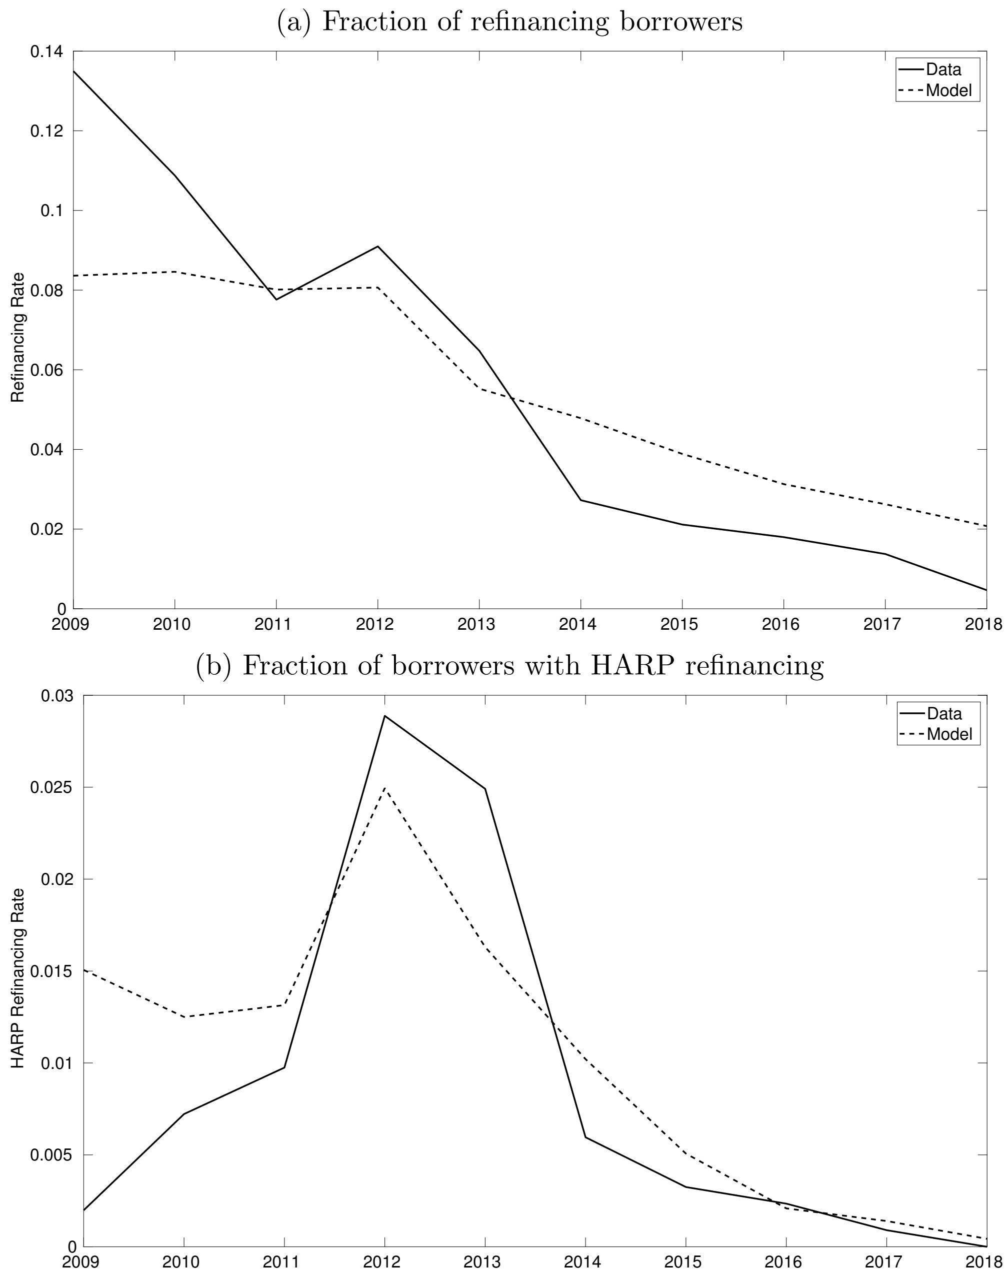

Figure 5: Model Fit

This figure shows the model predictions of refinancing decisions compared with the data. Panel (a) shows the fraction of refinancing borrowers in each year from 2009 to 2018, calculated as the number of borrowers who refinanced in a given year divided by the total number of borrowers. Panel (b) shows the fraction of borrowers with HARP refinancing among those whose initial LTV of the previous mortgage exceeded 80%. Within this subsample, the fraction is calculated as the number of borrowers choosing HARP refinancing in a given year divided by the total number of borrowers in the subsample.

outcome on the refinanced mortgage.

