The Roles of Borrower Private Information and Mortgage Relief Design in Foreclosure Prevention

Lauri Kytömaa*

November 4, 2022

Click here for the most recent version.

Abstract

I study frictions that prevent banks and loan servicers from granting debt relief to struggling borrowers in the U.S. residential mortgage market. I explore how asymmetric information, transaction costs and aid generosity associated with granting debt relief affect mortgage foreclosure outcomes. To disentangle these mechanisms, I introduce a structural model in which banks decide whether to offer debt relief to potentially distressed borrowers when processing relief is costly and borrowers hold private information about their financial well-being. Relative to full information, banks reduce the probability of granting relief to deter financially healthy borrowers from pretending to be distressed, leading to more foreclosures in equilibrium. I use my model to estimate the impact of the Federal Home Affordable Modification Program (HAMP) using the outcomes of mortgages that were originated before the 2008 financial crisis. I find that HAMP incentive payments offset bank costs enough to increase relief disbursement and to decrease realized foreclosures by 3%, or 200,000 properties nationally, over the decade from 2007 to 2016. Despite this, information frictions increased total foreclosures by 14%, or the equivalent of 1.1 million properties and $110 billion of lost value over the same time period. Finally, I find that the level of borrower relief prescribed under HAMP was insufficient for preventing 86% of foreclosures, highlighting the extent of borrower distress arising during 2008.

JEL classification: C57, D82, G21, G51

Keywords: Game theoretic modeling, asymmetric information, foreclosure, mortgage, loan forgiveness

*I thank Dan Ackerberg, Eugenio Miravete, Sam Kruger, Victoria Marone, Jorge Balat, Robert Town and all other participants of the Industrial Organization Writing Seminar for their guidance throughout this project. Additional thanks to Milena Almagro, James Brand, Dita Eckardt, Francisco Garrido, Carlos Garriga, Charles Hodgson, Mark Feldman, Anh Nguyen, Juan Sánchez and Shoshana Vasserman for feedback on my work. I acknowledge current and former UT PhDs Ivan Larsen, Cooper Howes, Justus Timmers, Richard Faltings, Ignacio Núñez and Lee Seltzer for their support during the darkest days of research amidst the COVID-19 pandemic. I need to give recognition and gratitude to Xiaohan Wu, who worked tirelessly as my research assistant. I am grateful to Eric Schulman who helped me learn how to use Python over the course of the PhD program. Finally, I acknowledge the Texas Advanced Computing Center (TACC) at the University of Texas at Austin for providing high performance computing resources that have been essential to the results reported in this paper. All errors are my own. Department of Economics, University of Texas at Austin, Email:lauri.kytomaa@utexas.edu

1 Introduction

Mortgage foreclosure has broad social costs that impact all participants in housing markets. Foreclosure involves significant financial and non-financial costs to a homeowner, and can affect the neighborhood of the property by depressing home values and increasing the probability of future foreclosures (Congressional Oversight Panel [2009], Anenberg and Kung [2014]). Past work also finds that foreclosure causes housing instability, reduced home-ownership, and increased rates of depression and divorce for property owners (Tsai [2015], Diamond et al. [2020]). Debt holders typically also take significant losses when a borrower's home is repossessed and sold through foreclosure; White [2009] estimates that the average foreclosure sale in November 2008 lost a debt holder 57% of the outstanding mortgage value. Despite the direct and indirect costs of foreclosure, United States loan servicers foreclosed 7.8 million homes in the decade between 2007 and 2016 (CoreLogic [2017]).

Due to the broad social costs of foreclosure and the reluctance of mortgage servicers to renegotiate with borrowers, federal and local governments in the U.S. often incentivize servicers to grant debt forgiveness to financially distressed borrowers, instead of foreclosing their homes. Servicers may be different from the original lender, and will be responsible for interacting with borrowers, revising debt terms, and initiating foreclosure proceedings. This paper focuses on debt forgiveness granted through loan modifications that reduce a borrower's monthly payments and increase liquidity, which Ganong and Noel [2020a] find meaningfully reduce borrower default. This forgiveness is predominately achieved through a mixture of maturity term extensions and interest rate reductions. The logic for encouraging loan modifications is straightforward: if a servicer can reduce the cost of a mortgage sufficiently, a borrower should be able to continue making payments. Although the loan can become less valuable to servicer, if reducing the debt produces a smaller loss than the debt holder would experience through foreclosure, relief should be beneficial for all parties.[1]

I study three mechanisms that can inhibit mortgage servicers from granting debt relief to borrowers: 1) asymmetric information, 2) transaction costs, and 3) the level of relief necessary. I look at these three mechanisms jointly because their interaction determines foreclosure outcomes and the effectiveness of public policy. First, the asymmetric information challenge arises from the fact that borrowers know more about their ability to make loan payments than their servicer. For example, an individual's impending divorce might eventually drive him into bankruptcy, whereas support from wealthy relatives could help keep him current on mortgage payments. Since servicers primarily make mortgage decisions based on a limited number of standard indicators, such as a borrower's credit score, monthly income and property value, they will miss important determinants of financial distress. Under this information asymmetry, servicers will also worry that a reputation for generous debt relief will lead individuals to pretend to be financially distressed in order to gain loan forgiveness. To deter this sort of strategic behavior, servicers offer less relief to borrowers relative to full information. Second, even if servicers have enough information to allocate relief correctly, the transaction costs of screening borrowers and modifying contractual terms can dampen incentives to award relief. Granting debt relief is a bespoke process that can incur significant costs for servicers. While mortgage agreements award loan servicers direct reimbursement for foreclosure procedures, servicers are not always similarly compensated for awarding borrowers with debt relief intended to reduce losses (Barclays [2008]). Past research has suggested that many servicing units lacked the resources to process the inflow of borrower defaults during the 2008 financial crisis, making it prohibitively costly to process additional debt relief to

[1] This project is going to assume that the servicer's objectives are aligned with the holder of the outstanding mortgage debt and I use the term "loan servicer" in what follows. If the language is more familiar to the reader, this combined entity could also be relabeled as "creditor", "bank" or "financial institution".

borrowers (Maturana [2017]). Other research has also highlighted the role of restrictive servicing agreements that prevented mortgage servicers from granting debt relief through loan modifications (Piskorski et al. [2010], Agarwal et al. [2012], and Kruger [2018]). Third, the level of relief necessary to prevent a foreclosure may lead to a greater loss for the debt holder than foreclosure itself. The mortgage servicer would then prefer the loss associated with foreclosure to the loss from an overly generous loan reduction.

I create an empirical model to quantify how the interrelated mechanisms affect mortgage debt relief and foreclosure outcomes. My approach formalizes the intuition of Adelino et al. [2013] around how information asymmetry undermines debt relief schemes while capturing the interaction between borrowers and their servicers.[2] I model the strategic interaction between borrowers and loan servicers as a game of debt relief that ends in either repayment or foreclosure. I allow borrowers to hold private information about both their willingness to temporarily miss debt obligations (become *delinquent* on their loan) and willingness to stop paying altogether and enter foreclosure (*default* on their loan). The timing is as follows. A loan servicer first observes borrower characteristics and the distribution of borrower private information, and then commits to a conditional probability of awarding debt relief through a loan modification. A borrower then observes his relief probability and decides whether to go delinquent on his mortgage. Borrowers that go delinquent incur a cost of delinquency, but gain the possibility of earning a loan modification that reduces their debt burden. Borrowers that choose to go delinquent then realize a modification outcome based on the servicer's policy, and decide whether to resume repaying their loan or to default and enter foreclosure. The servicer finds it beneficial to commit to the modification probability ex-ante, because this allows it to deter strategic behavior by borrowers that do not require relief to avoid foreclosure.

I estimate my model using a sample of around 52,000 Fannie Mae-guaranteed loans originated in California between 2004 and 2007, which I track until the end of 2019. By using California, the state with the largest number of home mortgages, I can focus on the mechanisms of interest by avoiding differences in state housing regulations, foreclosure practices, and banking market structure. Because I can observe the final payment or foreclosure outcomes for essentially all loans in my sample, my dataset avoids the significant outcome censoring that exists in earlier papers studying the 2008 financial crisis. Improved visibility on loan outcomes is critical for studying past foreclosure: the long timelines involved with resolving borrower default meant that quarterly foreclosure rates in the U.S. after 2008 remained elevated well into the mid-2010s. Central to the identification of the model, my data allows me to observe servicer identities and the loan modification outcomes of borrowers. My plausibly exogenous variation comes from differences in the rates of loan modification across servicers that are unrelated to borrower characteristics. The mortgage literature clearly documents heterogeneity in servicer modification rates that cannot be explained by borrower characteristics (Korgaonkar [2020], AGA [2011], Agarwal et al. [2017]).

Consistent with recent work by Diamond et al. [2020], my results suggest that historical estimates for the social costs of foreclosure are too low, while also highlighting how information asymmetry can lead to the misallocation of debt relief. I find that there are significant benefits in correctly targeted relief for both borrowers and their servicers: on average an individual borrower places the same value on foreclosure avoidance as a 7.8% reduction in their outstanding principal balance (over $23,000 in value gain). Servicers gain an average of 24.4% of the outstanding principal balance per loan modification that prevents a foreclosure, due to the large expected losses from a foreclosure sale (over $73,000 gain). For context, a comprehensive study by the U.S. Department of Housing and Urban Development [2010] found that the social cost of fore-

[2] Adelino et al. [2013] discuss the role of information asymmetry and propose a theoretical model that suggests that lenders will be less likely to modify distressed loans if the probability that the borrower will start repaying without assistance is high (*self-cure*) or the probability of repayment with assistance is low (*re-default*). This dynamic also exists in my work.

closure was only \$51,061 and that foreclosed homeowners only bear costs of \$10,300 from this process, half the amount estimated in this paper. Unlike these past papers, my model also allows me to consider how information asymmetry forces the servicer to consider the costs of incorrectly awarded debt relief. Servicers lose an average of 15.9% of the outstanding principal balance per borrower who could have avoided foreclosure without assistance (over \$41,000 loss). From a market-level perspective, I find that the social loss from failing to grant a modification that would prevent a foreclosure is 36% larger than the private loss for a loan servicer. Conversely, the social loss from incorrectly granting a modification to a borrower that did not need it is 80% smaller than the private loss faced by the servicer. The relationship between private and social losses of modification highlight how servicer incentives can prevent socially beneficial debt relief to borrowers. My results suggest that servicers only prevent 76% of foreclosures that could have been stopped by granting debt relief because they cannot perfectly screen borrowers, and face losses from making mistakes.

I leverage my results to study the effectiveness of the Federal Government's Home Affordable Modification Program (HAMP), which was launched in 2009 in response to mounting U.S. foreclosures. My structural model allows me to identify the drivers for foreclosure under HAMP. I find that the HAMP incentive payments made to servicers sufficiently offset bank costs to increase relief disbursement and to decrease realized foreclosures by 3%, or 200,000 scaled to the national level, over the decade from 2007 to 2016. Despite this, information asymmetry increased total foreclosures by 14% or the equivalent of 1.1 million properties and \$110 billion of lost value over the same time period. Finally, I find that the level of borrower relief prescribed under HAMP was insufficient for preventing 86% of foreclosures, highlighting the extent of borrower distress coming out of the 2008 financial crisis.

These results augment past work on debt relief policy, and specific studies of the Federal HAMP, by jointly quantifying the mechanisms that affect foreclosure prevention. My headline findings are qualitatively consistent with previous findings. Agarwal et al. [2017] leverage program eligibility cut-offs to study the causal effect of HAMP on foreclosure outcomes and find that the program prevented 600,000 foreclosures nationally, but only reached 1/3 of the intended population of households. Hembre [2018] takes a more structural approach by proposing a single-agent dynamic model to capture borrower defaulting behavior and finds that HAMP prevented slightly over 500,000 foreclosures nationally and that program relief was too large to be socially beneficial. My headline numbers are not directly comparable to these two papers because of my focus on a Californian sample, but I similarly find that only one third of delinquent borrowers were awarded with loan modifications, program subsidies may have exceeded social welfare maximizing levels, and that about 200,000 foreclosures nationally were prevented due to the Federal Government's incentive payments made to loan servicers.[3]

By modeling information asymmetry in mortgage debt relief, I contribute to a rich economic literature on the effect of information on equilibrium outcomes. The topic of my work closely relates to Hendren [2013], Hendren [2017], and Herbst and Hendren [2021], studies that demonstrate how private information can lead to market unravelling in health insurance, private provision of job loss insurance and student loan markets respectively. Information asymmetry in healthcare markets has also received significant academic attention (Cardon and Hendel [2001], Einav et al. [2013], Marone and Sabety [2022]). Within household finance markets, structural approaches have also explored asymmetric information in annuities markets (Einav et al. [2007]), online credit markets (Xin [2020]), and credit card price regulation (Nelson [2022]). I add to the literature by providing a bespoke model for studying the effects of information in the context of mortgage

[3]My estimate of national prevented foreclosures is computed based on my California sample and then scaled to the national level. California experienced a more acute housing crash after 2008 than the national average, a crash that was also accompanied with more foreclosures.

debt relief and foreclosure prevention. My study of how these mechanisms can deter mortgage debt relief has not been explored with past structural methods, even though this approach is fundamental to disentangling their effects on equilibrium outcomes. Analogous to the previous literature, my paper demonstrates the potential for an "unravelling" in mortgage debt relief, in the sense that no relief is granted if households hold a sufficient amount of private information that loan servicers cannot observe.

My findings have clear policy relevance and no previous paper, to my knowledge, formally quantifies how information, transaction costs and relief sufficiency influence loan modification and foreclosure outcomes. I show that debt relief can be a powerful approach for preventing mortgage foreclosure, but that its effectiveness depends on a number of factors. Though I find that information asymmetry between borrowers and their servicers is a primary cause for inefficient allocation of relief, overcoming the information issue can be challenging. Governments can directly reduce the transaction costs faced by mortgage servicers through incentive payments or subsidies, but the the social benefit and the effect on foreclosure outcomes will depend on how well-informed servicers are about borrower default probabilities. Subsidies and financial incentives can certainly help alter servicer decisions, but without careful consideration, policymakers risk unintentionally subsidizing relief that would have occurred in the absence of incentives, or spending to the point that the cost of subsidies outweigh their benefits. Finally, policymakers designing debt relief programs must consider carefully the level of relief required to help borrowers avoid foreclosure. Insufficient relief will have no effect on foreclosure outcomes, while overly generous relief will make servicers reluctant to voluntarily participate in the program because of the losses associated with debt forgiveness.

I proceed as follows. In Section 2 I provide additional background on foreclosures in the United States and the government's response to them. In Section 3 I present my theoretical model of foreclosure prevention policy. In Section 4 I describe my data and present descriptive evidence to support my identifying assumptions. In Section 5 I describe the empirical implementation of my model and the identification strategy. In Section 6 I present model estimates and main results. In Section 7 I leverage the results to evaluate foreclosures outcomes and social welfare under alternative information levels and government subsidies. Finally, in Section 8 I conclude and discuss future research directions.

2 Background

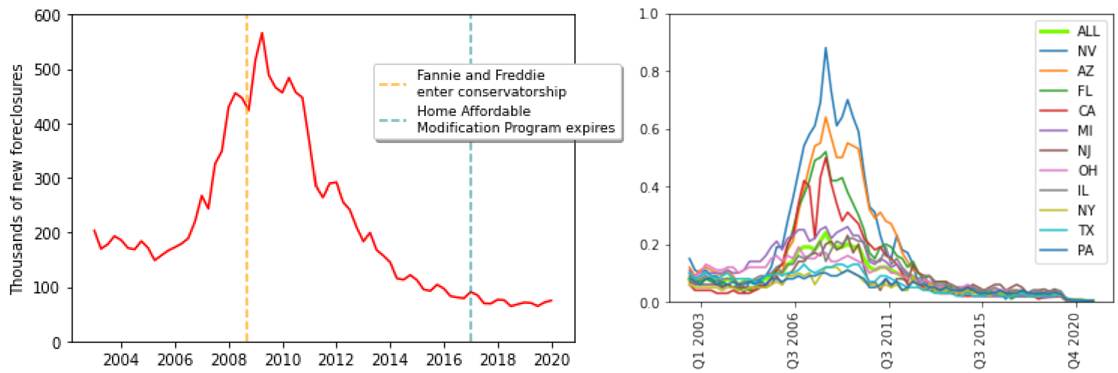

(a) New quarterly foreclosures in the United States

(b) % of consumers with new foreclosures, by state

Figure 1: Trends in United States foreclosures, *Source: New York Fed Consumer Credit Panel/Equifax*

The 2008 financial crisis and associated collapse in the United States residential housing market led to historic levels of residential mortgage foreclosure that persisted half a decade after the start of the downturn. Figure 1 above depicts the trends in national foreclosures (a) and the share of consumers entering foreclosure by state (b) between 2003 and the start of 2020. At its worst, the number of new quarterly foreclosures in the U.S. rose from around 170,000 in the fourth quarter of 2005 to a peak of 570,000 in the third quarter of 2009, a more than three-fold increase. The foreclosure crisis was particularly acute in a handful of states. Arizona, California, Florida and Nevada each experienced rates of quarterly foreclosure more than double the national average. To illustrate the severe nature of foreclosures, nearly 1% of consumers in Nevada from the New York Fed's Credit Panel faced a foreclosure *each quarter* at the peak of the financial crisis.

Economic distress and the dramatic rise in mortgage foreclosures quickly attracted policy interest, due to the risk of further market deterioration and the housing market's important role in financial markets. The U.S. Treasury publicly announced that Fannie Mae and Freddie Mac, which guarantee investors against losses associated with borrower default on mortgage properties, were going to be taken into Federal conservatorship on September 7, 2008 due to mounting financial losses.[4] The following month, Congress authorized $700 billion for the Troubled Asset Relief Program ("TARP") which established several wide-reaching initiatives relating to the U.S. financial system, economic growth and the housing market. As a part of TARP, the Federal Government launched the Making Home Affordable program in February 2009, with the explicit goal of assisting struggling homeowners and preventing foreclosure. Of these homeowner-facing initiatives, the single largest program was the Home Affordable Modification Program (HAMP) which was allocated $75 billion. The objective of HAMP was to encourage financial institutions to grant loan modifications to help reduce the debt burden for struggling borrowers.

Loan modification became the government's primary method of addressing the dramatic rise in residential foreclosures after 2008. Prior to the the economic downturn, private loan modification was uncommon and lacked standardization across financial institutions. Bank modifications would frequently take the form of a "capitalization" where borrowers were brought current on their loans by simply adding missed payments to

[4]For more on the history of Fannie Mae, Freddie Mac and conservatorship see: `https://www.fhfa.gov/Conservatorship/Pages/History-of-Fannie-Mae--Freddie-Conservatorships.aspx`

a borrower's outstanding mortgage balance, increasing monthly payments for borrowers already struggling to repay debt (Agarwal et al. [2017]). The Government introduced loan modification guidelines and provided participating mortgage institutions incentives to grant modifications. Federal guidelines established eligibility criteria and income targeting rules for modification that would reduce an eligible borrower's combined housing expenses to a set share of their income, to improve affordability for borrowers.[5] HAMP loan modifications were only available for mortgages that were originated on or before January 1, 2009.

Incentive compensation under HAMP featured a mix of upfront and on-going payments to servicers and debt holders. Servicers were eligible to receive a one-off payment up to $1,500 for each completed loan modification. After modification, servicers were also eligible for three to five annual "pay for success" payments which were equal to the lesser of $1,000 or one-half of the reduction in the borrower's annualized monthly payment. Servicers became ineligible for the on-going payments if a modified loan ceased to be in good standing.[6]

Table 1 presents an example loan modification taken from my dataset, with specific numbers rounded for exposition. The three primary levers by which a lender can modify a mortgage loan are the 1) annual interest rate, 2) the term of loan, and 3) the outstanding principal balance. Monthly payments can be reduced using any combination of reducing the annual interest rate, extending the term of the loan or forgiving some share of the outstanding principal balance. The most common form of loan modification following the 2008 financial crisis involved the combination of a reduced interest rate and an extended term length, leading to potentially large decreases in the monthly payments of borrowers.[7] In the example, the monthly loan payment is reduced by $858 through a combination of cutting the annual interest rate in half and increasing the term length by 30 years. This decrease in monthly payments translates into a meaningful 33% loss on the mortgage value for the creditor relative to full repayment without a modification.[8] Though this sort of modification presents a meaningful loss for the creditor, the relevant comparison was typically an even greater loss of 40-50% of the mortgage balance through foreclosure.

Though loan modification is mechanically similar to refinancing, the two ways of adjusting a mortgage differ in important ways. Delinquent borrowers are often ineligible for refinancing because being current on one's existing mortgage is a common prerequisite. Modification terms are often also more generous and involve smaller transaction costs than a refinance, mostly because they are used as a last resort to prevent foreclosure. Banks typically won't agree to a loan modification unless a borrower is at serious risk of foreclosure.[9]

To protect debt-holder interests, given the losses associated with modification, government guidelines required mortgage servicers to conduct net present value (NPV) tests to assess whether granting a loan modification would reduce expected payouts.[10] Under these rules, servicers were required to grant loan modification to borrowers whenever it yielded a positive net present value.

In general, borrowers needed to demonstrate financial hardship and a risk of "imminent default" to qualify for a loan modification.[11] Figure 12 in the Appendix displays the primary table from Fannie Mae and Freddie

[5] The precise target share varied by Federal program. Under HAMP, the target income share was 38% whereas the Federal Deposit Insurance Corporation's loan modification formula used a 31% target income share.

[6] Home Affordable Modification Program (2009), Supplemental Directive 09-01

[7] In practice, lenders are hesitant to write-down principal so most modifications involved adjusting only the interest rate and term of the loan.

[8] I provide more detail about the NPV calculation with examples in Appendix A.

[9] Experian (2020), 'How Can I Get a Mortgage Modification?'

[10] Congressional Oversight Panel, *October Oversight Report: An Assessment of Foreclosure Mitigation Efforts After Six Months.* October 9, 2009.

[11] For a detailed description of criteria under HAMP and the FDIC consider Agarwal et al. [2017] or Mulligan [2010].

	Annual interest	Term	Principal*	Monthly payments	NPV for creditor with full repayment**
Pre-modification	6%	10 years	$116k	$1,288	$116k
Post-modification	3%	40 years	$120k	$\underbrace{\$430}_{\text{-\$858 per month}}$	$\underbrace{\$78k}_{\text{33\% loss of value to creditor}}$

Note: * Past missed payments generally added back to principal once a loan has been modified.
** It is standard to calculate lender losses by discounting the original interest rate. For an example, see Maturana [2017].

Table 1: Example of a loan modification

Mac's current Mortgage Assistance Application form. The table demonstrates the wide variety of potentially viable arguments for financial hardship and even leaves room for borrowers to claim an alternative type of hardship not listed.

Uncertainty for borrowers and mortgage servicers arises from vagueness in the definition of financial hardship and imminent default, differences in hardship documentation requirements and the borrower's ability support their claim of hardship with evidence. Borrowers were often uncertain about their eligibility for loan modification and struggled with the processes of earning a loan modification, even when they were eligible. Despite standardized modification criteria and program participation by major mortgage institutions, loan servicers also differed widely in their willingness to grant loan modifications and their propensity to process foreclosures. In my model I assume that borrowers faced randomness in their likelihood of earning a loan modification.

On the servicer side, there was probably uncertainty about the degree and accuracy of borrower financial distress. While the majority of borrowers were severely impacted by the rapid economic downturn, the uncertainty created room for strategic behavior from more financially stable borrowers. In a statement to the U.S. House of Representatives, the Executive Vice President of Wells Fargo's loan servicing division emphasized this issue by stating the importance of: *[...] striking the delicate balance between providing aggressive solutions for those in need and guarding against moral hazard.*[12] Policy documents also explicitly forbid federal program participants from "intentionally" defaulting on their outstanding debt, but it is unclear how feasible it was to enforce such rules against individual borrowers.[13] Past academic literature has also provided evidence for strategic borrower behavior in the residential mortgage market (Guiso et al. [2013], Mayer et al. [2014], Gerardi et al. [2018]).

Despite an acute foreclosure crisis and mixed results from loan modification in the aftermath of the 2008 financial crisis, loan modification continues as a policy tool for foreclosure prevention today. After the Federal HAMP expired in 2016, Fannie Mae and Freddie Mac continued to offer loan modifications through their "Flex Modification" programs. Job losses associated with the COVID-19 pandemic created concern about a new wave of foreclosures and renewed interest in loan modification programs: on July 23rd, 2021 the Biden administration announced a federal initiative that would again help make mortgage debt more affordable by reducing interest rates and extending mortgage terms.[14] Given the ongoing relevance of loan modification policies, it is crucial to understand what can be improved from the past decade of foreclosure prevention efforts.

[12] Hearing before the subcommittee on Housing and Community Opportunity (February 2009), *Loan Modifications: Are Mortgage Servicers Assisting Borrowers with Unaffordable Mortgages?* Serial No. 111-6. p30.

[13] The Housing and Economic Recovery Act of 2008, Title IV-HOPE for Homeowners: *"The mortgagor shall provide certification to the [the Treasury] that the mortgagor has not intentionally defaulted on the mortgage or any other debt [...]"*

[14] See: The Wall Street Journal (2021), *New Aid Coming for Mortgage Borrowers at Risk of Foreclosure.* https://www.wsj.com/articles/new-aid-planned-for-mortgage-borrowers-at-risk-of-foreclosure-11627032601

3 Empirical model

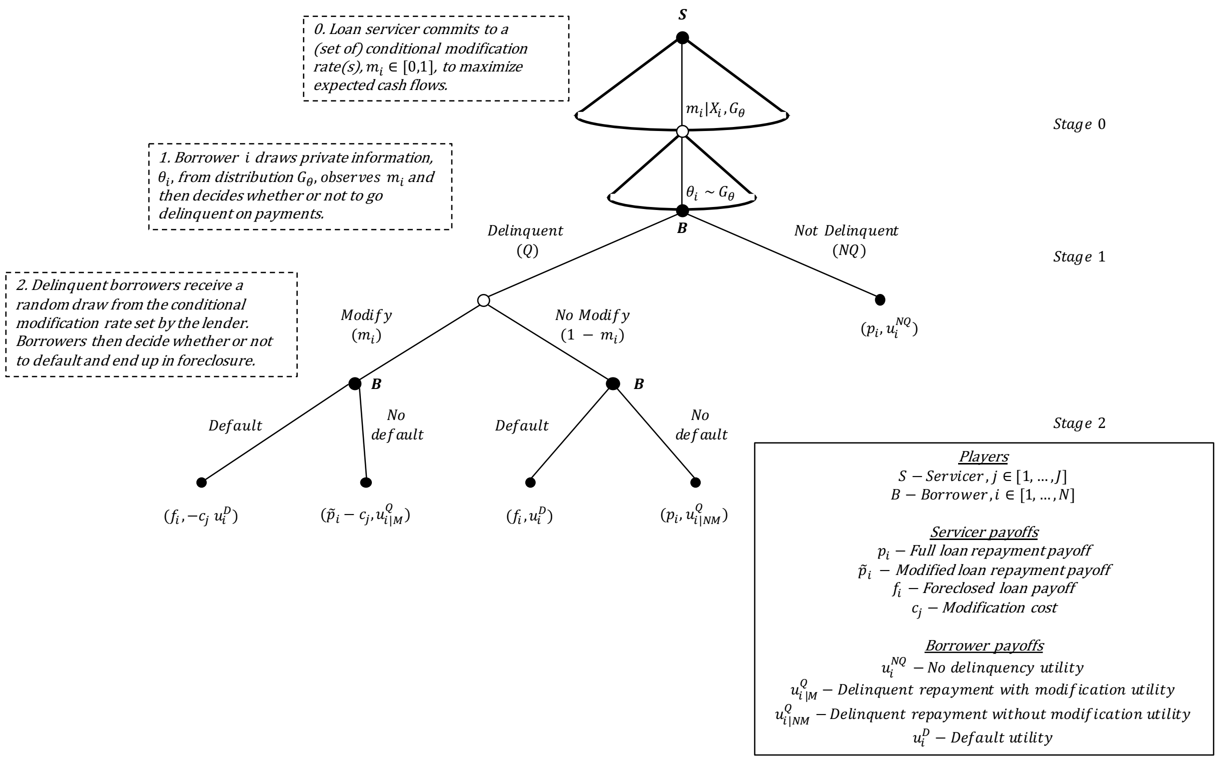

Figure 2: Loan modification game tree

I model the interaction between the servicer and a given borrower as a game of incomplete information with commitment using three decision stages. Figure 2 depicts the overall structure of this game. The timing is as follows:

- **Stage 0: Servicer sets modification policy** – The game begins with the servicer, S, committing to a set of conditional probabilities of modification for each pre-existing loan $i \in [1, \dots, N]$ that it faces. This probability, $m_i \in [0, 1]$, maximizes the expected future cash flows from loan i. These modification probabilities depend on a given borrower's observable characteristics, X_i, the amount of debt owed to the bank, p_i, debt owed after a loan modification, $\tilde{p}_i$, payoffs from foreclosure f_i, a joint distribution of unobservable characteristics, G_θ, and the servicer's cost of processing a modification, c_j. Borrowers are only eligible for modification after entering delinquency, so the servicer internalizes the fact that it will only be facing borrowers willing to miss payments. Since m_i reflects the probability that the borrower will receive a loan modification, if $m_i = 0$ then the borrower will never receive a modification and if $m_i = 1$ then the borrower is guaranteed one. A key underlying assumption is that S has a way to commit to this policy and it is worthwhile to commit to establishing a reputation with borrowers. The design of the optimal mechanism to achieve this is outside of the scope of this paper.

- **Stage 1: Borrower delinquency decision** – After the modification policy has been set, a given borrower i decides whether to go delinquent on loan payments after observing his m_i, observables, X_i, and set of unobservables, θ_i, drawn from the joint distribution G_θ. By going delinquent on payments, the borrower demonstrates financial distress and may be awarded a loan modification. By choosing not to be delinquent on payments, the borrower avoids any costs associated with late payment but also forgoes the possibility of earning a loan modification. If a borrower has chosen to go delinquent, he receives a Bernoulli draw with probability m_i and either realizes a loan modification or does not. If a borrower does not go delinquent on payments then the game ends with full repayment of the outstanding debt.

- **Stage 2: Borrower default decision** – If a borrower has been delinquent on his payments and has received a modification realization, he must then make a final decision about whether or not to default. Servicer and borrower payoffs will vary based on this decision. In my model the final default decision is synonymous with foreclosure and the borrower will lose his collateral/home.[15] Associating default with foreclosure in this way is relatively uncommon in the past mortgage literature, though it is the definition used in Ghent and Kudlyak [2011].

The model set-up simplifies an inherently dynamic process of mortgage debt relief in order to gain tractability in the interaction between a mortgage servicer and his borrowers. Bajari et al. [2013] and Fang et al. [2016] propose dynamic, single-agent models to study borrower default decisions and to learn about borrower discounting. By compressing mortgage resolution after financial distress into my three stage model of commitment, I trade-off model richness for improved tractability for estimating the role of information asymmetry in the loan servicer's policy setting decision. Avoiding a fully dynamic model comes with certain costs, for example, my model set-up cannot provide deep analysis of borrower expectations and learning, or servicer dynamic reputational considerations.

My set-up is also conditional on a set of loans that have already been made out to borrowers. This means that I focus on the effect of information on mortgage debt relief given a set of loans. Considering how information asymmetry and the loan renegotiation process can influence the set of loans that are originated to borrowers in equilibrium is outside the scope of this paper.[16]

3.1 Borrower behavior

Expected borrower behavior drives the loan servicer's cash flow maximization decision, so it is helpful to formally define the borrower's utility function and work backwards through the model stages before providing further detail on the servicer's problem.

3.1.1 Borrower utilities

An individual borrower i's payoffs will depend on player actions and the modification realization. I define borrower i's utility to depend on three key objects:

1. **Value of keeping home relative to defaulting and entering foreclosure (H_i):** H_i captures the value derived from a property, the benefit of avoiding foreclosure and the opportunity cost associated with paying off a mortgage. The term captures a variety of effects, including the value of the mortgaged property, and the transaction costs of moving to a different home. One would expect the benefit

[15] In practice, the servicer needs to initiate foreclosure proceedings after a borrower defaults but this generally occurs after a sufficient number of missed payments.

[16] This rules out a moral hazard effect where the occurrence of loan modification changes the overall distribution of borrower types in the population. I limit the potential effect of this moral hazard by focusing on loans made before 2008 in estimation given the unprecedented decline in the housing market and Federal Government aid for loan modification that followed. It is unlikely that borrowers would have been anticipating such outcomes when taking out mortgages at the start of the decade.

of avoiding foreclosure to play an important role here, because foreclosure significantly damages a borrower's credit score and ability to engage with financial markets in the future. Opportunity costs of remaining in the property will also be picked up by this term: for a borrower facing a decision between affording food and making mortgage payments, H_i will be low, all else equal. Alternatively, limited outside options for accommodations will drive H_i up for a given borrower.

2. **Delinquency cost of temporarily missing a mortgage payment (Q_i):** Q_i represents the collective costs of temporary payment delinquency, which include a decline in a borrower's credit score, late interest fees and the potential stress involved with interacting with a debt collector. This term will also capture the opportunity cost of on-time payment: a borrower facing a liquidity constraint will find it incredibly costly to remain current on payments, leading to a lower Q_i. A low Q_i can also capture a borrower's potentially naive perception of low delinquency costs.

3. **Payment disutility on the mortgage debt (unmodified p_i or modified $\tilde{p}_i$):** The payment disutility term measures the discounted stream of payments that need to be made on a mortgage to pay off outstanding debt. I denote unmodified payment disutility as p_i and modified payment disutility as $\tilde{p}_i$. It will always be the case that $p_i > \tilde{p}_i$. The payment disutility decreases the net utility value of remaining in a home since a loan balance must be paid off to retain the mortgaged property.

I allow both the H_i and Q_i terms to have heterogeneity across borrowers and to contain both observable and unobservable components, reflecting the fact that the borrower may hold some private information about these two objects. In addition to this, I allow the loan servicer to hold information about a borrower's probability of defaulting that is unavailable to the researcher/econometrician. Allowing the servicer to have additional information about the borrower reflects the fact that financial institutions request additional information from borrowers as a part of the loan modification process that is not observable in my dataset. The assumption of servicers with unobservable information about the borrower introduces an endogeneity problem that I need to address as a part of my identification argument. With a better-informed servicer, loans that receive modification will be unobservably different from loans that do not. Ignoring this aspect risks overstating the effectiveness of modifications in foreclosure prevention.

I specify the form of the borrower utility function for all potential outcomes of the game below.

Repay without delinquency

$$u_i^{NQ} = \underbrace{x_i'\beta + \xi_i + \varepsilon_i}_{\textit{Home utility relative to default } (H_i)} - \underbrace{p_i}_{\textit{Payment disutility}}$$

Repay with <u>delinquency</u> but without modification

$$u_i^{Q}|NM = x_i'\beta + \xi_i + \varepsilon_i - p_i - \underbrace{Q_i(w_i'\lambda, \eta_i)}_{\textit{Delinquency cost}}$$

Repay with <u>delinquency</u> and <u>modification</u>

$$u_i^{Q}|M = x_i'\beta + \xi_i + \varepsilon_i - \underbrace{\tilde{p}_i}_{\textit{Modification disutility}} - Q_i(w_i'\lambda, \eta_i)$$

Default

$$u_i^{D} = 0 - 0 - Q_i(w_i'\lambda, \eta_i)$$

■ Unknown to econometrician ■ Unknown to econometrician *and* loan servicer

- x_i is the set of observable characteristics that explains the mean value of keeping a home relative to defaulting. It contains characteristics of the borrower and the property, such as the value of the home, its geographic location and the income of the borrower. β is a vector of parameters for these observables.

- w_i is the set of observable characteristics that explains the mean delinquency cost. It is allowed to contain a different set of observables than x_i, including observables such as a borrower's credit score. λ is a vector of parameters for these observables.

- ξ_i explains the component of the variance of H_i that *is observed by the servicer* but is not known to the econometrician. This term is incorporated into the utility specification to reflect the fact that the loan servicer possesses additional information that a researcher doesn't have access to in conventional loan servicing data. ξ_i includes any information that a borrower would report to his servicer in order to be eligible for a loan modification. Examples of this sort of information include job losses, recent divorces and occurrences of major injury or death.

- ε_i is the private information of the borrower about H_i that further explains the variance of value of keeping the home. ε_i captures information that servicers cannot learn through typical means, including a borrower's emotional ties to a particular neighborhood or community, ability to access credit through informal markets and outside options for housing if they cannot remain at their property. The distribution of ε_i is known to the loan servicer.

- η_i is the private information of the borrower about their delinquency cost and it explains the variance in Q_i. Here I assume that the servicer does not observe individual heterogeneity beyond differences in w_i values. The η_i term captures factors such as an individual's financial sophistication, ability to negotiate with their bank and perception of being behind on debt.

- ξ_i, ε_i and η_i are drawn from a joint distribution G_θ, with known parameters. These drive heterogeneity between borrowers in their values of H_i and Q_i.

- p_i and $\tilde{p}_i$ are the net present values of future mortgage payments for an unmodified and modified loan respectively. These are observable for the econometrician and market participants. The parameter on payment disutility has been normalized to 1 so that all other model parameters can be interpreted in the units of mortgage payments.

The full set of model parameters Θ includes the parameters of G_θ, as well as the vectors β and λ. I discuss specific parametric assumptions in the estimation section.

3.1.2 Default decision

At the default stage, all borrower uncertainty is resolved and borrowers make a decision about whether to repay their mortgage balances or to default, given their modification outcomes. The borrower has already entered delinquency if he has reached this decision, so delinquency costs play no role in the choice; these costs must be incurred either way. A delinquent borrower chooses to default whenever the disutility of payment exceeds the value of keeping the home:

- **Default if modified when:**

$$\tilde{p}_i > H_i = x_i'\beta + \xi_i + \varepsilon_i \tag{1}$$

- **Default if not modified when:**

$$p_i > H_i = x_i'\beta + \xi_i + \varepsilon_i \tag{2}$$

Since $p_i > \tilde{p}_i$ the borrower is always (weakly) more likely to default without a modification.

Whenever a borrower defaults, I assume that the modification realization has no impact on payoffs: the servicer receives a foreclosure payout of f_i and the borrower receives a payoff of u_i^D. Modification outcomes impact all "no default" payoffs. If a modification was realized and the borrower chose not to default then the servicer receives a payoff $\tilde{p}_i < p_i$ and the borrower receives a payoff of $u_i^L|M$, associated with a late payment, modification and no default. Finally if no modification was realized and the borrower chose not to default then the servicer receives a payout of p_i and the borrower receives a payoff of $u_i^Q|NM$. It is the case that $u_i^Q|NM < u_i^Q|M$, meaning that a delinquent borrower at least weakly prefers to receive a modification.

3.1.3 Delinquency decision

At the delinquency stage, borrowers face uncertainty about whether or not they will realize a loan modification. A borrower i chooses to go delinquent if the expected gain from delinquency exceeds the utility from repayment:

$$\begin{aligned} & m_i \cdot Max\{u_i^D, {u_{i|M}}^Q\} + (1-m_i) \cdot Max\{u_i^D, {u_{i|NM}}^Q\} \geq u_i^{NQ} \\ & = \underbrace{m_i \cdot Max\{p_i - H_i, p_i - \tilde{p}_i\} + (1-m_i) \cdot Max\{p_i - H_i, 0\}}_{Delinquency\ expected\ utility\ gain} \geq \underbrace{Q_i}_{Delinquency\ cost} \end{aligned} \tag{3}$$

The second line in equation 3 eliminates common terms and provides intuition about the borrower's delinquency decision. The left-hand side of the inequality is the expected utility gain from entering delinquency and the right-hand side is the borrower-specific delinquency cost. A borrower will go delinquent whenever the expected gain from delinquency exceeds the cost of delinquency. If a borrower chooses to default, he gets out of mortgage payment p_i but suffers the loss of H_i yielding the term $p_i - H_i$. Meanwhile if the borrower chooses to repay he either gets a utility gain of $p_i - \tilde{p}_i$ with a modification, or a utility gain of 0 when modification does not occur. The max operator appears in this equation because the borrower observes the modification outcome before making the final default decision.

3.1.4 Borrower types

I define borrower types based on their choice if they were hypothetically at the final default decision. Borrower types are continuous, but a borrower will fall into one of three distinct groups based on how modification impacts their default outcome. Neither the loan servicer nor the researcher can directly observe these types due to the unobservable components of H_i and Q_i, but they can form probabilities for a given type of borrower given observable characteristics.