Using the simulated data, we calculate the fraction of borrowers who refinance in each year from 2009 to 2018 and compare this with the fraction calculated from the data, as depicted in Figure 5a. Although the model struggles to match the high refinancing uptake in 2009, it successfully mirrors the overall downward trend of refinancing activities, particularly the sharp decline after 2012. In general, the model's prediction of refinancing rate over time is smoother than the observed rates, yet the general trend aligns with the empirical pattern.

We then compare the refinancing rate for HARP, focusing on a subsample of borrowers with initial LTVs over 80%. This subset of borrowers is more prone to distress, making them the prime target for the program. The HARP refinancing rate is calculated by dividing the number of borrowers who opt for HARP refinancing each year by the total number of borrowers in the subsample. Figure 5b shows the comparison between the model prediction and the actual data. The model's prediction of HARP refinancing rate is higher than the data in 2009. Nevertheless, the model accurately captures the significant uptake of HARP since 2012 and the program's gradual decline during its latter half.

The average pre-2013 interest rate on HARP refinancing in the simulated data is 4.95%, compared to 4.51% in the data. In the post-2013 period, the model predicts an average interest rate of 4.28% for HARP refinancing compared to 4.11% in the data. About 51% of borrowers search when refinancing, close to the 49% from the NSMO survey. Among searchers, about 39% still choose the incumbent in the competition stage.

7 Counterfactual

Given the estimated model parameters, we conduct a series of counterfactual exercises to evaluate the effect of the asymmetric put-back risk exposure on the extensive and intensive margins of refinancing activities, compared with welfare effects from other sources. In Section 7.1 it is compared with a general risk reduction. In Section 7.2 we compare it with the welfare effect of the search friction and further explore the interaction between search friction and cost advantage. Finally, Section 7.3 examines the effects of HARP 2.0 modifications in comparison with the welfare effect of the asymmetric risk exposure.

7.1 The Effect of Asymmetric Put-Back Risk

7.1.1 Overall Effects

The asymmetry between the incumbent and competing lenders in terms of their exposure to put-back risk was removed by the new policy in 2013. In addition to the elimination of this asymmetry, the new policy also led to a general reduction in the put-back risk for every lender. In this section, we first consider the case where the risk exposure is symmetric from the beginning of HARP but the general reduction happens later in 2013. Then we consider the case where both the symmetric risk exposure and the general reduction occurs from the beginning. Specifically, we set $\Delta_P = 0$ if $year_t \geq 2009$ in the first exercise. The incumbent's expected put-back cost remains the same as the baseline model, while competing lenders now have a lower expected put-back cost than the baseline level due to the removal of differential risk exposure. In this setting with symmetric risk exposure, the incumbent lender has only the first-mover advantage but not a cost advantage. In the second exercise, we essentially move the policy change from 2013 to 2009. We refer to this one as the full reduction case.

Columns (2) and (3) in Table 5 summarize the outcome variables for the two counterfactual exercises. On the extensive margin, we calculate the overall refinancing rate as the number of borrowers who refinance before 2018 divided by the total number of borrowers. The refinancing rate increases from 86.3% to 86.6% in the case of symmetric risk exposure, and it further increases to 86.9% with a full reduction in put-back risk. Interestingly, the HARP refinancing rate hardly changes. In other words, the change in put-back policy leads to more regular refinancing activity rather than HARP refinancing. To assess the effect of the program on loan default, we calculate the 10-year default rate as the fraction of borrowers who default on their mortgage within 10 years of the loan origination. This includes both those who refinance and those who never refinance. Compared with the baseline model, there is hardly any change in the 10-year default rate with the symmetric risk exposure. With the full reduction case, the decline is 0.1% relative to the baseline model. As we zoom in on the high-LTV borrowers—that is, borrowers with initial LTV over 80%—the reduction in default rate is 0.1% in the case of symmetric risk exposure, with an additional 0.1% decline in the case of full reduction.