- **Low type - Always defaults**: For this type, $p_i - H_i > p_i - \tilde{p}_i$, so his delinquency decision simplifies to:

$$p_i - H_i \geq Q_i$$

The low-type defaults whether or not he receives a loan modification since he prefers to get out of payment disutility than to repay in all circumstances. This type will avoid delinquency (and thus default) only if the delinquency cost outweighs the gain of entering foreclosure. In practice, one would expect the magnitude of $p_i - H_i$ to greatly exceed Q_i. Offering modification to a low-type will not prevent a default, so a servicer would prefer not to grant a modification for this type of borrower. The always defaulting type can be interpreted as a borrower in significant financial distress or as a borrower willingly walking away from a home mortgage.

- **Medium type - Requires modification to avoid default**: For this type, $p_i - \tilde{p} > p_i - H_i > 0$, so his delinquency decision is given by:

$$m_i \cdot (p_i - \tilde{p}_i) + (1 - m_i) \cdot (p_i - H_i) \geq Q_i$$

The medium-type defaults only if he does not receive a loan modification. This is the type of borrower that the loan servicer finds optimal to grant loan modifications to in order to avoid foreclosure. This is the only borrower type that is "efficient" to modify because a modification helps prevent foreclosure.

- **High type - Never defaults**: For this type, $0 > p_i - H_i$ so his delinquency decision simplifies to:

$$m_i \cdot (p_i - \tilde{p}_i) \geq Q_i$$

The high-type's home value never enters his delinquency decision because he will keep his home in all circumstances. This type's delinquency decision depends on whether or not the expected decrease in mortgage payments from modification outweighs the cost of delinquency. Though the high-type will not enter foreclosure, not all borrowers of this type can be interpreted as "strategic". As an example, individuals that face short-term liquidity constraints that force them into delinquency fall in this category even if they are not attempting to fake distress in order to obtain a modification.

3.2 The servicer's problem

A loan servicer sets a modification policy, m_i, for all possible combinations of borrower observables, given a known distribution of unobservables. The servicer first forms the probabilities for borrower late payment and default for any given probability of modification. Using these probabilities, the servicer can calculate the expected cash flows from a given borrower i, conditional on m_i, borrower observable characteristics and model parameters.

I allow for multiple loan servicers indexed by $j \in [1, \ldots, J]$. To introduce heterogeneity across different loan servicers in their willingness to grant loan modifications, consistent with the findings in AGA [2011], Agarwal et al. [2017] and Korgaonkar [2020], I incorporate a servicer-specific transaction cost of processing a modification, c_j. This term reflects differences in servicing companies' ability to process modifications and screen borrowers, which can be driven by staff availability and training. Servicers have no interaction between one another and a loan i is always affiliated with a specific servicer j.

[16] AGA [2011] finds that adding servicer fixed-effects to their baseline regressions for loan modification increases explanatory power by around 40%.

Formally, the expected cash flow maximization problem for a servicer j facing loan i is given by:

$$\begin{aligned}
&\underset{m_i}{Max} \quad E[\pi_{ij}|X_i, \xi_i, p_i, m_i; \Theta] \\
&= \underset{m_i}{Max} \quad Pr(DLQ_i|X_i, \xi_i, p_i, m_i; \Theta) \times \Big[m_i(E[\pi_{ij}|X_i, \xi_i, p_i, DLQ_i, MOD_i; \Theta] - c_j) \\
&\qquad\qquad\qquad + (1 - m_i)E[\pi_{ij}|X_i, \xi_i, p_i, DLQ_i, NO\ MOD_i; \Theta]\Big] \\
&\qquad\qquad\qquad + (1 - Pr(DLQ_i|X_i, \xi_i, p_i, m_i; \Theta)) \quad \times \quad \underbrace{p_i}_{E[\pi_i|NO\ DLQ_i]}
\end{aligned} \tag{4}$$

Where:

- $X_i = \{x_i, w_i\}$: The set of borrower observable characteristics that affect the value of keeping the home and the delinquency cost, respectively.
- ξ_i : Borrower heterogeneity in home utility relative to default that is observable to the servicer but not the econometrican.
- p_i: Discounted stream of outstanding mortgage payments. $\tilde{p}_i$ is a function of p_i so does not appear in the conditioning of Equation 4. I also assume that the foreclosure payoff, f_i, is a function of p_i.
- c_j : Servicer specific transaction cost of processing a modification.
- DLQ_i, MOD_i: Realizations of delinquency and a modification, respectively.
- Θ : Set of model parameters, which includes β, λ and the parameters of the joint unobservable variable distribution, G_θ.

Under this set-up I assume that the servicer knows the outcomes $\tilde{p}_i$ and f_i when he is setting his modification probability m_i. I justify this assumption by having the loan servicer set the m_i modification probability based on the mean expected losses from $\tilde{p}_i$ and f_i rather than the exact value of these future outcomes for a given borrower. The interpretation is that the servicer commits to a policy based on expected losses, even if he doesn't know the exact losses for a given borrower when committing to a policy.

3.3 Equilibrium of the model

In equilibrium, each conditional modification policy, m_i maximizes the associated servicer j's expected cash flow $\forall i, j$ and each borrower, at each stage of the game, is maximizing his expected utility. Since borrower types are continuous and servicer policy is set ex-ante, a given borrower will never be indifferent in his delinquency and default decisions. The cost of modification term, c_j, also helps break indifference for the lender in cases where modification has no effect on a borrower's outcome: the servicer will always prefer not to modify if modification has no impact on borrower outcome probabilities.

Existence of equilibrium in the game of commitment is relatively straight forward. The servicer conditional modification probability m_i lies on the set $[0, 1]$ by definition of being a probability, so there must exist $m_i^* \in [0, 1]$ such that π_i is maximized (though it may not be unique). Once a modification policy has been set by the servicer, then borrowers make delinquency and potential default decisions conditional on this probability. The final default decision is essentially a static, single-agent problem with full information. Borrowers know the payoffs associated with defaulting or not defaulting and servicer behavior no longer has an effect on borrower payoffs. Given a delinquency decision and a realized modification outcome, borrowers will never be indifferent between defaulting or not defaulting by the continuity of H_i, Q_i and p_i. At the delinquency stage, a given borrower faces uncertainty over the modification outcome but will always have a unique utility maximizing option. Again, given the continuity of H_i, Q_i and p_i and the fact that $m_i \in [0, 1]$, a given borrower will never be indifferent between entering delinquency and not entering delinquency.

Uniqueness of equilibrium depends entirely on uniqueness in the servicer's conditional expected cash flow maximization problem. Unlike the borrower decisions, uniqueness of the servicer's optimal m_i does not immediately follow from continuity of borrower characteristics. For example, a servicer will be indifferent between different modification probabilities whenever the borrower's probability of entering delinquency remains near zero for all $m_i \in [0, 1]$. This indifference arises from the fact that if a borrower never becomes delinquent, then the servicer never risks incurring the cost c_j, associated with a modification. When delinquency does occur, the modification cost ensures that the servicer is never indifferent between modification and no modification if financial relief fails to shift default probabilities.

To ensure uniqueness in the servicer's problem, and thus model equilibrium, I make the following assumption about unobservable servicer behavior:

- **Assumption E1 (Servicer tie-break rule):** Define the set of conditional modification rates M_i as all $m_i \in [0, 1]$ that maximize the servicer's expected cash flow problem for borrower i. Whenever the cardinality of the set exceeds 1, $|M_i| > 1$, then assume $m_i^* = Max(M_i)$.

Assumption $E1$ states that whenever multiple conditional modification probabilities yield identical expected cash flows for the servicer, it will set the highest rate of this set as the equilibrium policy. I make this assumption based the requirement that Fannie Mae-affiliated servicers were required to grant loan modifications whenever a modification had a positive expected net present value. This tie-break rule simply implies that the indifferent servicer sets a policy that benefits borrowers.[17]

4 Data and Descriptive Evidence

The primary data for my project comes from Fannie Mae's *Single-Family Loan Performance* dataset.[18] This resource provides access to monthly loan performance information for around 45 million U.S. mortgages originated from the year 2000 up until present day.[19] All mortgages in the sample are fully amortizing, full documentation, conventional fixed-rate loans.[20] Fannie Mae is the single largest government-sponsored enterprise that purchases home mortgages and then securitizes them into residential mortgage-backed securities. The total value of Agency Mortgage-Backed Securities was around $4 trillion at the start of 2007; Fannie Mae loans accounted for half of this. [21]

Fannie Mae servicing data is an excellent sample for estimating my empirical model because it allows me to track individual mortgage performance from origination of a loan to its termination. For each month of a loan's existence, I can observe financial characteristics such as principal outstanding, monthly interest rate, term remaining and a borrower's payment delinquency status. When a loan receives a modification, I can

[17]Results aren't sensitive to this assumption. Alternative tie-break rules yield similar parameter estimates.

[18]See: `https://capitalmarkets.fanniemae.com/credit-risk-transfer/single-family-credit-risk-transfer/fannie-mae-single-family-loan-performance-data`

[19]Quarterly datasets are continuously added to the dataset by Fannie Mae.

[20]Focusing on conventional, full documentation loans helps avoid issues involved with misreporting and fraud that took place prior to the 2008 financial crisis (see for example: Kruger and Maturana [2019]).

[21]Urban Institute (2021), Housing Finance at a Glance, July 2021.

observe any adjustments made to its characteristics. At the termination of the loan, I can observe whether it was paid off or foreclosed upon and the net proceeds from a foreclosure sale.[22] The dataset also contains a rich set of borrower characteristics at loan origination, including borrower credit score, debt-to-income ratio, loan-to-value ratio, geographic location and intended property usage.

I merge regional home price indices and unemployment rates with my dataset to expand the set of observable characteristics. The house price index comes from the Federal Housing Finance Agency and the unemployment data comes from the Bureau of Labor Statistics. These two measures track changes in home values and regional unemployment over the life of a mortgage. The measure for home prices is particularly critical because it enables me to compute updated home values on a monthly basis. A limitation with the Fannie Mae data is that a property's geographic region is restricted to a 3-digit ZIP code and an MSA code to protect a borrower's identity. In practice, this is similar to following county-level trends.[23]

I make four primary restrictions when creating my estimation sample from the raw data. Firstly, I focus on loans originated in California to avoid challenges associated with cross-state differences in mortgage laws, regulations and borrower composition. Ghent and Kudlyak [2011] document important differences in state foreclosure laws and recourse policy that complicate a cross-state study of bank modification policy. National servicing organizations also behave differently given different rules across states. I chose California because it is the single largest state for Fannie Mae mortgages in my sample, and the most populous state in the United States. I address cross-state differences in Kytömaa [2022].

Secondly, I divide the data based on the origination quarter of the mortgage loans to allow model parameters to differ based on when a loan was made. Numerous papers have documented the decline in lending standards leading into the 2008 financial crisis (for example: Bajari et al. [2008] or Demyanyk and Van Hemert [2009]). In my set-up, changes in the composition of borrowers cannot simply be accounted for by controlling for origination timing in the observables: different distributions of borrowers respond differently to loan characteristics and draw their unobservables from entirely different distributions. If I were to pool all loans into a single sample, I would need to account for origination timing not only in observable borrower characteristics, but also in the unobservable distributions governing delinquency and default decisions. Separating the data by origination-quarter enables more flexible and straightforward estimation, while allowing me to quantify differences in the pool of borrowers based on when they received their loans.

Thirdly, I restrict my sample to loans that were originated between the beginning of 2004 and the end of 2007, to avoid the period of declining mortgage interest rates between 2000 and 2004 and the effects of the 2008 financial crisis on loan payback and origination, and to focus on loans that were eligible for HAMP. Figure 3 depicts the 30-year fixed rate mortgage average in the United States from 2000 to the end of 2021. The sharp decline in interest rates between 2000 and 2004 means that the majority of borrowers who bought a home in this period refinanced their fixed rate mortgages before the start of the 2008 financial crisis. Borrowers who did not refinance after such a substantial decline in interest rates are likely to be unobservably different from borrowers who did, possibly because they are financially worse off or otherwise ineligible for refinancing. I do not consider loans originated after 2007, because the ensuing economic downturn affected lending practices and borrower decisions to purchase homes. Furthermore, under HAMP, only loans that had been originated on or before January 1, 2009 were eligible. Finally, prior to 2008 there was little precedent for the net-present-value reducing loan modifications that were introduced to help prevent foreclosure.

[22] The vast majority of non-foreclosed mortgages are paid off through refinancing.

[23] As an example, there are 59 unique 3-digit ZIP codes in California and 58 unique counties. Figure 13 in the Appendix compares the maps for California's 3-digit ZIP codes and its counties.

Table 2: Summary statistics for pooled estimation sample

	Full sample		Delinquent loans		Modified loans		Foreclosed loans	
Observations	51,884		7,843		2,203		5,439	
Share of loans entering delinquency	15.1%		-		-		-	
Share of loans receiving modification	4.3%		28.1%		-		13.0%	
Share of loans entering foreclosure	10.5%		69.4%		32.0%		-	
At origination								
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Principal balance ($000s)	244.82	87.48	264.37	83.34	280.51	79.2	263.53	83.05
Home value, ($000s)	430.88	156.35	384.75	126.04	416.22	125.27	365.12	115.54
Loan-to-value	0.59	0.17	0.7	0.13	0.69	0.12	0.73	0.11
Debt-to-income ratio	0.37	0.13	0.42	0.11	0.43	0.1	0.42	0.11
Credit Score	730.91	53.79	699.27	53.59	690.58	53.36	701.56	53.23
In October 2009								
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Principal balance ($000s)	227.50	84.46	251.82	81.11	268.49	77.08	251.68	81.00
Home value ($000s)	338.91	149.31	259.13	105.40	278.86	105.60	236.06	87.47
Loan-to-value mean (std.)	0.75	0.31	1.04	0.30	1.03	0.29	1.11	0.28

Figure 3: Mortgage interest rate trends

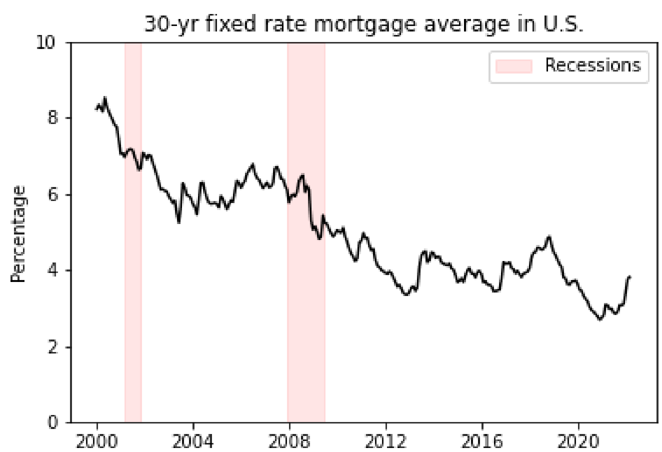

Figure 4: Delinquency timing

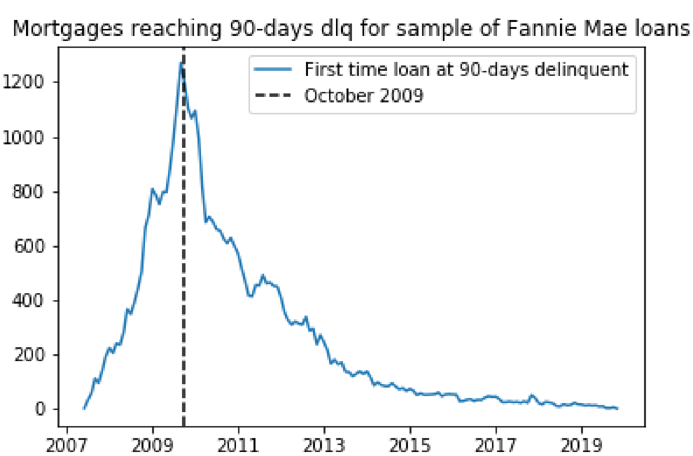

Figure 5: Realized loan proceeds as a share of outstanding principal

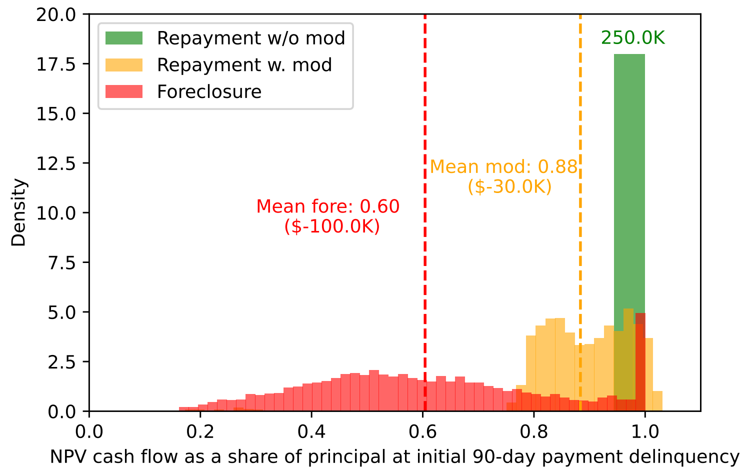

Finally, I remove loans that were paid off prior to October 2009, which is the peak month for mortgage delinquencies in my sample. My three-stage model requires an assumption about which month's data should be chosen to inform the borrower's decisions. Though my sample loans are *originated* between 2004 and 2007, loan information used before the onset of the 2008 financial crisis misses critical variation in home values, relative loan-to-value ratios and changes in national unemployment that help explain borrower decisions. I chose October 2009 to focus on loans that were exposed to the financial hardship of the financial crisis and about which decisions had to be made regarding mortgage delinquency and default, as home values declined.

My final sample looks at 51,884 loans originated in the second quarters of 2004, 2005, 2006 and 2007, tracking their monthly performance until the end of 2019.[24] Table 2 provides summary statistics for the four quarters of data pooled together. In the pooled sample, borrowers end up in all possible outcomes outlined in the model, but the majority of them avoid delinquency altogether, where I define delinquency as 90 days behind on payments. Of the 15.1% of borrowers who become three-months late on their mortgage debt, the majority lost their home (69.4% of delinquent loans).[25] Loans that entered delinquency and were awarded a loan modification tended to perform better than non-modified loans: only 32.0% of modified loans ended in foreclosure while about 84% of delinquent, non-modified loans ended in foreclosure. Despite these trends, no single condition appears deterministic: loans that received modifications sometimes still ended in foreclosure while un-modified loans sometimes "self-cured" and started repaying without any debt forgiveness.

These statistics suggest some role for modification in foreclosure prevention, but causality is not immediately obvious. If observably (or unobservably) financially healthier borrowers receive modifications then the fact that modified loans tend to avoid foreclosure can be explained by borrower characteristics rather than the effectiveness of modification on its own. The summary statistics show that modified loans tended to have higher home values and lower loan-to-value ratios at origination and in October 2009 than loans that ultimately ended in foreclosure. Conversely, modification recipients tended to have lower credit scores and higher debt-to-income ratios at origination than the average foreclosed borrower, showing that they were not observably less risky by all measures.

[24] Estimation uses the second quarters from these years because April, May and June tend to be busy months for U.S. home sales. Model estimation is also computationally costly so it was necessary to select a sub-set of all potential Fannie Mae loans.

[25] I group foreclosure and foreclosure alternatives together. Alternatives include short sales, third-party sales or deeds in-lieu-of foreclosure.

There are large differences in loan outcomes across the four origination groups and pooled summary statistics conceal important trends that influence my results. Mortgages originated closer to 2008 were more likely to enter delinquency and foreclosure than mortgages made earlier in the 2000s. For the 2004 cohort, only 7.76% of mortgages ever reach 90-days delinquent on debt payments and 44.41% of these delinquent loans ultimately enter foreclosure. Delinquency rates climb to 17.57% for the 2005 cohort and then over 27% for both 2006 and 2007. Foreclosure rates climb in a similar way following the 2004 cohort: for delinquent loans these were 55.62%, 58% and 60.00% for 2005, 2006 and 2007 samples respectively. Difficulties with later repayment is partly explained by the fact that origination cohorts take on progressively more debt over time. Delinquent loans in the 2004 cohort hold a mean principal balance of $203k in October 2009, while the other three sample cohorts hold mean balances of $234k, $265k, and $273k in the same month respectively. Earlier cohorts also tended to have lower debt-to-income-ratios, lower loan-to-value ratios, and higher credit scores at the time of origination.

4.1 Measures of p_i, $\tilde{p}_i$ and f_i

The choice of borrower-specific payment measures for unmodified and modified loans plays a central role in estimating my empirical model. The two values of p_i and $\tilde{p}_i$ reflect the payment disutility in the borrowers' reduced-form utility functions. Due to a necessary scale-normalization in my discrete choice framework, all other parameters of the model can be interpreted relative to a change in this measure.

Calculating the net present value of future payments for a fully amortizing, fixed rate mortgage is straightforward when information on a loan's interest rate, outstanding balance, and term remaining is available. Equation 4.1 below depicts the form of the calculation:

$$p_i = \Sigma_{t=0}^{T} \delta^t \textit{Monthly payment}_i = \frac{1-\delta^{T+1}}{1-\delta}(\textit{Monthly payment}_i) \quad \delta : \textit{Discount factor}, T : \textit{Term remaining}$$

I discount using an individual mortgage's own interest rate, consistent with approaches for mortgage loss calculations (see Hu and Cantor [2004] or Maturana [2017]). Using a mortgage's own interest rate means that the outstanding principal balance coincides with the net present value of future payments. Given this, I directly use a mortgage's outstanding principal balance on October 2009 for p_i.[26] I also assume that borrowers and the mortgage servicers have identical discount rates.

I use the average realized losses from modification and foreclosure to set my $\tilde{p}_i$ and f_i measures. Figure 5 provides information about the ex-post realized losses that the servicer faces from modification and foreclosure compared to full repayment. If full repayment is associated with a payback of 100% of expected future payments, then modification is on average associated with an 87% payback. The average realized payback from a foreclosure sale was only 61% of the original balance due, but the data suggests that there is a wide spread in how much of the outstanding debt the servicer can recoup from a foreclosed loan. My approach assumes that $\tilde{p}_i$ and f_i are deterministic based on p_i and the realized means of the distributions of the two outcomes. For example, a loan with a $100k outstanding balance has a p_i = $100k, $\tilde{p}_i$ = $87k and f_i = $61k.

The deterministic rules for $\tilde{p}_i$ and f_i are strong simplifications, but there are a few reasons why they offer a reasonable approximation to reality. With regards to loan modification losses, Fannie Mae affiliated loan servicers followed prescribed rules for the size of a modification and were required to award one whenever borrowers satisfied a net-present-value criteria. These guidelines focused on specific debt-to-income shares for borrowers, in essence limiting servicers' ability to adjust the size of modification. Unfortunately my data

[26] October 2009 is chosen because it represents the period of peak mortgage delinquencies in my sample (Figure 4).

Table 3: Share of delinquent loans receiving a modification, by servicer in pooled estimation sample

Servicer	Number of delinquent loans	Share receiving modification
Bank of America (BoA)	2,755	13.4%
Wells Fargo	1,297	48.0%
Fannie Mae/Seterus	1,164	47.9%
Citimortgage	831	19.6%
JP Morgan	658	35.6%
GMAC	625	16.6%
Green Tree	513	51.1%

does not contain updated information about borrower income around the time of mortgage delinquency, so using a fixed rule for modification size is an attempt to approximate the formal discount rule that servicers applied to awarded modifications.

On losses from foreclosure, one may worry about using the distribution of realized losses as a measure for loans that do not actually end up in foreclosure. As pointed out by Korgaonkar [2020], realized losses may not reflect the losses that would be faced by non-foreclosed homes were they to default counterfactually. Nevertheless, it is not obvious whether non-foreclosed homes would face a larger or smaller loss through a foreclosure sale relative to market value.

4.2 Servicer heterogeneity in loan modification and other outcomes

My sample focuses on the seven largest loan servicers, as measured by the number of loans that they manage. In order of size, the top seven Fannie Mae servicers in California from 2004 to 2007 were: Bank of America, Wells Fargo, Citimortgage, GMAC, JP Morgan, Fannie Mae/Seterus and Green Tree Servicing. The focus on this group is primarily motivated by practical reasons: the number of observations per servicer falls dramatically below the top seven, making it hard to precisely estimate cost parameters for smaller servicers. Across the origination cohorts used for estimation, the top seven servicers account for 76% of all loans serviced.

I report the share of delinquent loans that received a modification by servicer in Table 3. The raw data shows clear differences in the likelihood of receiving a modification between servicers. Bank of America and Well Fargo, the two organizations with the largest number of delinquent loans, had a nearly 35 percentage point difference in their shares of delinquent loans receiving modifications. Similarly, while GMAC, JP Morgan and Green Tree all faced a comparable number of delinquent loans in my sample, the servicers' shares of modified loans were dramatically different. GMAC only offered modifications to 16.6% of its delinquent borrowers, when JP Morgan and Green Tree modified 35.6% and 51.1% of their delinquent borrowers respectively.

Differences in servicer modification rates cannot be rationalized by borrower observable characteristics alone, and are likely to reflect differences in institutional willingness or ability to grant modifications. The figures in Appendix D show predicted probabilities from simple probit models for borrower delinquency, modification and default before and after controlling for observable characteristics. Borrower controls include home value, principal balance, loan-to-value ratio, debt-to-income ratio, credit score and the ZIP-code of the the borrower at origination. I also control for the home value, principal balance and loan-to-value ratio in October 2009. Conditional on these observables, the probability of borrower 90 day delinquency

[26] I compute predicted probabilities for a mortgage in which the borrower is the primary resident of the property at the

is statistically similar across loan servicers within all estimation samples (though different across samples). Although these predicted rates suggest that borrowers had comparable financial health across servicers, the institutions display significant differences in their willingness to grant loan modifications to delinquent borrowers. The difference between the most generous and least generous servicer highlights the large gap in the modification rates between particular institutions. These descriptive findings are consistent with AGA [2011] and Agarwal et al. [2017]. Agarwal et al. [2017] point out that modification rates "cannot be accounted for by differences in contract, borrower or regional characteristics of mortgages across servicers". In their paper, the authors suggest that servicer-specific factors such as the quality and size of servicing staff played a role in overall modification rates. The cross-sectional heterogeneity in servicer modification rates provides the primary source of identifying variation in this paper. I discuss this further in section 5.2.

Differences in the conditional default probabilities predicted by the probit models are largely explained by bank modification practices. The banks least likely to offer loan modification tend to have the highest rate of defaults for loans that ultimately do receive a modification. This suggests that the servicers only offer debt relief to borrowers that are the most likely to end in default, relative to other servicers that are more generous with modification offers. For example, Bank of America and GMAC are the two least likely servicers to offer modifications to their delinquent borrowers across all four samples, and in turn, the borrowers that do receive modifications from them are the most likely to end up in default.

The conditional delinquency and default *without* modification probabilities provide further evidence that the probability of default *with* modification is driven by bank behavior, rather than borrower characteristics. Neither Bank of America nor GMAC borrowers enter delinquency at a different rate from borrowers of other servicers. Borrowers also enter default without modification at nearly identical rates to other servicers. Jointly these facts suggest that borrowers had similar unobservable financial health across servicers and that differences in observed outcomes were primarily driven by the modification practices of financial institutions.

5 Estimation and Identification

This section outlines the central features of my estimation procedure. Appendix F provides additional details on the model estimation algorithm that are omitted here. I estimate my model using a simulated maximum likelihood ("SMLE") approach that jointly matches the probabilities of loan outcomes to those observed in the data. The potential outcomes for a given loan are: 1) No delinquency, 2) Delinquency, modification and foreclosure, 3) Delinquency, no modification and foreclosure, 4) Delinquency, no modification and no foreclosure, and 5) Delinquency, modification and no foreclosure.

I form the probabilities of the five potential outcomes by conditioning on observables and integrating out the unobservable variables. The probability of a given outcome j for borrower i is given by:

$$\Pr(y_i = j|\mathcal{X}_i;\Theta) = \int_{\varepsilon,\xi,\eta} \Pr(y_i = j|\mathcal{X}_i, \varepsilon_i, \xi_i, \eta; \Theta) dG(\varepsilon, \xi, \eta) \quad j \in \{1, \dots, 5\}$$

- Where $\mathcal{X}_i$ includes all observables, x_i, w_i, p_i and $\tilde{p}_i$.
- Θ is the full set of model parameters.
- G is the joint distribution of the unobservables governing home utility and delinquency cost.

medians for all continuous borrower characteristics.

These probabilities are combined together to form the likelihood:

$$\mathcal{L}\big(Y, X; \Theta\big) = \prod_{i=1}^{n} \prod_{j=1}^{5} \big[\Pr(y_i = j | \mathcal{X}_i, \Theta)^{\mathbb{1}(y_i = j)}\big] \tag{5}$$

Taking the natural log of the likelihood yields the log-likelihood function:

$$\log \mathcal{L}\big(Y, X; \Theta\big) = \sum_{i=1}^{n} \sum_{j=1}^{5} \big[\mathbb{1}(y_i = j) \log\{\Pr(y_i = j | \mathcal{X}_i, \Theta)\}\big] \tag{6}$$

My SMLE estimation routine searches across the Θ space in order to maximize the log-likelihood reported above.

5.1 Parametric assumptions

To make estimation tractable, I make parametric assumptions on the distributions of the unobservable variables that govern home utility relative to default and delinquency cost.

Assumption 1. *Parametric form of borrower unobservables.*

$$\xi_i \sim \text{Normal}(0, \sigma_\xi^2)$$
$$\varepsilon_i \sim \text{Normal}(0, \sigma_\varepsilon^2)$$
$$\eta_i \sim \text{Log Normal}(0, \sigma_\eta^2)$$

I intentionally allow the home utility terms, ξ_i and ε_i, to take on both positive and negative values to reflect the fact that an individual will find a low relative value of repaying a loan in situations where they are financially distressed and the opportunity cost of repayment is high. The idiosyncratic delinquency cost, η_i, is assumed to be log-normal to rule out individuals who benefit from missing loan payments or acquire some financial value that is unrelated to a loan modification.[27]

Assumption 2. *Independence of borrower unobservables.*

$$\xi_i \perp \varepsilon_i \perp \eta_i$$

Here I assume that the unobservable components of H_i and Q_i are independent of one another. Correlation between H_i and Q_i is allowed through observable characteristics, x_i and w_i.

With these parametric assumptions in place, I can express the distribution of the borrower's value of keeping their home and delinquency respectively by:

$$H_i \sim \text{Normal}(x_i'\beta, \sigma_\xi^2 + \sigma_\varepsilon^2)$$

$$Q_i \sim \text{Log Normal}(w_i'\lambda, \sigma_\eta^2)$$

[27] In reality, liquidity constrained borrowers may benefit from being able to temporarily stop paying their mortgage. I relax the log-normal assumption as a part of my model extensions.

5.2 Model identification and necessary assumptions

The primary challenge to identification in the model arises from the presence of the ξ_i term, which drives heterogeneity in borrower home utility relative to default, and is assumed to be observed by borrowers and servicers but not the econometrican. The separate identification of the two home utility terms, ξ_i and ε_i, is not trivial because they affect borrower default outcomes in an identical way. The only difference between the two unobservables is that a servicer is assumed to set its modification policy with knowledge of a particular borrower's ξ_i, but without knowing the private information governing home utility, given by ε_i. Overcoming this challenge relies on differences in servicer modification rates that are unrelated to borrower characteristics. For cross-servicer variation to enable identification, I make two key assumptions.

Assumption 3. *No borrower sorting to servicers on unobservables.*

The no sorting assumption states that borrowers are not assigned to servicers on the basis of their *unobservable* characteristics. It implies that individual borrowers draw their home utility and delinquency cost heterogeneity from common distributions, irrespective of servicer affiliation.

The assumption does not rule out borrowers from sorting to servicers on the basis of *observables*. I allow certain servicers to cater more to quantifiably riskier borrowers, such as those with lower credit scores or higher interest rates. My sorting assumption leverages the fact that Fannie Mae's loan servicing data is rich at the time of loan origination, allowing me to control for almost all measures of riskiness known to the bank when the loan was first made.[28] The interpretation of this assumption is that the distribution of idiosyncratic shocks governing borrower delinquency and default following the 2008 financial crisis was independent of the factors influencing original mortgage origination and servicer assignment.

The no-sorting assumption is required for the model to be identified. The central aspect is that multiple servicers face different modification costs, c_j, but a common joint distribution of unobservable characteristics, G_θ. The borrower characteristics and servicers' identities will create cross-sectional differences in the share of delinquent loans receiving modification. After controlling for observable characteristics and servicer identity, the remaining differences in the share of loans receiving modifications for each servicer must be jointly rationalized by the common distribution of home utility heterogeneity known to the servicer (ξ_i) because servicers do not observe any other source of borrower heterogeneity. The no-sorting assumption can be relaxed by grouping servicers based on their borrower characteristics and assuming a common distribution within these smaller sub-groups, but servicer groups must still contain multiple institutions and each servicer within a group must face different costs of modification.[29]

Assumption 4. *Servicer modification costs do not depend on borrower characteristics.*

The cost assumption says that the servicer's cost of modification, c_j, does not depend on borrower or loan characteristics. This implies that a borrower's financial health or his loan conditions do not complicate the loan modification process. Assumption 4 is less fundamental to identification than the no-sorting assumption, and can be relaxed easily by putting structure on the relationship between modification costs and borrower characteristics.

Using the assumptions above, I can identify the distributions of individual heterogeneity in home utility and delinquency costs introduced by ξ_i, ε_i and η_i. I start by thinking about the servicers' modification

[28] At time of origination I know a borrower's credit score, debt-to-income ratio (allowing me to impute income), approximate geographic information, the number of borrowers on the loan and numerous other relevant loan characteristics.

[29] Time variation for a particular servicer may also be helpful here, but the model would have to be adjusted to leverage this source of variation.

decisions, which are driven by borrower characteristics (x_i and w_i), costs of modification (c_j) and the ξ_i draw. The no-sorting assumption and the fact that servicers offer different probabilities of modification due to differences in c_j lead to exogenous variation in modification outcomes, which allows me to learn about the distribution of servicer observed borrower heterogeneity.[30] This follows from:

1. After controlling for borrower observables, modification cost parameters, $c_j \in [1, \ldots, J]$, capture cross-sectional differences in modification outcomes across servicers. In the absence of a separate cost parameter for each mortgage institution, two servicers facing observably identical borrowers would award the same share of these borrowers with modifications.

2. From the no-sorting assumption, borrowers across different servicers draw their ξ_i from a common distribution that does not depend on the servicer's identity.

3. By the assumption that costs do not depend on borrower characteristics, differences in c_js will lead to different conditional modification rates across servicers for observably identical borrowers. Differences in servicer modification policies will also lead to different ex-post ξ_i distributions after conditioning on delinquency.

4. The *share* of loans receiving modifications across different servicers must be jointly rationalized by the common distribution of ξ_i draws. Neither ε_i nor η_i can explain differences in the shares of delinquent loans receiving modifications across servicers because servicers cannot observe either term when setting the probabilities of loan modification for borrowers.

With σ_ξ^2 pinned down, identifying the remain variance terms becomes more straightforward. As stated above, only the borrower heterogeneity observed by the servicer, σ_ξ^2, can jointly rationalize observed shares of modifications across all loan servicers in the data. The remaining differences in default outcomes for identically observable borrowers can only be explained the borrower heterogeneity that is unknown to the servicer, σ_ε^2. The delinquency cost plays no role in the default outcomes so σ_η^2 cannot explain observed differences in borrower default outcomes. With the distributions of ξ_i and ε_i pinned down by modification and default outcomes respectively, I can look at the remaining variation in delinquency outcomes for observably identical borrowers to identify the delinquency cost heterogeneity parameter, σ_η^2.

5.3 Conditional modification probabilities

I must solve for the conditional modification probabilities as part of the estimation. The probability that a borrower i, with observable characteristics X_i, will receive a modification, m_i, is a crucial input into the likelihood but cannot be seen directly in the data because it is a function of unobservables. These probabilities establish model equilibria for a given set of observables and model parameters, but are computationally costly to solve for in every likelihood iteration. To overcome this computational challenge, I pre-solve the loan servicer's expected cashflow maximization problem on a six-dimensional grid that accounts for borrower characteristics and unobservable parameters. The maximization routine loads in the grid and then draws modification probabilities from it for borrowers given a guess of model parameters. For points between grid values, I use linear interpolation.

Solving for a grid of conditional modification probabilities relies on evaluating the servicer's expected cashflow function with many different inputs. Specifically, the modification rate is a function of six variables:

[30] "Exogenous" in the sense that some aspect of the modification probabilities is totally unrelated to borrower characteristics.

Table 4: Estimated home utility parameters, by origination cohort

	Q2 2004 Sample		**Q2 2005 Sample**		**Q2 2006 Sample**		**Q2 2007 Sample**	
	Ests	Std. Err.	Ests	Std. Err.	Ests	Std. Err.	Ests	Std. Err.
β_0	\$33,556	144.8	\$26,740	193.2	\$-33,380	471.9	\$38,639	225.8
$\beta_{Home\ Value}$	\$1,549	1.2	\$2,758	3.1	\$2,636	3.9	\$3,205	2.7
β_{LTV}	\$-36,067	225.6	\$-71,855	5,899.1	\$-55,364	473.0	\$-81,662	241.9
β_{DTI}	\$-32,630	292.3	\$-17,025	967.6	\$-32,350	961.7	\$1,286	207.4
$\beta_{Credit\ Score}$	\$583	1.0	\$450	0.5	\$578	0.9	\$402	0.6
$\beta_{\Delta UE}$	\$-74,189	146.4	\$-77,875	111.8	\$-73,471	285.9	\$-101,997	107.3
$\beta_{Log\text{-}income}$	\$-810	12.9	\$228	6.9	\$564	35.3	\$1,867	31.1
σ_ξ	\$144,949	99.4	\$164,112	183.4	\$181,845	306.7	\$171,845	242.2
σ_ε	\$13,467	287.1	\$10,916	454.3	\$12,021	187.2	\$11,672	827.6
N	21,310		10,864		8,928		10,782	
Log LL	6,661		5,875		6,649		7,664	

Note: This table presents estimated model parameters that jointly explain the borrower home utility, H_i, for single-family Fannie Mae loans in the state of California. Parameters have been re-scaled so that they can be interpreted in terms of nominal dollar amounts. Estimation was conducted for each loan origination group separately, and there are no restrictions that would force parameters to be consistent across origination groups. Information asymmetry comes from the σ_ξ and σ_ε terms. These two parameters represent the standard deviations of unobservable borrower information that are either known to the servicer (ξ_i) or entirely private information to the borrower (ε_i). Un-scaled results used directly in estimation are reported in Appendix G.

$x_i'\beta+\xi_i, w_i'\lambda, p_i, \sigma_\varepsilon, \sigma_\eta$ and c_j.[31] The first three terms are observable characteristics of an individual borrower to the servicer. The fourth and fifth terms are the standard deviations of borrower private information relating to home utility (ε_i) and delinquency cost (η_i), while the c_j is the servicer-specific cost of processing a modification. The computation of cash flow maximizing m_i is highly conducive to parallel computing since the optimum for a given set of inputs can be computed independently from the optimum for a different set of arguments. I use the Texas Advanced Computing Center's Stampede2 cluster to calculate the optimal modification policy grid. My estimation routine employs a grid with 30 points in the $X_i'\beta+\xi_i$ dimension, 26 points in the p_i and c_j dimensions, 17 points in σ_ε dimension and 15 points for the $w_i'\lambda$ and σ_η dimensions for a total grid with 77.6 million points. Forming this grid takes approximately 20 minutes with 78 computing nodes.[32] Additional details on the formation of the loan modification grid are available in Appendix E.

6 Results

Table 4 reports financial equivalents for my borrower home utility parameters and Table 6 reports financial equivalents for estimated servicer costs. I report financial equivalents because model parameters are difficult to interpret directly. All estimated coefficients can be interpreted in terms of nominal dollar values because my implicit scaling assumption means that parameters are always relative to changes in the mortgage principal, p_i. The log-normal parameterization of the borrower delinquency cost distribution makes conversion of specific delinquency cost parameters into nominal dollar amounts less straightforward than for the home utility and servicer cost parameters.[33] As a result, I present the un-scaled delinquency parameter estimates in Table 18 of the Appendix. All un-scaled parameter estimates can be found in Appendix G.[34]

[31] To reduce dimensionality, I assume that $\tilde{p}_i$ and f_i are functions of p_i.

[32] For information about Stampede2, see: `https://www.tacc.utexas.edu/systems/stampede2`

[33] In part this is because the μ and σ parameters of a log-normal jointly affect the mean *and* the variance of the distribution.

[34] 'Scaling' here just means dividing dollar denominated amounts by 3,000 to reduce overflow errors in estimation. This is innocuous: for example, a \$300,000 home is entered as having a value of 100 in my code. To get back to nominal dollar amounts I then multiply my parameter estimates by this amount.