The interest savings from the counterfactual put-back policy is economically significant. With the symmetric exposure case, the average rate reduction through refinancing increases

Table 5: Mean of Outcome Variables from Counterfactual and Baseline Models

	Baseline	Sym. risk	Sym. risk + Risk reduction
	(1)	(2)	(3)
Refinancing rate (%)			
All	86.3	86.6	86.9
HARP	6.4	6.4	6.4
Default rate (%)			
All	7.9	7.9	7.8
High-LTV borrowers	12.2	12.1	12.0
For borrowers who refinance:			
Δr (bps)	0.7	19.9	37.7
Δ annual payment	1.0	1.3	1.5
Total payments	203.1	200.1	197.3
Borrower welfare			
All	511.1	513.2	515.1
High-LTV borrowers	387.2	390.1	392.8

This table summarizes the means of outcome variables from counterfactual scenarios (columns (2) and (3)) and the baseline model (column (1)). Column (1) shows the baseline scenario with the policy change in 2013. Column (2) assumes a partial implementation of the new policy in 2009 with a symmetric exposure to put-back risk for the incumbent and competing lenders. Column (3) assumes a full implementation of the new policy in 2009, with both symmetric exposure and general reduction in put-back risk. Refinancing rate (all) is the percentage of borrowers who refinance before 2018. HARP refinancing rate is the fraction of HARP borrowers. Default rate is measured by the percentage of borrowers who default on their mortgage within 10 years of loan origination. High-LTV borrowers are those whose initial LTV is over 80%. Borrower welfare is the discounted sum of lifetime consumption in units of $1,000. Δr is the the difference between the original interest rate and the new interest rate on the refinanced mortgage in basis points. Δ annual payment is the difference between the original annual mortgage payment and the new annual payment in units of $1,000. Total payments is the discounted sum of all mortgage payments throughout the borrower's life with a discount factor of 0.95. $N = 2124700$.

by 19.2 bps. The rate reduction increases by another 17.8 bps with the full reduction in put-back risk. This translates into a reduction in annual mortgage payments by \$300 and \$200, respectively. Accounting for the amortization period, the present value of total mortgage payments over a borrower's lifetime decreases by \$3,000 on average as the cost asymmetry is removed, with an additional decrease of \$2,800 with the general reduction in put-back risk.

Finally, we calculate the total welfare of a borrower as the discounted sum of lifetime consumption, taking into account any default outcomes and refinancing activities. The overall borrower welfare increases by \$2,100, or 0.4%, with the elimination of the asymmetric risk exposure alone, and the increase doubles with the full installation of the new policy in 2009. The effect mainly comes from the intensive margin of greater interest savings.

7.1.2 Effects on the Intensive Margin

Table 6 focuses on the subsample of borrowers who refinance, especially those with HARP refinancing. In the baseline model, HARP borrowers are significantly riskier than other borrowers who refinance, with a larger initial loan balance and LTV. The average initial loan balance and LTV become larger in the counterfactual scenarios, for both HARP borrowers and other borrowers who refinance. This suggests that the increase in overall financing activity in the counterfactual scenarios are driven by higher-risk borrowers.

How does the timing of refinancing change in the counterfactual scenarios? Figure 6a and 6b plot the refinancing rate and HARP refinancing rate during 2009–2018 in the baseline and counterfactual scenarios. Compared to the baseline model, the overall refinancing rate becomes more front-loaded in the counterfactual scenarios, with higher a refinancing rate before 2012 and a lower refinancing rate afterward. A similar pattern is also present with the HARP refinancing rate. In other words, some borrowers who refinance later in the baseline model would refinance earlier in the counterfactual scenarios with less waiting time. This is also reflected by the change in loan balance at refinancing from Table 6. In the baseline model, borrowers who refinance typically wait until the loan balance drops by 10.8% before refinancing, while in the counterfactual scenarios the average decreases in loan balance are 10.7% and 10.6%, respectively. The same pattern is shown in the subsample of HARP borrowers.

In general, borrowers who refinance do so when their LTV decreases by 13.9% from origination. However, for HARP borrowers, their LTV at the time of refinancing is typically

Table 6: Refinancing Outcomes for Borrowers Who Refinance and for HARP Borrowers

	Baseline		Sym. risk		Sym. risk + Risk reduction	
	All refi (1)	HARP (2)	All refi (3)	HARP (4)	All refi (5)	HARP (6)
At origination:						
Loan balance	171.1	187.4	171.2	186.9	171.2	186.5
LTV (%)	77.8	84.6	77.9	84.6	77.9	84.6
Default risk (%)	9.0	16.3	9.1	16.4	9.2	16.4
At the time of refinancing:						
Δ%loan balance	-10.8	-10.0	-10.7	-9.9	-10.6	-9.8
Δ LTV	-13.9	10.6	-13.8	10.5	-13.6	10.4
Housing shock (%)	0.3	-12.8	0.3	-12.7	0.3	-12.6
Δ default risk	-5.2	-10.3	-5.2	-10.2	-5.3	-10.1
Total payments	203.0	207.9	199.8	200.3	196.9	193.2
N	1985972	136003	1989461	135591	1992429	135163