Table 5: Mean utility response to a 10% rise in home value and outstanding principal, by origination cohort

	Q2 2004	**Q2 2005**	**Q2 2006**	**Q2 2007**
Mean(10% of home value)	37,546	31,453	29,703	30,656
$E[\Delta u_i{}^{NQ}]$ w. 10% increase in home value	\$21,044	\$33,769	\$30,513	\$39,714
Mean home value elasticity, $E\left(\frac{\Delta 10\% u_i^{NQ}}{\Delta 10\% Home\ value}\right)$	0.82	2.17	2.22	2.77
Mean(10% of outstanding principal)	\$19,808	\$21,933	\$24,969	\$26,391
$E[\Delta u_i{}^{NQ}]$ w. 10% increase in outstanding principal	-\$22,703	-\$27,884	-\$30,949	-\$33,995
Mean principal elasticity, $E\left(\frac{\Delta 10\% u_i^{NQ}}{\Delta 10\% Outstanding\ principal}\right)$	-0.95	-2.11	-2.59	-2.46

Note: This table presents the average response of borrower utility to a 10% rise in October 2009 home value and outstanding principal balance, by origination cohort. Data for each origination cohort comes from the same month so nominal dollar amounts are comparable across years. Borrowers in later origination cohorts tend to have lower value homes and higher levels of debt by October 2009 than earlier cohorts. In turn, the model results suggest that borrowers in later cohorts behave more sensitively to changes in home value and outstanding principal. The $E[\Delta u_i{}^{NQ}]$ term reflects the average change in net utility for a borrower that avoids delinquency and repays debt; because the means of both ξ_i and ε_i are assumed to be zero, these terms have no effect on this expectation. "Utility elasticities" with respect to either a 10% change in home value or principal outstanding are elastic for all cohorts after 2004 and tend to become higher in later years. With the exception of 2007, borrowers are more sensitive to changes in outstanding principal than home equity.

Table 6: Estimated costs of loan modification, by servicer and origination cohort

	Q2 2004 Sample		**Q2 2005 Sample**		**Q2 2006 Sample**		**Q2 2007 Sample**	
	Ests	Std. Err.	Ests	Std. Err.	Ests	Std. Err.	Ests	Std. Err.
c_{BoA}	\$3,081	\$140.34	\$2,998	\$136.34	\$2,337	\$177.01	\$1,329	\$76.05
$c_{Wells\ Fargo}$	\$670	\$39.92	\$665	\$33.00	\$706	\$17.97	\$720	\$4.31
$c_{Citimortgage}$	\$3,188	\$223.40	\$3,077	\$161.05	\$2,327	\$287.53	\$2,052	\$366.86
c_{GMAC}	\$3,270	\$2,682.23	\$6,174	\$482.52	\$7,153	\$509.32	\$9,750	\$1,294.40
$c_{JP\ Morgan}$	\$1,856	\$421.62	\$679	\$27.52	\$6,480	\$752.34	\$5,055	\$729.44
$c_{Fannie/Seterus}$	\$1,212	\$347.44	\$646	\$19.50	\$595	\$30.74	\$521	\$35.98
$c_{Green\ Tree}$	\$572	\$88.81	\$616	\$46.19	\$514	\$65.37	\$538	\$56.98
N	21,310		10,864		8,928		10,782	
Log LL	6,661		5,875		6,649		7,664	

Note: This table presents estimated parameters for the costs of modification for single-family Fannie Mae loans in the state of California. Parameters have been re-scaled so that they can be interpreted in terms of nominal dollar amounts. Estimation has been conducted for each loan origination group separately and there are no restrictions that would force estimated costs of modification for the same servicer to be consistent across origination groups. In general, there are notable differences in the costs of modification between servicers and these differences tend to persist across origination groups. Higher cost servicers, such as servicers GMAC and JP Morgan, tend to be high cost across all estimation groups. Low costs servicers, such as servicers Green Tree and Wells Fargo tend to exhibit lower estimated costs to process loans. Differences in the estimated cost of modification *for the same servicer*, such as servicer JP Morgan, may reflect the servicer's differing treatment of mortgages based on their origination timing. These modification cost results highlight the role of servicer-specific factors in foreclosure prevention efforts. Un-scaled results used directly in estimation are reported in Appendix G.

My results suggest that borrowers put significant value on avoiding foreclosure. Using the full datasets of borrowers, the expected value of home utility minus payment disutility, $u_i^{NQ} = H_i - p_i$, for a single, non-delinquent borrower monotonically declines from $275,121 in 2004 to $153,120 in 2007, suggesting rising probabilities of borrower default for loans originated closer to the start of the 2008 financial crisis. These amounts can be interpreted as the mean values of avoiding foreclosure and paying to remain in a home for different borrower cohorts. There is significant value to borrowers from avoiding foreclosure: a foreclosed borrower in the U.S. cannot qualify for a conventional mortgage for seven years and will have a significantly reduced credit score that will restrict access to non-housing credit markets. Studies have also linked foreclosure to elevated levels of depression, divorce and suicide.[35] Applying the headline estimates from Bhutta et al. [2017] to my data suggests that a $300k home would need to have a $522k valued mortgage before the median borrower would choose to "walk away" from their mortgage obligation: this $222k gap is consistent with the magnitude of my estimates. It is also worth commenting on the decline in the net home utility over time, a result that is driven by both differences in β parameter estimates and increases in p_i as Fannie Mae increased its conforming loan limits.[36] This finding suggests that financially less healthy borrowers were taking on larger mortgage balances than financially more healthy borrowers in the lead up to 2008, even in the prime mortgage market.

The estimated results also suggest that borrower utility is more responsive to increases in debt burden than increases in home value, and that borrowers with loans originated closer to 2008 tended to be more sensitive to changes in home value and outstanding principal. Table 5 summarizes the estimated mean change in net borrower utility for a non-delinquent, repaying borrower associated with a 10% increase in home value versus a 10% increase in outstanding principal balance. For home value, a $1 increase tends to result in a less than a $1 increase in utility for the average borrower. On the other hand, a $1 increase in debt typically leads to a greater than $1 decline in expected utility. These patterns seem reasonable because home equity gains tend to be illiquid and involve transaction costs before they can be used for consumption. Increases in debt burden directly impact monthly spending and reduce a borrower's ownership of a property. It is worth highlighting that borrowers in earlier origination groups appear relatively less sensitive to home value increases but more sensitive to increases in their total outstanding debt. These differences in responsiveness provide evidence for the fact that loans originated closer to the 2008 financial crisis were also made to borrowers who were less risk averse towards debt than earlier borrowers.[37] Mean elasticities of utility with respect to home value and outstanding principal also reveal that borrowers are more sensitive to debt level changes than home value changes, with the exception of the 2007 cohort. For borrowers with loans originated closer to the 2008 financial crisis, the utility-elasticities of an increase in debt burden or home value become more similar.

The estimated standard deviation parameters in the home utility, σ_ξ (known to servicer) and σ_ε (unknown to servicer), suggest significant unobservable heterogeneity across borrowers, much of which is known to loan servicers. The large and statistically significant combined dollar measures for the two unobservable parameters mean that unobservable factors can ultimately drive borrower delinquency and default decisions. ξ_i likely captures much of this variation because a borrower often informs his bank about adverse life events,

[35] Tsai [2015] conducts a meta-analysis on 42 papers that explore the linkages between mental health outcomes and foreclosure. Fowler et al. [2015] provide an original study on the link between foreclosure and suicide rates.

[36] Between 2004 and 2007 single family, Fannie Mae loan limits rose from $333,700 to $417,000. See: `https://www.hsh.com/mortgage/a-history-of-conforming-fanniefreddie-loan-limits.html`

[37] Less risk averse towards debt in the sense that later borrowers perceive gains in home value similarly to equal-sized declines in outstanding balance.

even if it does not show up in my servicing data.[38] The importance of unobservables supports recent mortgage literature that assigns the vast majority of mortgage defaults to adverse life events rather than negative equity concerns (Ganong and Noel [2020b], Low [2022]). Though the scale of servicer-known heterogeneity is an order of magnitude larger than servicer-unknown heterogeneity, servicers still face economically meaningful uncertainty that impacts their decision making. Subsection 6.3 below explores the effectiveness of loan modification allocations in more detail.

While loan, property, and borrower characteristics all play an important role in shifting the likelihood of default, my results suggest that delinquency costs play a relatively insignificant role for the typical borrower. At the point estimates for delinquency cost parameters, a borrower with the mean credit score behaves as if the costs of delinquency are $1,014, $2,406, $998 and $1,394 respectively across origination groups, with standard deviations of $1,024, $4,417, $1,805 and $2,288.[39] These measures of delinquency cost are greatly outweighed by potential changes in payment disutility through modification and the cost of default, which tend to be upwards of tens of thousands of dollars. Low mean delinquency costs may reflect the fact that borrowers have a high opportunity cost of making payments on time and are not willing to forgo other spending to remain current on their payments. This sort of interpretation is consistent with thinking about temporary delinquency as a form of insurance (Bhutta et al. [2017]).[40]

As expected from observable modification activity and past papers on loan renegotiation, there are statistically significant differences in the cost of processing loan modifications across servicers and this heterogeneity plays an important role in modification and foreclosure outcomes. The lowest estimated modification costs are associated with Fannie Mae, which behaves as if the cost of processing modifications is negative in all four samples, albeit with noisy parameter estimates. Negative costs of modification should not be disqualified as unrealistic: as discussed in Section 2, HAMP servicers were eligible to receive thousands of dollars of subsidies for individual loan modifications. Institutions best setup to take advantage of these programs likely behaved as if they faced negative transaction costs of processing loan modification. In my results GMAC faces the largest costs of loan modification, with estimates that exceed $3,0000 in all four samples. In my model, these high costs rationalize GMAC's reluctance to grant loan modifications to its borrowers.

The modification cost parameters capture factors beyond the mortgage servicer's literal transaction costs. First, the modeling set-up forces all mortgage servicers to have the same expectation about future mortgage losses on delinquent loans. Due to this restriction, differences between servicer housing market expectations would most likely be captured in estimated modification costs. Second, the model does not allow servicers to differ in their ability to assess borrowers as a part of the loan modification process. Screening ability between banks would also be captured in the modification cost estimates. Third, large costs of modification could indicate that a servicer is understaffed or otherwise unable to service large quantities of loans: a servicer facing modification capacity constraints could be perceived as having an infinite cost of granting modification if constraints are not carefully accounted for.

Figure 6: Model fit, outcome probabilities

(a) Actual outcomes (b) Predicted outcomes

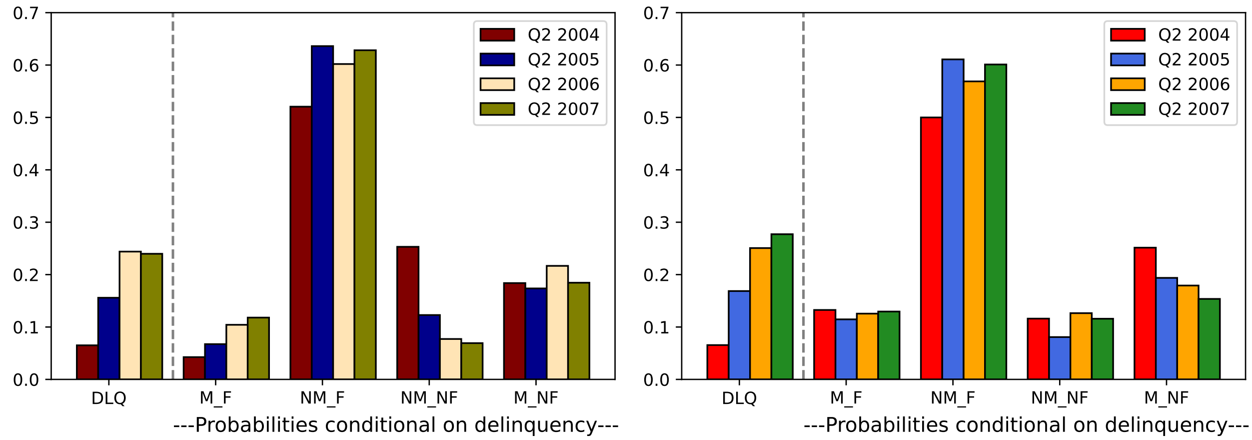

Note:
This graphic presents actual (left) and model predicted (right) outcome probabilities across origination cohort samples. The DLQ bars are the share of loans that become at least three months late on their mortgage payments. The other four groups represent outcome probabilities *conditional on delinquency*. Each delinquent outcome is a combination of a modification, M, and foreclosure, F, outcome. An outcome with an N signifies either *Not modified* or *Not foreclosed*. The model implied probabilities are ordinally correct. The model fits delinquency probabilities across origination cohorts almost exactly and predicts the correct magnitude for non-modified foreclosures.

6.1 Model fit

Figure 6 presents (a) the actual borrower outcome probabilities in the data, relative to (b) the model predicted borrower outcome probabilities. The model generally preserves both relative cardinality and ordinality of outcome probabilities: parameter estimates suggest that the majority of borrowers avoid delinquency all together and that delinquent borrowers are mostly likely to end in foreclosure, without having received a modification. The estimates slightly under-predict the probability of avoiding foreclosure given a modification and over-predict the probability of a foreclosure after a modification. Most obviously, the model suggests that servicers make more modification errors than is supported by the raw data.

There are various factors that affect the current model fit. Perhaps most importantly, the fact that over 80% of borrowers never enter delinquency means that estimation relies on a relatively small fraction of loans to fit the model to all outcome probabilities. The model performs well for the outcomes with the most observations: predicted probabilities of delinquency (DLQ) match the data almost exactly and predicted probabilities of foreclosure without modification (NM_F) also appear closely matched to observed outcomes. Discrepancies primarily arise with the two least likely outcomes: foreclosure after modification (M_F) and no foreclosure without a modification (NM_NF).

Discrepancies between observed loan outcome probabilities and the model predicted outcome probabilities are also at least partially explained by two key restrictions: 1. setting modification and foreclosure losses as a fixed share of outstanding principal and 2. the use of a sparse grid in estimating model parameters.

[38] To earn a loan modification, borrowers were explicitly required to provide evidence of adverse life events. Fannie Mae's Single Family Loan Performance data does not contain information about these interactions.

[39] The credit score parameter is insignificant at conventional significance levels for all origination groups, so changes in borrower credit score barely shift the mean of this distribution.

[40] I separate delinquency and default more explicitly than much of the mortgage literature, which is often focused on a measure similar to my measure of delinquency (60 or 90 days behind on mortgage payments). It seems credible that borrowers are willing to go delinquent when facing a liquidity constraint but would still try to avoid the high costs of foreclosure.

Figure 7: Observed modification rates and predicted modification rates, by servicer and cohort

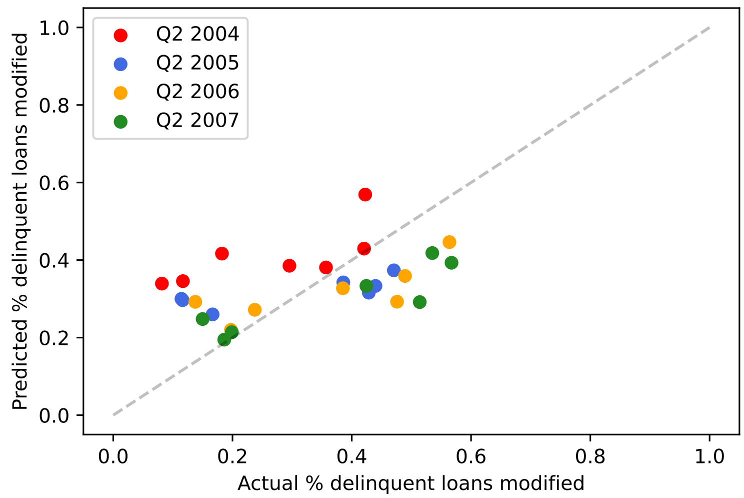

Note:
This graphic presents a scatter plot of the actual share of loans modified (x-axis) against the model predicted share of loans modified (y-axis), by individual servicer. As demonstrated by the grey 45-degree line: model predicted modification probabilities tend to align closely with actual modification rates, with no clear pattern of over-prediction or under-prediction.

Firstly, as previously discussed, Figure 5 highlights the fact that financial institutions realized a wide range of potential losses through modification and foreclosure and the current model does not allow for this. This is not as extreme as it may seem at face value: servicers would not have known future loss realization when setting modification policy.[41] Secondly, the sparse grid used in estimation may introduce error due to the linear interpolation used to approximate the servicer's policy function. I mitigate this second potential challenge by solving the servicer's problem at tens of millions of potential states of borrower characteristics and model parameters.

Counteracting potential concerns about the use of the modification rate grid, the model's predicted probabilities are closely aligned with observed data at the servicer level. Figure 7 presents a scatter plot of the actual percentage of delinquent loans that received a modification relative to the the predicted share of delinquent loans that were modified, at the individual servicer level. Predicted modification probabilities are highly correlated with observed shares of modification for all servicers.

6.2 Benefits of modification

Model estimates suggest that there are meaningful benefits from appropriately targeted modification from both the private and social perspective. My results suggest that servicers have a strong incentive to grant modifications to borrowers that would enter foreclosure otherwise. This incentive is counteracted by meaningful, though relatively smaller, losses if modification is awarded to borrower types whose foreclosure statuses are unchanged by financial assistance. Borrowers also meaningfully gain if they are able to avoid foreclosure. Since modifications function as a transfer between the borrower and his servicer, social losses from incorrect allocation of modifications tend to be relatively lower than the private losses considered by a loan servicer.

[41] There is scope to relax this assumption by allowing foreclosure losses to depend on borrower observable characteristics.

Table 7: Mean gains/losses from modification as percentage of outstanding principal

Result	Q2 2004	Q2 2005	Q2 2006	Q2 2007
Borrowers				
Gain - Medium borrower type avoids foreclosure	\$8.03%	\$7.75%	\$7.55%	\$7.78%
Servicers				
Loss - Low borrower receives modification	\$-1.25%	\$-1.23%	\$-1.05%	\$-0.79%
Gain - Medium borrower avoids foreclosure	\$23.90%	\$23.87%	\$24.04%	\$24.35%
Loss - High borrower modification	\$-16.50%	\$-16.46%	\$-16.34%	\$-15.91%

Note:
This table presents mean gains and losses experienced by both borrowers and servicers through modification outcomes, as a percentage of outstanding principal. The model estimates suggest that high borrower types may receive more relative utility from modification than medium types, and this is primarily driven by the fact that medium borrower types have lower home values relative to default, as well as lower outstanding principal balances. Servicers tend to take small relative losses when modifications are incorrectly rewarded to low borrower types, but financial institutions take relatively large losses from incorrectly granting aid to a borrower that could have paid without assistance (high types). Benefits from awarding modifications to medium types that avoid foreclosure outweigh the losses of an incorrect modification to either a low or high borrower type. My model assumes that high borrower types never enter foreclosure so their utility gain is always equal to the 15% of outstanding principal that is forgiven through modification. I also assume that the low borrower type receives no benefit from modification and always ends up in foreclosure. As a result of this assumption, modifications have no effect on this borrower type's welfare. The loss to the servicer from modifying a low borrower type arises from the cost of processing the modification, c_j.

Private benefits and losses

Using my estimates I can quantify the private benefits of foreclosure prevention. Table 7 summarizes the mean gains and losses from modification as a percentage of the outstanding loan principal. The average gain in utility for a borrower who successfully avoids foreclosure will be given by: $E(H_i - \tilde{p}_i | H_i \leq p_i, H_i > \tilde{p}_i)$, which is the average home value relative to default minus payment disutility for the medium-type borrower. I find on average that this benefit is between 7.5% and 8.0% as a percentage of outstanding borrower principal, depending on the origination sample. In nominal dollar terms this is equivalent to about \$23,000.

Average servicer gains for the medium-type borrower can also be calculated as: $E(\tilde{p} - f_i - c_j | H_i \leq p_i, H_i > \tilde{p}_i)$, which I find to be over 20% of the value of the outstanding principal due to large expected losses from foreclosure.[42] These benefits must be weighed against the losses from incorrectly awarded debt relief. Creditors will on average lose $E(p_i - \tilde{p}_i | H_i > p_i)$ when they award relief to the high-type borrower. I estimate that servicers lose an average of 15.9% to 16.5% of the outstanding mortgage principal per borrower who could have avoided foreclosure without assistance, across origination samples. Again in nominal dollars, this mean loss upwards of \$41,000.

Social benefits

The social benefits of modification greatly outweigh the private benefits to financial institutions, a result driven by the fact that modification acts as a transfer between individual and servicer. Figure 8 depicts both servicer and social losses from (a) a "modification error" in which a modification was awarded that did not affect ultimate foreclosure statues, and (b) a "no modification error" in which a modification would have shifted a borrower away from foreclosure but was not offered. Mean servicer losses from modification errors greatly outweigh the mean social losses from modification, because any reduction in servicer cashflow

[42] Over \$70,000 in nominal dollar terms for the servicer.

Figure 8: Mean servicer error losses, private vs. social

(a) Mean losses from *modification* error (b) Mean losses from *no modification* error

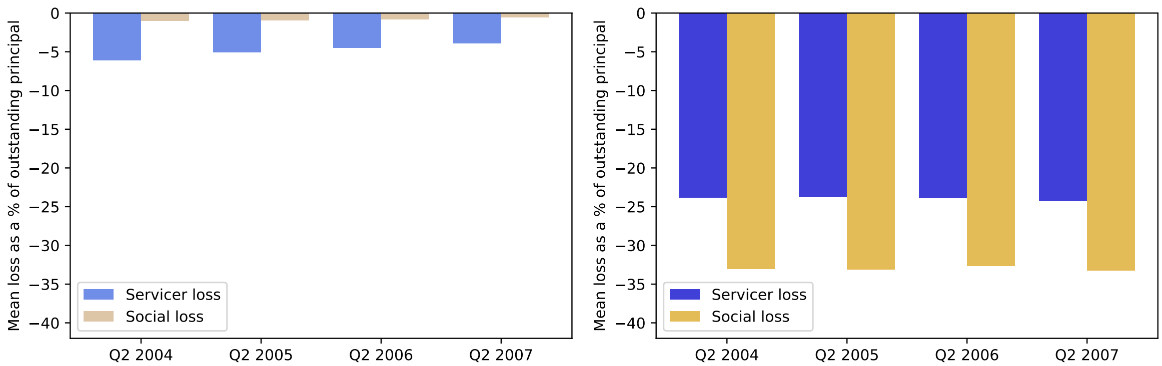

Note:
This graphic presents stimulated mean modification error losses (left) and no modification error losses (right) across origination cohorts in the samples based on model estimates. The main take-away is that the mean losses of "no modification" error, or failing to modify a borrower that could have avoided foreclosure, greatly outweigh the losses from a "modification error", modifying a borrower whose foreclosure status is unaffected by assistance. Though servicers face large private losses from a no modification error, the wider social loss is even greater, since borrower welfare declines meaningfully through foreclosure. Social losses from a modification error also tend to be smaller than the servicer's private losses from incorrectly awarding financial assistance. Taken together, this implies that a social planner maximizing social welfare would prefer to offer even more modifications than the servicer, because the relative losses from modification error are low and the losses from no modification error are relatively higher.

from modification directly improves the borrower's welfare. Social losses from modification error are driven exclusively by the avoidable c_j costs. The model predicts that a servicer will perceive modification errors as a loss of about 5% of the value of the outstanding mortgage principal. For comparison, the social cost of a modification error is generally below 1%.

While servicers face greater losses from modification errors than a social planner, they also absorb smaller losses from "no modification errors" than the social loss. Figure 10b shows the average servicer and social loss as a share of outstanding mortgage principal. The servicer perceives the mean loss from a no modification error to be about 22% of the value of the outstanding mortgage balance while the total social loss is in the order of 33%. This gap is explained by the loss in welfare faced by the borrower that is forced into foreclosure. This loss of welfare is not internalized by the servicer. Taken together with greater direct losses from modification errors, it is clear why a servicer will offer fewer loan modifications than a social planner who looks to maximize the combined welfare of the servicer and the borrower.

6.3 Accuracy of the servicers

In addition to quantifying the losses arising from misallocated modifications, my model estimates allow me to simulate the accuracy of servicer decisions. "Accuracy" here simply reflects whether the servicer made the correct modification decision ex-post. Servicers granting modifications to a large number of borrowers who cannot avoid foreclosure are "inaccurate", while servicers that only grant modifications that succeed in preventing a foreclosure are "accurate".

Table 8 reports the model predicted accuracy of servicers using my Q2 2007 sample. Each cell of the table contains the percentage of delinquent loans based on whether modifications would have been successful in changing a foreclosure outcome (rows) and whether they actually received a modification (columns). Estimates show that servicers tend to make the correct modification decisions and prefer to grant too many

Table 8: Predicted servicer modification accuracy, Q2 2007

		Servicer action	
		No modification	Modification
State of	Modification has no effect	66.5%	16.6% *(Error: Modification)*
the world	Modification is effective	5.2% *(Error: No modification)*	11.7%

Note:
This table presents model-predicted servicer accuracy using the estimated results from my Q2 2007 sample. Each percentage in the table is the fraction of delinquent loans with the particular realized outcome. Around 72% of the time the servicer makes the correct modification decision: either not offering a modification when it has no effect or offering one when it will help avoid a foreclosure. In equilibrium, servicers prefer to make more modification errors than non-modification errors. This is seen in the fact that 21.8% of delinquent loans receive a modification even when it makes no difference to their ex-post foreclosure outcome (recalling that servicers face information asymmetry when setting their modification policy). This behavior is rationalized by the fact that servicers gain significantly when a modification succeeds in preventing a foreclosure. The losses from a modification error are smaller than from a no modification error. The model also suggests that only around a quarter of delinquent loans stand to benefit from modification, meaning that many foreclosures would not have been preventable without additional modification generosity. For reference, Appendix I depicts simulated modification accuracy shares for all four origination cohorts. Model-predicted servicer modification effectiveness tends to be consistent across all four origination cohorts.

modifications rather than too few. I find that servicers make the correct modification decision 78% of the time; the rest of the time they either grant a modification when it has no effect, or fail to grant a modification when it would have prevented foreclosure. The servicer preference for modification can be seen from the fact that they make many more "modification" errors (16.6% of delinquent loans) than "no modification" errors (5.2%). I report the predicted servicer modification accuracy for other estimation samples in Appendix I. The predicted servicer modification accuracy tends to be consistent across all four origination cohorts used for estimation.

The expected losses from misallocation of modifications rationalize servicer modification practices. As shown in the previous section, servicers face much larger expected losses from failing to grant modifications that could prevent foreclosure, than from awarding a modification that was unnecessary. Though servicers tend to commit many modification errors, the large social losses from failing to prevent foreclosure suggest value in granting even more modifications to borrowers than the servicer finds optimal. The next section considers additional counterfactual policies using my model results.

7 Counterfactual analysis

The sizeable differences between private and social losses discussed in Section 6 suggest scope for policy intervention in foreclosure prevention. The gap between losses is driven by the fact that mortgage servicing institutions do not internalize borrower welfare losses from foreclosure, which can lead them to have an insufficient incentive to offer modifications from a social welfare perspective. The inclusion of negative foreclosure externalities on local housing markets (Campbell et al. [2011], Anenberg and Kung [2014]) would further widen the gap between private and social losses, the analysis presented here does not include these externalities.

In this section of the paper I use my estimated model to test counterfactual policy solutions that could have been used to achieve greater levels of social welfare. I start by comparing the model-predicted servicer behavior to two informative benchmarks: (a) the "First Best" case in which there is no information asymmetry and no inefficient foreclosure occurs, and (b) the "social planner" problem in which a hypothetical social planner sets modification policy to maximize total social welfare, taking as given information asymmetry

Table 9: Modifications and foreclosures, relative to First Best

	% Modifications	% Foreclosures
First Best	100.00	100.00
Social planner	298.48	104.93
Baseline: Servicer	153.74	111.52

Note:
This table summarizes the model-predicted number of modifications and foreclosures using the estimated results from my Q2 2007 sample. The table depicts the baseline case and the social planner's problem as a percentage of the First Best benchmark. In the First Best case, mortgage servicers have full information about borrower financial health and only award modifications when they are effective in preventing foreclosure. Modification rates are significantly higher under both the baseline and social planner's problem since modifications cannot be perfectly targeted to those borrowers who need them to avoid foreclosure. The social planner would award significantly more modifications to borrowers than servicers in order to reduce the social cost of foreclosures that could have been prevented.

Table 10: Borrower and servicer surplus, relative to First Best

	% Total surplus	% Borrower surplus	% Servicer surplus	% Medium-type surplus	% High-type surplus
First Best	100.00	100.00	100.00	100.00	100.00
Social planner	99.26	100.97	98.20	82.31	102.33
Baseline: Servicer	99.11	98.82	99.28	57.81	100.34

Note:
This table summarizes the model-predicted borrower and servicer surplus using the estimated results from my Q2 2007 sample. The table depicts the baseline case and the social planner's problem as a percentage of the First Best benchmark. In the First Best case, mortgage servicers have full information about borrower financial health and only award modifications when they are effective in preventing foreclosure. Medium-type surplus tends to improve significantly with a decline in foreclosures. High-type borrowers benefit from higher rates of loan modification when servicers cannot perfectly screen: they play a significant role in driving up total borrower surplus in the social planner case.

and servicer modification costs.[43] I then predict equilibrium outcomes under counterfactual circumstances where: 1) servicers receive further subsidies that lower loan modification costs, and 2) servicers hold more information about borrower financial positions.

7.1 Policy benchmarks

Table 9 and Table 10 summarize the relative welfare of servicers and borrowers when comparing the predicted model outcomes with the servicer to the First Best and social planner solutions.

In the First Best case, servicers have complete information about borrower outcomes and are able to stop all preventable foreclosures. With full information, servicers offer no modifications to the low- or high-type borrowers for whom a modification will not change equilibrium outcomes. Based on my model estimates for Q2 2007, there would be 53.7% more modifications but also 11.5% more foreclosures under the base case, in which servicers maximize their expected cashflows, relative to the First Best case. Modifications decline in the First Best case because servicers face no uncertainty about which borrowers would benefit from such assistance: each modification offered successfully prevents a foreclosure. Medium-type borrowers are made significantly better off under this hypothetical scenario: their combined welfare improves by 44.0% relative to the base case in which a servicer is not fully informed about borrowers. Conversely, the high-type borrowers lose out because they never receive modifications when a servicer is fully informed; their combined welfare falls by around 0.3%. This decline is small because the vast majority of high-type borrowers never enter delinquency at all.

Under the social planner's problem, a theoretical planner attempts to maximize joint welfare of borrowers

[43] Again, the social planner's problem does not consider the negative externalities associated with foreclosure. Neither of the benchmarks are achievable in and of themselves, but still offer a useful comparison to benchmark equilibrium outcomes and outcomes under alternative policy proposals.

and servicers, but faces the same information asymmetry and modification costs as the servicers themselves. Due to the larger social costs of foreclosures, the planner offers significantly more modifications than the servicers offer in practice: there would be around two times more modifications than the servicers would offer on their own and around three times more modifications than in the First Best case. The dramatically increased rates of modification help drive foreclosures down; under the social planner's solution there are only 5.0% more foreclosures than in the First Best case. The sharp increase in modifications comes at the cost of modifications being misallocated to borrowers who do not need them. High type borrowers enjoy total welfare that is 2.3% higher because many of them receive modifications that decrease their debt burden.

These two policy benchmarks suggest a few important aspects of equilibrium outcomes. Firstly, given a fixed level of modification generosity, even a fully informed servicer can only prevent around 10.3% of foreclosures that occur in the model-implied equilibrium. This means that many foreclosures simply cannot be prevented without increased generosity of modification. Some borrowers could certainly have been under sufficient financial duress that no reasonable level of forgiveness would have allowed them to keep their homes. In a few other cases, borrowers may have also had low attachment to their properties and been interested in walking away from an underwater mortgage.[44] The benchmark results also highlight the tension between incorrectly awarding modifications and successfully preventing foreclosures. Though total borrower welfare under the First Best, Servicer and social planner outcomes appear similar, this can be misleading because it arises from the fact that high-type borrowers benefit from loan modifications. In practice, a policy-maker wants to carefully consider the incorrect allocation of modifications to borrowers who do not need them because this hurts financial institutions without providing meaningful gains in social welfare.[45]

In addition to thinking about theoretical benchmarks, my model enables an analysis of equilibrium outcomes in the mortgage market under alternative foreclosure prevention policies. I consider these in the next section.

7.2 Counterfactual policies

To inform policy on foreclosure prevention, I use my model to explore the roles of modification subsidies and servicer information about borrowers on equilibrium outcomes. I consider the effects of subsidies and information together to develop a more nuanced understanding of the borrower screening and debt-relief allocation. The explicit treatment of information as a part of equilibrium outcomes is a unique contribution of this paper.

Modification subsidies reduce the effective transaction costs that servicers face when rewarding modifications and can help reduce the losses associated with debt forgiveness. Though subsidies will reduce transaction costs to the servicer, they will not guarantee more modifications. Servicers will still consider their expected cash flows with and without modification and then choose a policy that maximizes their expected cash flows. In my model a subsidy, s, reduces the size of c_j in the following way:[46]

$$\hat{c}_j = c_j - s$$

A servicer's information about his borrower's financial position can also play a crucial role in modification

[44] Studies have generally disproved underwater mortgages as a primary driver of borrower default, but there is still evidence that this occurred from time to time (Guiso et al. [2013]).

[45] It's also important to remember that in practice excessive forgiveness impacts not only servicers, but also investors and guarantors of mortgage returns like Fannie Mae and Freddie Mac.

[46] A *reduction* in subsidies can be expressed with a negative s, which increases the cost of processing modifications.

Figure 9: Effect of modification subsidies and servicer information on equilibrium outcomes

(a) Delinquency

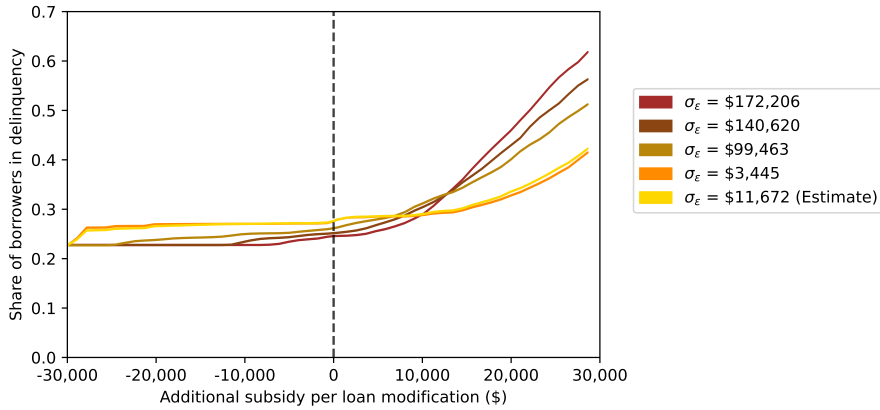

(b) Modification

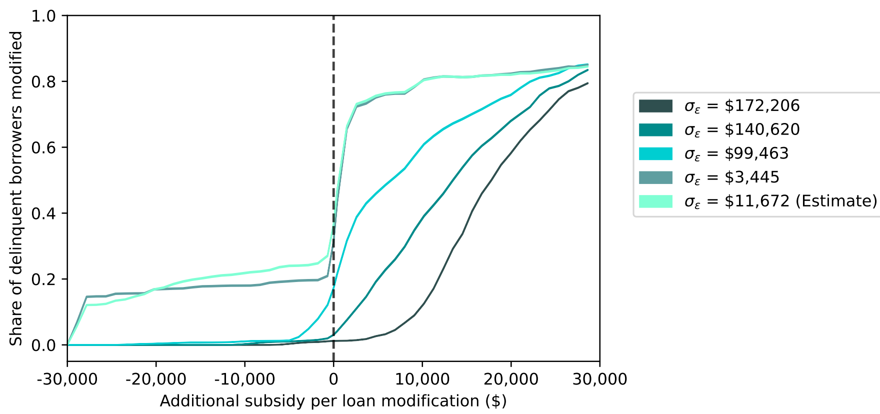

(c) Foreclosure

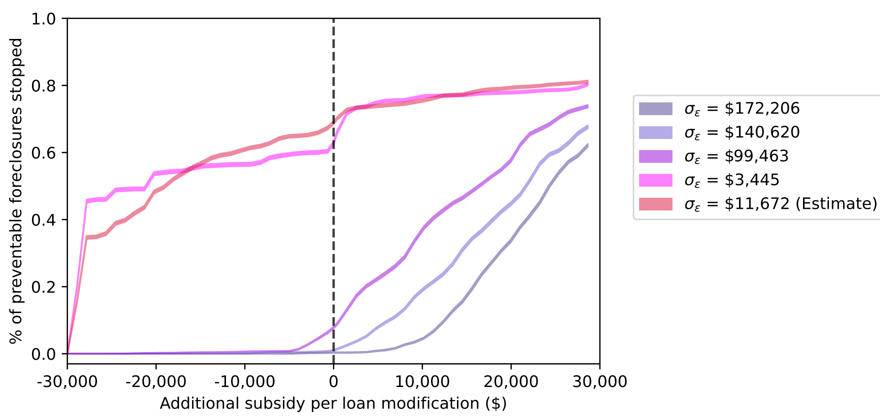

Note:
This graphic presents model-predicted effects of loan modification subsidies and lender information on equilibrium modifications, delinquencies and foreclosures for the Q2 2007 origination cohort. Each line represents a different assumed level of servicer information. When servicers observe a greater share of the $Var(H_i)$, they are more informed. In the baseline estimates the standard deviation of borrower private information in their home utility was $11,672. Each line represents the mean outcome for 100 Monte Carlo replications of the model given a particular assumed subsidy and servicer level of information. The shaded regions around each line are the 95% confidence intervals derived from these replications. Similar graphics for all other cohorts are presented in Appendix K.

decisions. Servicers who are fully-informed about borrowers' financial health can prevent all inefficient foreclosure and have a financial incentive to do so. In fact, perfectly-informed servicers don't need to be subsidized because they have sufficient private incentive to avoid the mean losses associated with making *no modification errors*. Policy only starts to play a role precisely when servicers face information asymmetry and have insufficient private incentive to grant modifications that would be socially beneficial.

To change the amount of information that loan servicers have about borrower unobservable financial information in the model, I shift the share of variance in H_i that is known to the servicer when he makes loan modification decisions. Information about H_i is central to the loan servicer's problem because a borrower's permanent default and foreclosure will drastically impact payouts; temporary delinquency has negligible effects on the servicer if a borrower ultimately resumes paying.[47] I vary the share of $\sigma_\xi^2 + \sigma_\varepsilon^2$ that is explained by σ_ξ^2. As the share explained by σ_ξ^2 rises, the servicer becomes more informed about borrower private information while the overall distribution of $\xi_i + \varepsilon_i$ is preserved. I provide further detail for this approach in Appendix J.

The model predictions from shifting the per-modification subsidies and servicers' level of information reveal the importance of effective borrower screening for subsidy policy. Figure 9 presents the model predicted effects of loan modification subsidies and lender information on equilibrium modifications, delinquencies and foreclosures for the Q2 2007 estimation sample. Information alone dramatically shifts equilibrium outcomes, which can be seen by comparing the model predictions along the dotted vertical line where zero additional subsidies are offered. As servicers become more informed about borrower financial health, they are better able to identify borrowers who require debt relief in order to avoid foreclosure. More informed servicers increase the number of modifications available and reduce foreclosures, even in the absence of government subsidies. An uninformed servicer will offer nearly zero loan modifications because it is unlikely that they will be allocated to the correct borrowers. As servicers observe a greater share of borrower heterogeneity, they dramatically increase the number of modifications offered. At the estimated level of borrower private information, servicers award modifications to around 35% of delinquent borrowers and this results in a substantial decline in foreclosures: they fall by about 20% relative to a servicer that observes close to 0% of borrower heterogeneity driven by $\sigma_\xi^2 + \sigma_\varepsilon^2$, the variance of home utility relative to default.

Equilibrium delinquencies also shift with changes in servicer information. As servicers become more informed and offer a greater number of modifications, delinquencies rise as financially healthy borrowers attempt to earn loan modifications. This pattern reverses once servicers become more highly informed; the closer servicers get to perfect information, the less likely they are to mistakenly offer a modification to a borrower going delinquent strategically. Strategic borrowers respond to more informed servicers by choosing not to enter delinquency, in order to avoid the costs of late payment. Equilibrium levels of delinquency lie above the First Best level precisely because the perfectly-informed servicer will never make any modification errors, completely eliminating the benefits of strategic delinquency.