This table is generated from the subsample of borrowers who refinance and the subsample of borrowers with HARP refinancing under each scenario. Loan balance, reduction in annual payment, total mortgage payments, and borrower welfare are in units of \$1,000. Change in loan balance at the time of refinancing is expressed as the percentage change since loan origination. Change in LTV is the difference between LTV at the time of refinancing and origination. Idiosyncratic housing risk is calculated as $e^q - 1$, where q is the idiosyncratic housing shock variable at the time of refinancing. Change in default risk is the difference between the new 10-year default rate after refinancing and the 10-year default rate without refinancing. Total mortgage payments is the discounted sum of all mortgage payments throughout the borrower's lifetime.

(a) Fraction of refinancing borrowers

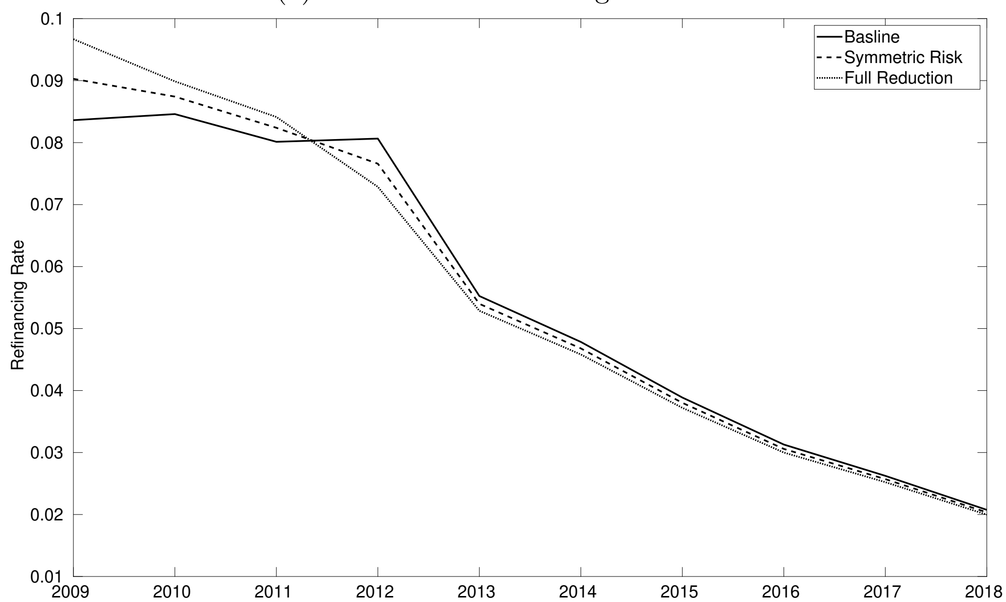

(b) Fraction of borrowers with HARP refinancing

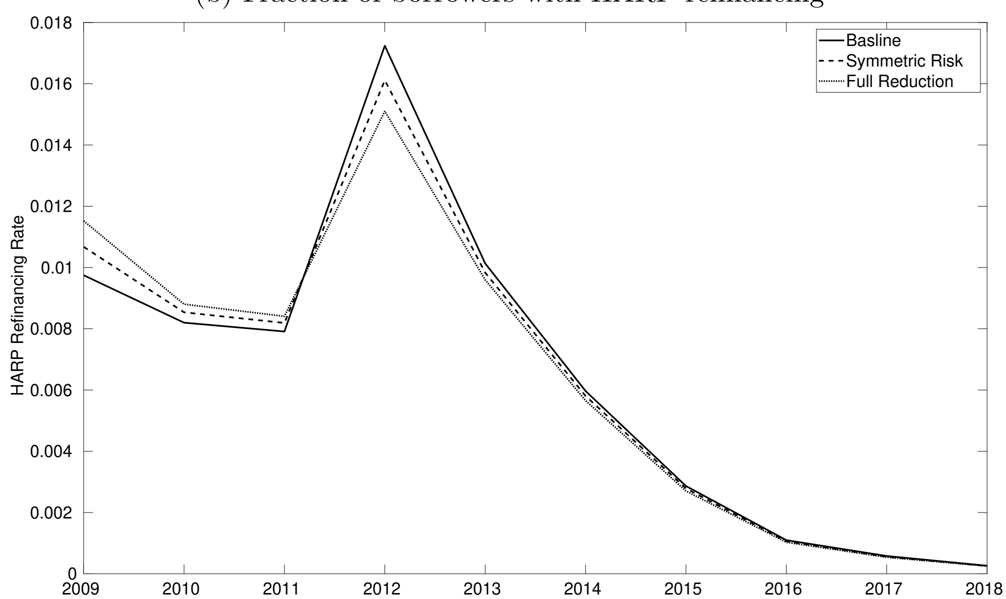

Figure 6: Timing of Refinancing Decisions from the Baseline and Counterfactual Models

This figure shows refinancing decisions from the baseline model and two counterfactual models. The solid line corresponds to the baseline model, the dashed line corresponds to the counterfactual model with symmetric exposure to put-back risk, and the dotted line corresponds to the counterfactual model with a full reduction in put-back risk. Panel (a) shows the fraction of refinancing borrowers in each year from 2009 to 2018, calculated as the number of borrowers who refinance in a given year divided by the total number of borrowers. Panel (b) shows the fraction of borrowers with HARP refinancing.

higher than the initial condition by an average of 10.6% in the baseline model. This is due to their adversarial individual housing condition: On average, their house price is 12.8% lower than the market average. In the counterfactual scenarios, a HARP borrower's LTV is still higher than her initial LTV, but the difference is slightly smaller compared to the baseline model.

The effect of refinancing on default risk is also divergent between HARP borrowers and other borrowers who refinance. For borrowers who refinance, their default risk of the new mortgage is generally lower, and the reduction is slightly larger in the counterfactual scenarios. For HARP borrowers, the risk reduction effect is twice as large (10.3% versus 5.2%), but becomes smaller in the counterfactual scenarios (10.1% versus 10.3%). This can be explained by the relatively larger loan balance at the time of refinancing in the counterfactual scenarios compared to the baseline.