Given the effects that the servicer's information has on equilibrium outcomes, knowledge about borrowers also dictates the effectiveness of subsidy policy. In Figure 9, movement along the horizontal axis represents different levels of per-modification subsidies granted to loan servicers. As servicers become more informed about borrower financial health the outcome lines steepen, which reflects that servicers are more responsive to modification subsidies. An uninformed servicer ($\sigma_\varepsilon =$ \$172,206) essentially cannot be incentivized to grant loan modifications because of the large losses associated with awarding them to borrowers who do not need

[47] Knowledge about Q_i allows the servicer to set a stricter modification policy for borderline borrowers for whom a reduction in the modification probability would push them out of delinquency entirely, but this only explains a negligible group of borrowers.

aid. For the uninformed servicer, modifications only increase to around 10% of delinquent loans, even if they are receiving $10,000 of additional subsidy *per modification*, and foreclosure levels are effectively unchanged. The story differs for an informed servicer: at the estimated level of servicer informed-ness, subsidies can increase the number of modifications from around 20% of delinquent borrowers to over 70% of them, and most of this increase occurs with additional subsidies between $0 and $4,000 per modification. I find that the increase in modifications for the more informed servicer would lead to as much as a 2.4% decrease in the total number of foreclosures. Delinquency behavior also reflects how effective subsidies are in reaching the target population. If servicers are relatively uninformed, modification subsidies will tend to jointly drive up modification and delinquency, leading to a muted effect on foreclosures, as modifications are awarded to borrowers who do not need them or do not benefit from them. For highly informed servicers, modification subsidies have diminishing effects on borrower delinquency and a larger share of debt relief goes to preventing foreclosures.

Two main factors drive the result that foreclosure prevention requires a substantial shift in total modifications. Firstly, observed modifications took place under active Federal subsidy policies, meaning that any subsidies discussed above were *additional subsidies* on top of those that were granted to financial institutions in reality (and they were labeled as such in the previous section) and incentive subsidies are likely to suffer from diminishing returns. Secondly, model estimates suggest that the majority of borrowers can be categorized as either Low Types or High Types, which means that modification is unable to shift outcomes for the a significant share of borrowers. The combination of imperfect information about borrower financial health and a small share of Medium Type borrowers makes it challenging for servicers to grant loans to precisely the borrowers that need them to avoid foreclosure.[48]

7.3 Assessing subsidy policy and the costs of foreclosure prevention

The empirical model enables a back-of-the-envelope calculation on the effectiveness of true foreclosure prevention policy. Loans in my sample would have generally been eligible for HAMP assistance since they were originated before the start of 2009. I consider the equilibrium response of modifications and foreclosures in a hypothetical, no-subsidy scenario by increasing the transaction costs by the amount servicers were entitled to under policy rules.[49] The same assumption also allows me to quantify the total subsidy cost of foreclosure prevention to the government.

I simulate loan outcomes for each origination sample assuming that no subsidies are offered to mortgage servicers, and then compare them to the baseline case in which subsidies were available. To do this, I make the assumption that the average loan subsidy awards $5,000 directly to a loan servicer for processing a modification. In practice, this means that I simulate the model with new costs of $c_j^{No\ sub} = c_j + \$5,000$ and compare equilibrium outcomes to baseline model simulations. The dollar amount is taken from Hembre (2018), which notes that HAMP yielded $5,000 of direct payments to financial institutions. As a part of the process, I simulate the model for each separate origination cohort with and without subsidies using 100 Monte Carlo replications.

[48] In fact, there will also be many Low Type and High Type borrowers that appear identically observable to Medium Types to the servicer. This clearly challenges efficient debt-relief allocation.

[49] The calculation is "back-of-the-envelope" because I cannot observe the subsidy that a servicer received for a specific modification. Subsidies under HAMP were standardized across borrowers and servicers but subsidy entitlements were conditional on future borrower repayment and the type of loan being managed. My counterfactual exercise avoids making assumptions about a bank's subsidy expectations by simply using a flat $5,000 payment to approximate reality.

Figure 10: Estimated effect of modification subsidies on equilibrium outcomes

(a) Change in modifications (b) Change in foreclosures

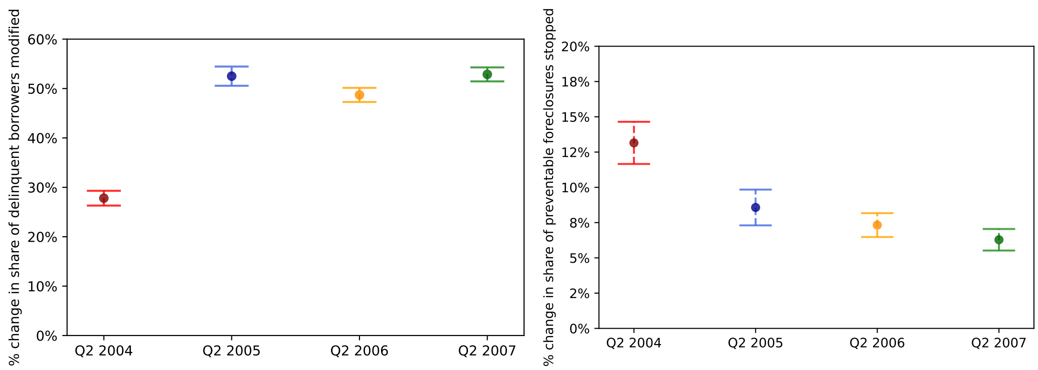

7.3.1 Effects on modifications and foreclosures

The counterfactual exercise suggests that modification subsidies did increase modifications and reduce foreclosures, but that the magnitude of effect varied across origination cohorts. Figure 10 reports the means and standard errors for the the percentage changes in the share of delinquent borrowers receiving a modification and the share of preventable foreclosures stopped. For loans originated in the second quarter of 2004, the mean result suggests that subsidies caused the share of delinquent borrowers receiving modifications to rise by 27.8% and that this in turn helped increase the share of preventable foreclosures stopped by about 13.2%. The share of delinquent loans that receive modification rises by even more in later cohorts, but the foreclosure impact is less suggesting that modifications are less effective at preventing foreclosure for loans originated closer to 2008. For Q2 2005, the counterfactual analysis suggests that there was a rise of 52.5% in the share of delinquent borrowers that received modification due to Federal subsidies but that there were was an only 8.6% rise in the share of preventable foreclosures stopped. For both Q2 2006 and Q2 2007, the change in the share of modifications rises by more than 45% and the share of preventable foreclosures stopped rises by less than 8%.

The differences between origination cohort outcomes are driven by both the underlying borrower financial health and mortgage servicer behavior. As discussed in the results section, the mean value of keeping a home relative to defaulting progressively declines for loans originated later in the 2000's. More financial distress means that there will be greater shares of Low- and Medium- Type borrowers relative to the High- Type borrowers. Servicers only want to offer assistance to Medium-Type borrowers that can avoid foreclosure through debt relief, if only the share of Low-Type borrowers grows, then they have no increased incentive to grant modifications. Estimated transaction costs of loan modification differ within-servicer across estimation samples. For example, the results for Bank of America and Citimortgage suggest that two of the biggest loan servicers were more willing to offer loan modifications to the 2007 origination cohort than to any of the earlier ones. At face value, this suggests that processing modifications was somehow more labor-intensive or otherwise costly to complete for older mortgages. In reality, it may also suggest differences in bank expectations about particular cohorts or specific knowledge that they held about borrower groups. High loan modification costs in the model may also reflect bank capacity constraints to process loan modifications driven by factors such as staff-shortages. The trend of decreasing modification costs for later cohorts was not present for all other loan servicers, potentially highlighting a difference between financial institutions.

A key takeaway here is that subsidies alone may not be sufficient to influence servicer behavior if they are dealing with non-cost related impediments to processing modifications.

7.3.2 Welfare effects of policy

Figure 11: Optimal subsides and welfare without foreclosure externalities

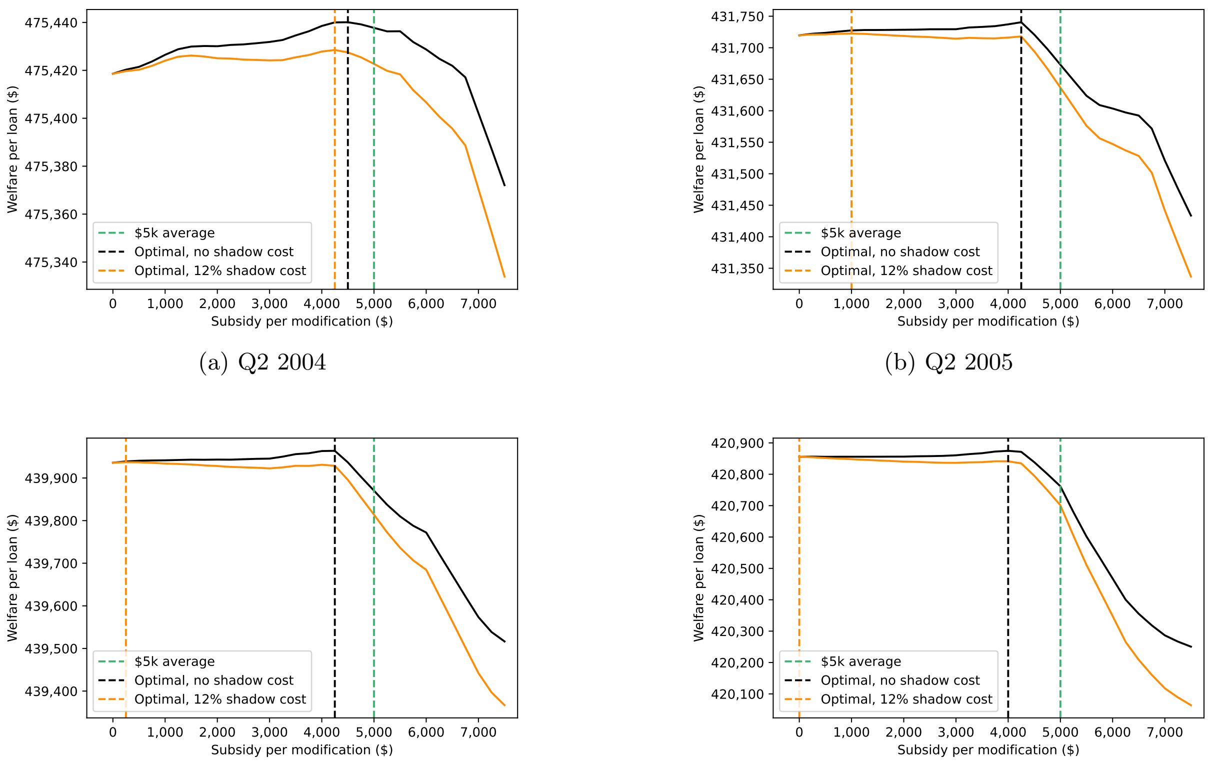

(a) Q2 2004

(b) Q2 2005

(c) Q2 2006

(d) Q2 2007

A welfare calculation becomes necessary to assess the appropriateness of loan subsidy policy as a means of foreclosure prevention. Leveraging the model estimates, I am able to measure the joint welfare of mortgage borrowers and servicing institutions and weigh it against government subsidy spending.

Figure 11 presents the welfare per loan for all borrowers at different levels of subsidy-per-modification offers and with different assumed social costs of government spending. The results suggest that some subsidy spending was worthwhile from a social welfare perspective, but that subsidy spending had a relatively small effect on overall social welfare. The assumed $5,000 per-modification level exceeds the socially optimal level for all origination cohorts, but welfare per loan does not appear sensitive to policy until around $4,000 per modification. At this point average welfare begins to decline due to a rise in unnecessary transaction costs associated with processing modifications that fail to prevent foreclosure. I find that the optimal level of subsidy spending tends to be higher in earlier origination cohorts: the optimal level exceeds $1,000 per-modification in both 2004 and 2005, but subsidy offers only slightly above $0 would have been socially optimal in 2006 and 2007. Again, these results reflect the fact that borrower distributions and servicer behavior across origination years influenced the effectiveness of Federal policy.

The structural model only provides a partial welfare analysis because it lacks several participants in the housing market. The model only provides meaningful quantification for existing borrower welfare and servicer cashflows. One population that might be affected by foreclosure policy is the households located close to pending foreclosures that are negatively impacted by proximity to reclaimed homes. Negative foreclosure externalities will generally increase the social costs of foreclosure, driving up the marginal welfare benefits from increased modification activity and reduced foreclosure. A different group, that may actually benefit from foreclosure, would be future home buyers who benefit from large discounts offered by banks through foreclosure auctions.

As suggested in the graphics above, the shadow cost of government spending may also play an important role in assessing the effectiveness of a subsidy policy. In a seminal work, Browning [1976] finds that the marginal cost of public funds for taxes on labor income ranges between \$1.09 to \$1.16 per dollar of tax revenue. This means that every dollar of government spending must be 9-16% more productive than private sector spending due to the costs of administering taxation. In the context of my welfare analysis, the marginal cost of public funds will shift the marginal cost curve for subsidy spending upwards, reducing the optimal level of per-modification subsidy. For the 2006 and 2007 origination cohorts, where marginal welfare benefits and marginal costs of subsidies are of a similar magnitude, accounting for the marginal cost of public funds may push the optimal subsidy towards zero because the cost of public spending may ultimately exceed the welfare benefit for market participants.

It is challenging to speculate on the change in optimal subsidy policy in a model that considers all potential market participants, negative externalities and the marginal cost of public funds for taxes. Some of these effects will drive up the marginal benefits of subsidies, while others will drive up the marginal costs of public spending. Ex-ante it is unclear which effect will dominate.

8 Conclusion

I present a new model to quantify the role of information asymmetry, transaction costs and relief sufficiency in foreclosure prevention efforts. By reducing an inherently dynamic interaction between borrowers and their loan servicers into a three-period game of commitment, I am able to reduce the computational burden involved with a multi-period, dynamic game with multiple players. I can use this empirical model to study how the distribution of the unobservable financial health of borrowers explains observed loan modification and foreclosure outcomes in my sample of Fannie Mae loans from their origination to the end of 2019. My results offer a comprehensive insight into how borrowers' financial wellbeing and the behavior of financial institutions influence debt relief policy.

I contribute several empirical findings on the role of information in foreclosure prevention. The central result is that financial institutions need to be highly informed about borrower default probabilities in order to offer loan modifications; greater information asymmetry between borrowers and their servicers leads to lower rates of modification and more foreclosures. I find that Fannie Mae loan servicers tended to be highly informed about borrower outcome probabilities in the wake of the 2008 financial crisis and that subsidies offered to banks under the program were more effective at preventing foreclosures in loans originated earlier in the 2000s, even though banks tended to be equally well informed about borrower financial health in all sample cohorts. In my welfare calculation, I find that Federal Government subsidies may have been well above the social welfare maximizing level but that the subsidies did not not have a dramatic effect on aggregate social welfare.

My work also contributes valuable findings about mortgage borrower behavior. I find that borrower foreclosure costs are significant and larger than previously estimated in past research. My estimates suggest that borrowers are more responsive to changes in their debt burden than changes in the value of their home. I also also find that there is meaningful variation in borrower financial health which is not explained by borrower or loan characteristics. This supports the more recent mortgage literature that assigns the vast majority of mortgage defaults to adverse life events (Ganong and Noel [2020b], Low [2022]).

My findings open up several exciting new directions for academic inquiry. Debt relief programs typically face a trade-off between speedy delivery versus effective targeting. Faster delivery of aid often requires less screening of borrowers, which comes at the cost of poorly targeted relief that ends up being rewarded to borrowers who are less likely to need it. Better targeting of relief leads to better allocation of resources but also inherently requires more screening, leading to longer waits and potentially greater cost for organizations analyzing borrower credit-worthiness. Payroll Protection Program loans made out by the U.S. government in response to the COVID-19 pandemic tended to prioritize rapid disbursement, but have often been criticized for being poorly targeted across organizations (Bartik et al. [2020] and Autor et al. [2022]).[50] My study highlights and quantifies the importance of a lender's information about its borrowers. When policymakers rely on the voluntary participation of lenders to grant relief, then they must factor in the role of information for decision-making within these organizations. Less-informed organizations may not have a sufficient private incentive to grant loan modifications, leading to costlier and less effective incentive programs. Future projects can also look to better understand the source of servicer information and how best to improve the quantity of efficient modifications. The exact causes of modification costs are also left unexplored. There are many potential culprits for high costs of processing loan modifications which include servicer capacity constraints, poorly trained staff, and insufficient resources.

[50] Popular media also gave significant coverage to poor loan targeting, see for example: New York Times (2020), *How Bad Was Virus Aid Fraud? One Banker Was 'Frustrated With Humanity'*, `https://www.nytimes.com/2020/12/09/business/ppp-fraud-paycheck-protection-program.html`.

References

The role of securitization in mortgage renegotiation. *Journal of Financial Economics*, 102(3):559–578, 2011.

M. Adelino, K. Gerardi, and P. Willen. Why don't lenders renegotiate more home mortgages? Redefaults, self-cures and securitization. *Journal of Monetary Economics*, 60(7):835–853, 2013.

S. Agarwal, Y. Chang, and A. Yavas. Adverse selection in mortgage securitization. *Journal of Financial Economics*, 105(3):640–660, 2012.

S. Agarwal, G. Amromin, I. Ben-David, S. Chomsisengphet, T. Piskorski, and A. Seru. Policy intervention in debt renegotiation: Evidence from the Home Affordable Modification Program. *Journal of Political Economy*, 125(3):654–712, 2017.

E. Anenberg and E. Kung. Estimates of the size and source of price declines due to nearby foreclosures. *American Economic Review*, 104(8):2527–51, August 2014.

D. Autor, D. Cho, L. D. Crane, M. Goldar, B. Lutz, J. K. Montes, W. B. Peterman, D. D. Ratner, D. Villar Vallenas, and A. Yildirmaz. The $800 billion paycheck protection program: Where did the money go and why did it go there? Working Paper 29669, National Bureau of Economic Research, January 2022.

P. Bajari, C. S. Chu, and M. Park. An empirical model of subprime mortgage default from 2000 to 2007. Working Paper 14625, National Bureau of Economic Research, December 2008.

P. Bajari, C. S. Chu, D. Nekipelov, and M. Park. A dynamic model of subprime mortgage default: Estimation and policy implications. Working Paper 18850, National Bureau of Economic Research, February 2013.

Barclays. 2008 Global Securitisation Annual 2. Technical report, January 2008.

A. W. Bartik, Z. B. Cullen, E. L. Glaeser, M. Luca, C. T. Stanton, and A. Sunderam. The targeting and impact of paycheck protection program loans to small businesses. Working Paper 27623, National Bureau of Economic Research, July 2020.

C. Bhat. Quasi-random maximum simulated likelihood estimation of the mixed multinomial logit model. *Transportation Research Part B: Methodological*, 35(7):677–693, 2001.

N. Bhutta, J. Dokko, and H. Shan. Consumer ruthlessness and mortgage default during the 2007 to 2009 housing bust. *The Journal of Finance*, 72(6):2433–2466, 2017.

E. K. Browning. The marginal cost of public funds. *Journal of Political Economy*, 84(2):283–298, 1976.

J. Y. Campbell, S. Giglio, and P. Pathak. Forced sales and house prices. *American Economic Review*, 101 (5):2108–31, August 2011.

J. H. Cardon and I. Hendel. Asymmetric information in health insurance: Evidence from the national medical expenditure survey. *The Rand journal of economics*, 32 3:408–27, 2001.

Congressional Oversight Panel. Foreclosure crisis: Working towards a solution. *March Oversight Report*, 2009.

CoreLogic. United States residential foreclosure crisis: Ten years later March 2017. Technical report, 2017.

Y. Demyanyk and O. Van Hemert. Understanding the Subprime Mortgage Crisis. *The Review of Financial Studies*, 24(6):1848–1880, May 2009.

R. Diamond, A. Guren, and R. Tan. The effect of foreclosures on homeowners, tenants, and landlords. Working Paper 27358, National Bureau of Economic Research, June 2020.

L. Einav, A. Finkelstein, and P. Schrimpf. The welfare cost of asymmetric information: Evidence from the U.K. annuity market. Working Paper 13228, National Bureau of Economic Research, July 2007.

L. Einav, A. Finkelstein, S. P. Ryan, P. Schrimpf, and M. R. Cullen. Selection on moral hazard in health insurance. *American Economic Review*, 103(1):178–219, February 2013.

H. Fang, Y. S. Kim, and W. Li. The dynamics of subprime adjustable-rate mortgage default: A structural estimation. *Working Paper*, January 2016.

K. A. Fowler, R. M. Gladden, K. J. Vagi, J. Barnes, and L. Frazier. Increase in suicides associated with home eviction and foreclosure during the us housing crisis: findings from 16 national violent death reporting system states, 2005-2010. *American Journal of Public Health*, 105(2):311–316, February 2015.