The total mortgage payments for HARP borrowers decrease more with the removal of cost asymmetry than other borrowers. HARP borrowers' total mortgage payments decrease by $7,600 with the removal of cost asymmetry, with an additional $7,100 reduction with the reduction in put-back risk. These effects are twice as large for the average borrower with refinancing activities.

In sum, we find that if the new policy on mortgage put-back were implemented in 2009 instead of 2013, 0.6% more borrowers would have refinanced before 2018, and the timing of refinancing decisions would be earlier. Although HARP take-up rate hardly change compared to the baseline, HARP borrowers benefit more from the program due to a 7.1% decrease in total mortgage payments over their lifetime. Eliminating the incumbent-competing differential in put-back risk alone can achieve about half of the total benefits. This is despite the fact that the marginal effect of the incumbent-competing differential on put-back probabilities is only less than half of the general reduction, as we find in Section 3.3.1.

7.2 Search Friction and Cost Advantage

The asymmetric risk exposure to mortgage put-back leads to welfare loss not just because of higher average cost, but more importantly the competitive frictions associated with it. Our results also show considerate search friction in this market, in which the incumbent lenders have a first-mover advantage. How does the first-mover advantage interact with the

Table 7: Mean of Outcome Variables from Counterfactual and Baseline Models

	Baseline (1)	No search friction (2)	Sym. risk (3)	No search friction + Sym. risk (4)
Refinancing rate (%)				
All	86.3	86.6	86.6	86.9
HARP	6.4	6.4	6.4	6.4
Default rate (%)				
All	7.9	7.9	7.9	7.9
High-LTV borrowers	12.2	12.2	12.1	12.1
For borrowers who refinance:				
Δr (bps)	0.7	20.0	19.9	37.9
Δ annual payment	1.0	1.2	1.3	1.4
Total payments	203.1	201.3	200.1	198.5
Borrower welfare				
All	511.1	512.6	513.2	514.5
High-LTV borrowers	387.2	388.8	390.1	391.4

This table summarizes the means of outcome variables from the counterfactual scenarios (columns (2)–(4)) and the baseline model (column (1)). Column (2) assumes no search friction in the counterfactual scenario. Column (3) assumes symmetric exposure to put-back risk since 2009. Column (4) assumes no search friction and symmetric exposure. The definition of the variables is the same as in Table 5. $N = 2124700$.

cost advantage? Does one exacerbate the other? We conduct two additional counterfactual experiments to answer these questions.

The first counterfactual experiment shuts down the search friction by removing the incumbent lender's first-mover advantage. In this case, interest rates are generated directly from an English auction where lenders have potentially heterogeneous costs. Column (2) of Table 7 summarizes the outcomes from this experiment. Search friction hardly changes the external margin or default rate, but its effects on the intensive margin is economically significant. For borrowers who refinance, the absence of search friction boosts interest savings by about 13.3 bps, or \$200 in annual mortgage payments. Over a borrower's lifetime, it helps to save \$1,800 on mortgage payments in terms of present value. Overall, borrower welfare increases by \$1,500 in the absence of search cost, which is smaller than the welfare increase associated with the removal of asymmetric risk exposure.

To evaluate the change in welfare when the risk asymmetry is removed in an environment without search friction, we conduct another counterfactual experiment in which either search friction and risk asymmetry are present. This is shown in column (4) of Table 7. The external margin increase is 0.6%, with a 0.1% reduction in default risks for high-LTV borrowers compared to the baseline model. On the intensive margin, the average interest savings of this case is double that of the previous case with no search friction. The overall welfare effect is \$3,400, with a higher effect for high-LTV borrowers at \$4,200.

By comparing column (4) with column (2), we find that the welfare implication of the risk asymmetry in absence of search friction is \$1,900, lower than the welfare effect in an environment with search friction. In other words, the presence of search friction exacerbated the welfare loss from the risk asymmetry. Notice that the opposite is also true. The risk asymmetry also aggravates the inefficiencies from search friction. Therefore, the overall market power of the incumbent is not a simple sum of the two sources; they interact and amplify the individual effects.

7.3 The Effect of HARP 2.0

As another interesting comparison, we evaluate the effectiveness of HARP and the subsequent modifications to HARP (HARP 2.0). Table 8 presents a summary of the key outcome variables in a series of counterfactual scenarios where we remove the entire HARP 2.0 modi-

Table 8: Mean of Outcome Variables from Counterfactual and Baseline Models

	Counterfactual				Baseline
	HARP 1.0	HARP 1.0 + partial HARP 2.0			HARP 1.0+ HARP 2.0
		ϕ^H	ξ	LTV cap	
	(1)	(2)	(3)	(4)	(5)
Refinancing rate (%)					
All	86.1	86.2	86.1	86.1	86.3
HARP	3.9	4.0	5.7	4.0	6.4
Default rate (%)					
All	8.0	8.0	8.0	8.0	7.9
High-LTV borrowers	12.3	12.3	12.3	12.3	12.2
For borrowers who refinance:					
Δr (bps)	0.4	0.5	0.5	0.4	0.7
Δ annual payment	1.0	1.0	1.0	1.0	1.0
Total payments	203.2	203.1	203.2	203.2	203.1
Borrower welfare					
All	509.9	510.3	509.9	509.9	511.1
High-LTV borrowers	385.4	386.1	385.4	385.4	387.2

This table summarizes the average outcome variables from counterfactual scenarios (columns (1)–(4)) and the baseline model (column (5)). Column (1) corresponds to the case with only HARP 1.0 throughout 2009 to 2018, respectively. Columns (2)–(4) show the scenario where HARP 1.0 is implemented through 2009 to 2011, followed by changes in the fixed cost (ϕ^H), awareness (ξ), and LTV cap, respectively. Column (5) shows the baseline model with HARP 1.0 during 2009–2011 and HARP 2.0 afterwards, with changes in all three above-mentioned variables. Definitions of the variables are the same as in Table 5. $N = 2124700$.

fications or specific measures in the HARP 2.0 modification. In column (1), only HARP 1.0 is available throughout the 2009–2018 period. In columns (2)–(4), HARP 1.0 is implemented during the initial phase (2009–2011) followed by only one modification to a certain aspect of the program. Column (2) shows the case where the modification targets only the fixed cost of refinancing by setting ϕ^H to zero, while other aspects are left unchanged. Column (3) corresponds to the case where only the program awareness is changed in the second phase of the program. In column (4) the only modification is eliminating the LTV cap requirement. Lastly, column (5) provides the baseline case where HARP 2.0 encompasses all three measures.

In terms of extensive margin, HARP 2.0 leads to a 0.2% increase in the overall refinancing rate, with the change in fixed cost contributing the most to the increase in the overall refinancing rate. The HARP refinancing rate is 3.9% without the HARP 2.0 modifications. Given the baseline HARP refinancing rate of 6.4%, HARP 2.0 raises the utilization of the program by 64% (6.401-3.904)/3.904), with the awareness of HARP as the main contributor.

In the absence of HARP 2.0, the average 10-year default rate would be 0.1% higher. In terms of the interest cost of borrowing, we find that HARP 2.0 only had a marginal effect in reducing total mortgage payments. Overall, HARP 2.0 boosts the average borrower welfare by $1,200. For borrowers with an initial LTV over 80%, the welfare effect is $1,800, or 0.5%. The reduction of fixed effects plays the most prominent role among the three factors, accounting for 40% of the effect.

Although economically meaningful, the welfare effect of HARP 2.0 modifications is still smaller than that of the cost asymmetry from the put-back risk. Note that the welfare impact of HARP 2.0 modifications comes from other channels rather than interest savings, namely reduced fixed costs, higher refinancing rates, and lower default rates. By comparison, the welfare implications of the cost asymmetry is mostly from the intensive margin of greater interest savings and earlier refinance timing.

8 Conclusion

This paper quantifies the welfare implications associated with the incumbent cost advantage related to HARP. The existence of this advantage is not intentional; however, it leads to non-negligible welfare effects, larger than the inefficiencies from the more well-known market friction, search friction. The potential welfare gain, if it was corrected earlier, is comparable to the large-scale program enhancements made under HARP 2.0. This is because the program-granted advantage interacts with pre-existing market frictions, exacerbating the incumbent's market power. This leads to more surplus extracted by the incumbent rather than flowing to the borrowers, impeding the pass-through of the program's benefits.

Insights from this paper apply to programs beyond HARP. Anti-trust is not the only way for the government to promote competition in the financial sector; instead, the consideration of competition needs to be present in the design of relevant policies that apply directly to financial intermediaries. It is important to understand the pre-existing competitive frictions in the market and analyze whether a certain policy or program contains rules that explicitly or implicitly treat participants unequally and therefore give rise to advantageous positions to some over others. These distributional effects of policy are not always second-order, especially when the final goal of the program is to reach financial consumers.

References

Joshua Abel and Andreas Fuster. How do mortgage refinances affect debt, default, and spending? Evidence from HARP. *Evidence from HARP (August 5, 2019). FRB of New York Staff Report*, (841), 2019.

Sumit Agarwal, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru. Policy intervention in debt renegotiation: Evidence from the Home Affordable Modification Program. *Journal of Political Economy*, 125(3):654–712, 2017.

Sumit Agarwal, John Grigsby, Ali Hortaçsu, Gregor Matvos, Amit Seru, and Vincent Yao. Searching for approval. Technical report, National Bureau of Economic Research, 2020.

Sumit Agarwal, Gene Amromin, Souphala Chomsisengphet, Tim Landvoigt, Tomasz Piskorski, Amit Seru, and Vincent Yao. Mortgage refinancing, consumer spending, and competition: Evidence from the home affordable refinance program. *The Review of Economic Studies*, 90(2):499–537, 2023.

Alexei Alexandrov and Sergei Koulayev. No shopping in the US mortgage market: Direct and strategic effects of providing information. *Consumer Financial Protection Bureau Office of Research Working Paper*, (2017-01), 2018.

Jason Allen and Shaoteng Li. Dynamic competition in negotiated price markets. Technical report, Bank of Canada, 2020.

Jason Allen, Robert Clark, and Jean-François Houde. Search frictions and market power in negotiated-price markets. *Journal of Political Economy*, 127(4):1550–1598, 2019.

Sumedh Ambokar and Kian Samaee. Inaction, search costs, and market power in the US mortgage market. 2019.

Gene Amromin and Caitlin Kearns. Access to refinancing and mortgage interest rates: Harping on the importance of competition. 2014.

Neil Bhutta, Andreas Fuster, and Aurel Hizmo. Paying too much? Price dispersion in the US mortgage market. 2020.

Lance Brannman and Luke M Froeb. Mergers, cartels, set-asides, and bidding preferences in asymmetric oral auctions. *Review of Economics and Statistics*, 82(2):283–290, 2000.

Anthony A DeFusco and John Mondragon. No job, no money, no refi: Frictions to refinancing in a recession. *The Journal of Finance*, 75(5):2327–2376, 2020.

Leland E. Farmer and Alexis Akira Toda. Discretizing nonlinear, non-Gaussian Markov processes with exact conditional moments. *Quantitative Economics*, 8(2):651–683, 2017. doi: https://doi.org/10.3982/QE737. URL `https://onlinelibrary.wiley.com/doi/abs/10.3982/QE737`.

Andreas Fuster, Laurie S Goodman, David O Lucca, Laurel Madar, Linsey Molloy, and Paul Willen. The rising gap between primary and secondary mortgage rates. *Economic Policy Review*, 19(2), 2013.

Umit G Gurun, Gregor Matvos, and Amit Seru. Advertising expensive mortgages. *The Journal of Finance*, 71(5):2371–2416, 2016.

Elisabeth Honka. Quantifying search and switching costs in the US auto insurance industry. *The RAND Journal of Economics*, 45(4):847–884, 2014.

Benjamin J Keys, Devin G Pope, and Jaren C Pope. Failure to refinance. *Journal of Financial Economics*, 122(3):482–499, 2016.

Doug McManus, Liyi Liu, and Mingzhe Yi. Why are consumers leaving money on the table? *Freddie Mac Economic & Housing Research Insight, http://www. freddiemac. com/research/insight/20180417_ consumers_ leaving_ money. page*, 2018.

Tomasz Piskorski, Amit Seru, and Vikrant Vig. Securitization and distressed loan renegotiation: Evidence from the subprime mortgage crisis. *Journal of Financial Economics*, 97 (3):369–397, 2010.

David Scharfstein and Adi Sunderam. Market power in mortgage lending and the transmission of monetary policy. *Unpublished working paper. Harvard University*, 2016.

Susan E Woodward and Robert E Hall. Diagnosing consumer confusion and sub-optimal shopping effort: Theory and mortgage-market evidence. *American Economic Review*, 102 (7):3249–3276, 2012.

Appendix

A.1 Survival Analysis

Table 9: Survival Analysis

	(1)	(2)
Exit Event:	Default	Default or Prepay
Model:	Log-logistic	Log-normal
FICO	0.003***	−0.000***
	(0.000)	(0.000)
LTV	−0.014***	0.003***
	(0.000)	(0.000)
Interest Rate	−0.316***	−0.328***
	(0.004)	(0.002)
log(Balance)	−0.308***	−0.254***
	(0.005)	(0.002)
log(Income)	0.182***	0.013***
	(0.004)	(0.001)
Market FE	Yes	Yes
Year FE	Yes	Yes
Observations	2,079,763	2,079,763

Column (1) reports the results of survival probability where the exit event is default using a log-logistic model, while column (2) reports the results of survival probability where the exit event is either default or prepay using a log-normal model. The figures in parentheses are standard errors, with 1, 2, and 3 asterisks indicating statistical significance at 10%, 5%, and 1%, respectively.

A.2 Idiosyncratic Housing Shock

In the first stage, we estimate the binary decision to take HARP refinancing using a probit model. This stage contains borrower and loan characteristics (FICO, income, interest rate, principal, LTV, whether first-time buyer, insurance coverage, occupancy type, number of borrowers) that affect their refinance decision but should not affect the house value (exclusion restriction). Only for those who choose to take HARP refinancing in the first stage do we observe their new home value at the time of refinance, and thus Δh_{it}. The main regression in Equation (20) is estimated in the second stage. Table 10 shows the regression results from both stages.

Table 10: Idiosyncratic Housing Shock

	(1)		(2)
	First Stage HARP Refinance		Second Stage House Value
log(FICO)	0.419^{***}		
	(0.020)		
Prev. Rate	0.007		
	(0.005)		
log(Income)	-0.082^{***}		
	(0.004)		
log(Balance)	0.275^{***}		
	(0.005)		
LTV	0.046^{***}		
	(0.000)		
$\log(\Delta HV_t)$	-4.313^{***}		1.108^{***}
	(0.012)		(0.004)
ρ		-0.634	
σ^2		0.189	
Observations	2,146,151		208.075

This table reports the results from a Heckman two-step selection model. The first stage is a probit regression where the the dependent variable is whether a household refinanced under HARP. The second stage estimates the main regression as in Equation (20). The figures in parentheses are standard errors, with 1, 2, and 3 asterisks indicating statistical significance at 10%, 5%, and 1%, respectively.