P. Ganong and P. Noel. Liquidity versus wealth in household debt obligations: Evidence from housing policy in the Great Recession. *American Economic Review*, 110(10):3100–3138, October 2020a.

P. Ganong and P. J. Noel. Why do borrowers default on mortgages? A new method for causal attribution. Working Paper 27585, National Bureau of Economic Research, July 2020b.

K. Gerardi, K. F. Herkenhoff, L. E. Ohanian, and P. S. Willen. Can't Pay or Won't Pay? Unemployment, Negative Equity, and Strategic Default. *The Review of Financial Studies*, 31(3):1098–1131, March 2018.

A. C. Ghent and M. Kudlyak. Recourse and Residential Mortgage Default: Evidence from US States. *The Review of Financial Studies*, 24(9):3139–3186, June 2011.

L. Guiso, P. Sapienza, and L. Zingales. The determinants of attitudes toward strategic default on mortgages. *The Journal of Finance*, 68(4):1473–1515, 2013.

E. Hembre. HAMP, home attachment, and mortgage default. *Working Paper*, February 2018.

N. Hendren. Private information and insurance rejections. *Econometrica*, 81(5):1713–1762, 2013.

N. Hendren. Knowledge of future job loss and implications for unemployment insurance. *American Economic Review*, 107(7):1778–1823, July 2017.

D. Herbst and N. Hendren. Opportunity unraveled: Private information and the missing markets for financing human capital. Working Paper 29214, National Bureau of Economic Research, September 2021.

J. Hu and R. Cantor. Measuring loss severity rates of defaulted residential mortgage-backed securities: A methodology. Special Comment 86769, Moody's Investors Servicers, April 2004.

S. Korgaonkar. The limited benefits of mortgage renegotiation. *Working Paper*, November 2020.

S. Kruger. The effect of mortgage securitization on foreclosure and modification. *Journal of Financial Economics*, 129(3):586–607, 2018.

S. Kruger and G. Maturana. Collateral misreporting in the residential mortgage-backed security market. *Management Science*, 67(3):2729–2750, 2019.

L. Kytömaa. Understanding cross-state differences in loan modification policy. Working paper, August 2022.

D. Low. What triggers mortgage default? New evidence from linked administrative and survey data. Working paper, January 2022.

V. R. Marone and A. Sabety. When should there be vertical choice in health insurance markets? *American Economic Review*, 112(1):304–42, January 2022.

G. Maturana. When Are Modifications of Securitized Loans Beneficial to Investors? *The Review of Financial Studies*, 30(11):3824–3857, May 2017.

C. Mayer, E. Morrison, T. Piskorski, and A. Gupta. Mortgage modification and strategic behavior: Evidence from a legal settlement with countrywide. *American Economic Review*, 104(9):2830–57, September 2014.

C. B. Mulligan. Foreclosures, enforcement, and collections under the federal mortgage modification guidelines. Working Paper 15777, National Bureau of Economic Research, February 2010.

S. Nelson. Private information and price regulation in the US credit card market. Working paper, April 2022.

T. Piskorski, A. Seru, and V. Vig. Securitization and distressed loan renegotiation: Evidence from the subprime mortgage crisis. *Journal of Financial Economics*, 97(3):369–397, 2010. The 2007-8 financial crisis: Lessons from corporate finance.

K. Train. *Discrete Choice Methods with Simulation*. Cambridge University Press, 2009.

A. C. Tsai. Home foreclosure, health, and mental health: a systematic review of individual, aggregate, and contextual associations. *PLoS One*, 10(4), April 2015.

U.S. Department of Housing and Urban Development. Economic Impact Analysis of the FHA Refinance Program for Borrowers in Negative Equity Positions. 2010.

A. M. White. Deleveraging the American homeowner: The failure of 2008 voluntary mortgage contract modifications. *Connecticut Law Review*, 41(4):1107–1131, 2009.

Y. Xin. Asymmetric information, reputation, and welfare in online credit markets. *Working Paper*, August 2020.

A Net present value calculation for a mortgage loan

A.1 Illustrating net present value calculation

- Take a fully-amortizing, fixed-rate loan with $1,000 principal and 1% monthly interest rate.
- With fixed interest rate and principal, longer term lengths always yield the same t_0 (ignoring default risk).
- The t_0 value of the loan starts to fall if a modification reduces the interest rate or grants principal forgiveness.

Table 11: 1-month repayment

	$t = 1$	$t = 2$
Monthly payment	$1,010.00	-
Discount factor	1.01	-
Period discounted value	$1,000.00	-
t_0 value		*$1,000.00*

Table 12: 2-month repayment

	$t = 1$	$t = 2$
Monthly payment	$507.51	$507.51
Discount factor	1.01	1.02
Period discounted value	$502.49	$497.51
t_0 value		*$1,000.00*

Table 13: 1-month repayment, monthly interest reduced to 0.50%

	$t = 1$
Monthly payment	$1,005.00
Discount factor	1.01
t_0 value	*$995.05*
% of original princ	99.5%

Table 14: 12-month repayment, monthly interest reduced to 0.50%

Monthly payment	$86.07
Summed t_0 value	*$968.70*
% of original princ	96.7%

- If we assume a late fee then the t_0 value could exceed the original principal value. For example:

	$t = 1$	$t = 2$
Monthly payment due	$507.51	$507.51
Monthly payment made	$0	$507.51*(1.02) + $507.51
Discount factor	1.01	1.02
Period discounted value	$0	$1,004.97
t_0 value		*$1,004.97*
% of original princ		*100.5%*

A.2 Net present value examples from Fannie Mae data

A non-performing loan with capitalized principal of \$104,290 at three months delinquency and a 6.25% annual interest rate receives a modification and then repays the loan in the following way:

- Pays for 60 months with 2% annual interest
- Pays for 12 months with 3% annual interest
- Pays for 25 months with 4% annual interest
- Ends loan with prepayment

The table below presents the contractually expected total payments for this modified loan and the realized total payments in the data. Expected payments and realized payments are within \$50 of each other.

Table 15: Expected payments

Time window	Expected total payments	Total discounted value
60 months payments @ 2%	\$20,457	\$17,553
12 months payments @ 3%	\$4,359	\$3,054
25 months payments @ 4%	\$10,365	\$6,601
Principal prepayment	\$91,271	\$54,290
Sum	\$126,452	\$81,498
% of princ at first 3-mo dlq	*121%*	*78%*

Table 16: Realized payments

Time window	Realized total payments	Total discounted value
60 months payments @ 2%	\$23,025	\$19,634
12 months payments @ 3%	\$4,565	\$3,197
25 months payments @ 4%	\$10,694	\$6,810
Principal prepayment	\$88,124	\$52,418
Sum	\$126,408	\$82,059
% of princ at first 3-mo dlq	*121%*	*79%*

B Additional detail on loan modification eligibility

TYPE OF HARDSHIP (CHECK ALL THAT APPLY)	REQUIRED HARDSHIP DOCUMENTATION
☐ Unemployment	▪ Not required
☐ Reduction in income: a hardship that has caused a decrease in your income due to circumstances outside your control (e.g., elimination of overtime, reduction in regular working hours, a reduction in base pay)	▪ Not required
☐ Increase in housing-related expenses: a hardship that has caused an increase in your housing expenses due to circumstances outside your control (e.g., uninsured losses, increased property taxes, HOA special assessment)	▪ Not required
☐ Disaster (natural or man-made) impacting the property or borrower's place of employment	▪ Not required
☐ Long-term or permanent disability, or serious illness of a borrower/co-borrower or dependent family member	▪ Written statement from the borrower, or other documentation verifying disability or illness **Note:** Detailed medical information is not required, and information from a medical provider is not required
☐ Divorce or legal separation	▪ Final divorce decree or final separation agreement **OR** ▪ Recorded quitclaim deed
☐ Separation of borrowers unrelated by marriage, civil union, or similar domestic partnership under applicable law	▪ Recorded quitclaim deed **OR** ▪ Legally binding agreement evidencing that the non-occupying borrower or co-borrower has relinquished all rights to the property
☐ Death of borrower or death of either the primary or secondary wage earner	▪ Death certificate **OR** ▪ Obituary or newspaper article reporting the death
☐ Distant employment transfer/relocation	▪ **For active duty service members**: Permanent Change of Station (PCS) orders or letter showing transfer. ▪ **For employment transfers/new employment**: Copy of signed offer letter or notice from employer showing transfer to a new location or written explanation if employer documentation not applicable, **AND** ▪ Documentation that reflects the amount of any relocation assistance provided (not required for those with PCS orders)
☐ Other – hardship that is not covered above: ____________________	▪ Written explanation describing the details of the hardship and any relevant documentation

Figure 12: Fannie Mae/Freddie Mac Form 710

C 3-digit Zip codes and Counties

Fannie Mae's Single-Family loan performance dataset restricts geographic data to a 3-digit ZIP code and an MSA code for each individual loan. The figures above depict California's 3-digit ZIP code areas (left) and counties (right).

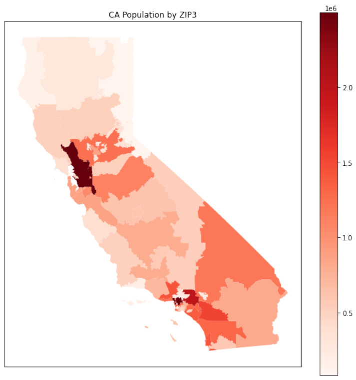

(a) California 3-Digit Zip Codes

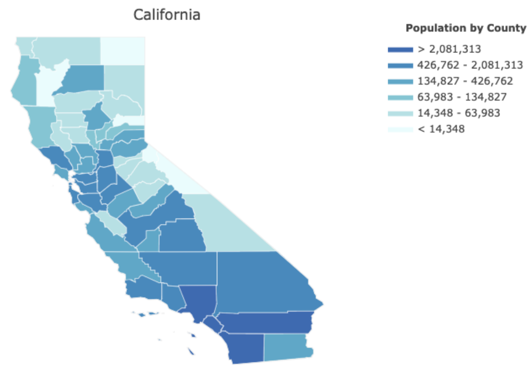

(b) California Counties

Figure 13: California 3-Digit Zip codes and Counties Counties[51]

D Servicer modification probability probit regressions

Figure 14: Predicted probabilities of delinquency, modification and default, by loan servicer in Q2 2004

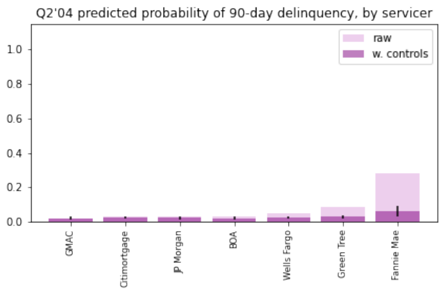

(a) Delinquency

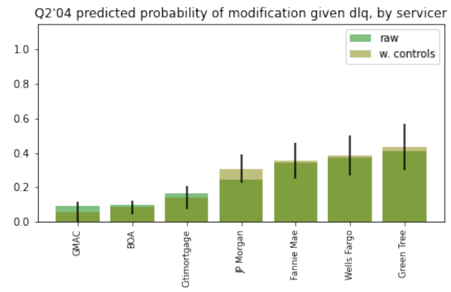

(b) Modification

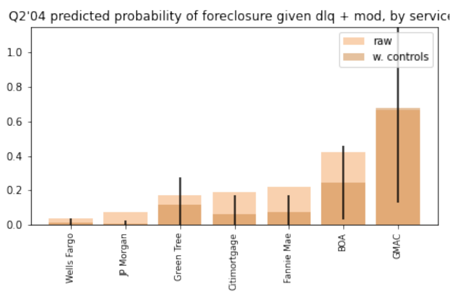

(c) Default with modification

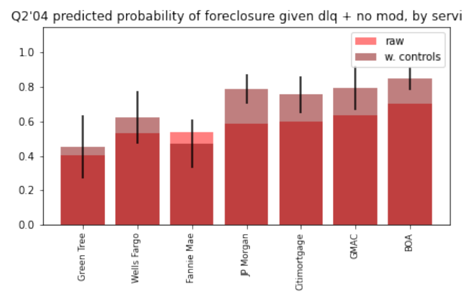

(d) Default without modification

Figure 15: Predicted probabilities of delinquency, modification and default, by loan servicer in Q2 2005

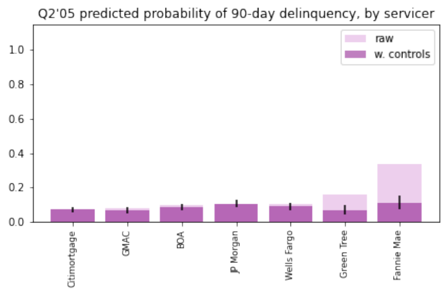

(a) Delinquency

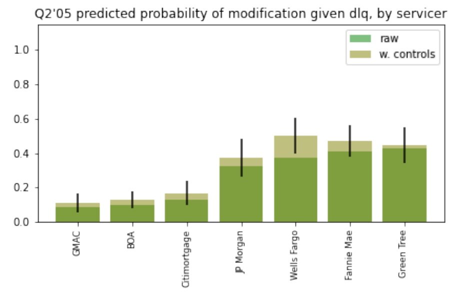

(b) Modification

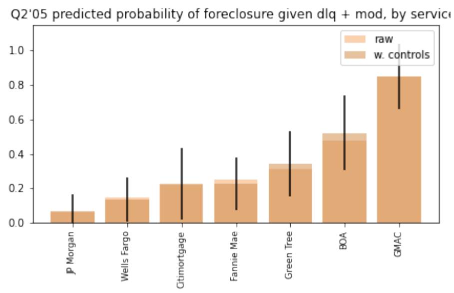

(c) Default with modification

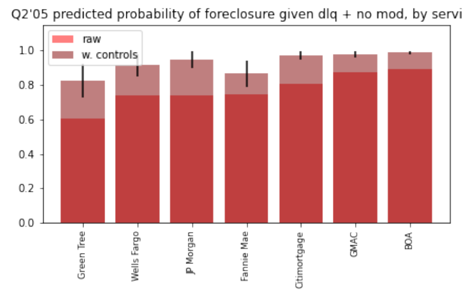

(d) Default without modification

Figure 16: Predicted probabilities of delinquency, modification and default, by loan servicer in Q2 2006

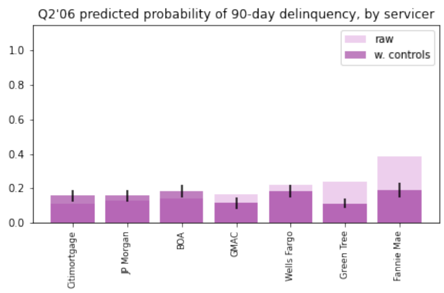

(a) Delinquency

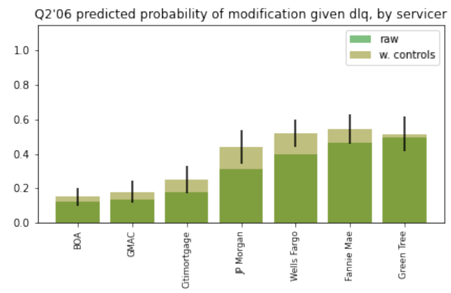

(b) Modification

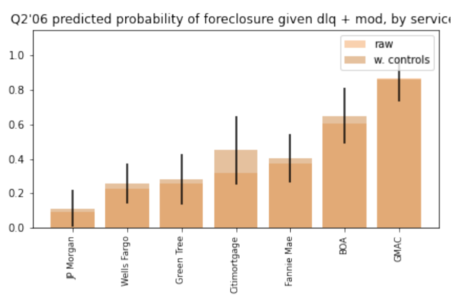

(c) Default with modification

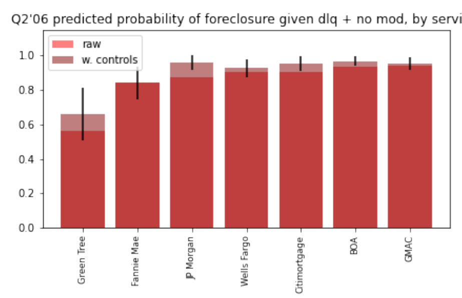

(d) Default without modification

Figure 17: Predicted probabilities of delinquency, modification and default, by loan servicer in Q2 2007

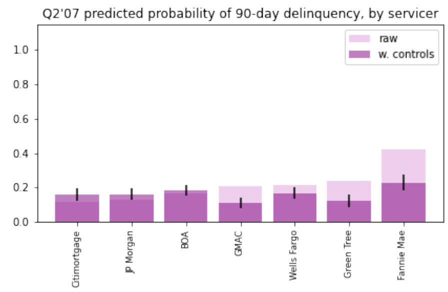

(a) Delinquency

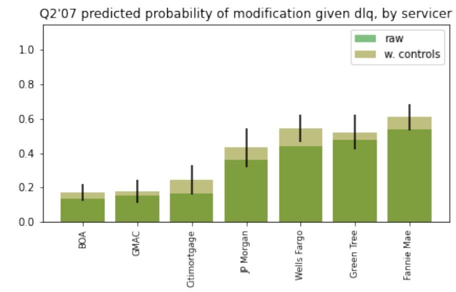

(b) Modification

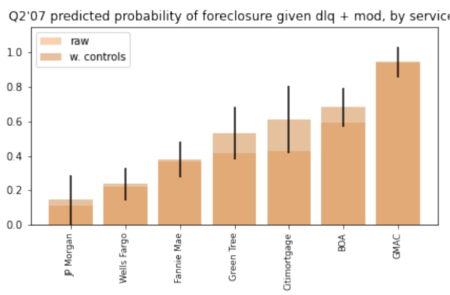

(c) Default with modification

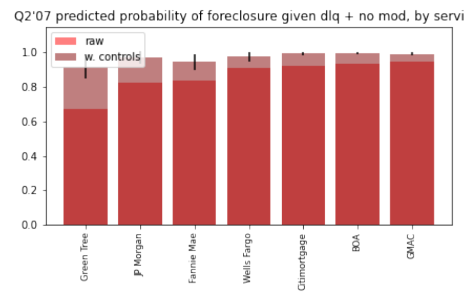

(d) Default without modification

E Construction of the equilibrium grid

This section provides additional detail on the construction of the equilibrium loan modification grid discussed in sub-section 5.3. The grid solves the servicer's cashflow maximization problem, and in doing so, also solves for equilibrium in the loan modification game for a large number of potential borrower characteristics, repayment expectations, unobservable variable distributions, and servicer modification costs. Solving for modification probabilities set by servicers, $m_i \in [0, 1]$, is central to estimating the model because these probabilities are not directly observable in the data. The probabilities rely on the structural model because they will depend on unobservables and because modification outcomes are censored for any borrower that avoids delinquency.

The simulated maximum likelihood approach used in estimation can become numerically demanding to implement because one must solve for model equilibrium for every observation in the dataset. Even if borrowers are grouped into discrete types based on their observable characteristics, one must still compute potentially many equilibrium probabilities across types. Within the maximum likelihood routine, one must then also re-solve the model for every borrower (or borrower type) for every potential guess of model parameters. To mitigate the computational burden involved with solving for model equilibria, I pre-solve them using a six-dimensional grid that can be loaded into the maximum likelihood routine. The six-dimensions are chosen to reduce the complexity of the servicer problem, which in practice can be very high dimensional. Without some form of simplification, the grid quickly suffers from a Curse of Dimensionality.

Dimensionality reduction takes place on borrower home utility, delinquency costs and expected repayments. The linearity assumption in the borrower's indirect utility function allows for a wide set of borrower observable characteristics to be reduced to the scalar value $x_i'\beta$. A benefit of this assumption is that additional observable characteristics and parameters can be added without affecting the dimensionality of the servicer's maximization problem. The ξ_i term is assumed to be known to the servicer and additively separable in the indirect utility so servicers set policy based on one dimension given by $x_i'\beta + \xi_i$. A similar linearity assumption in the borrower's delinquency cost also means that many observable characteristics collapse into the scalar $w_i'\lambda$ in the μ parameter of the delinquency cost. On repayment expectations, my model assumes that $\tilde{p}_i$ and f_i are both functions of p_i and rescaling factors ρ_m and ρ_f. Explicitly that: $\tilde{p}_i = \rho_m \cdot p_i$ and $f_i = \rho_f \cdot p_i$. Given the assumptions above, the six dimensions of the grid are:

1. $x_i'\beta + \xi_i$ - Home utility parameter

2. σ_ε - Unobserved standard deviation of home utility

3. $w_i'\lambda$ - Mean of delinquency cost

4. σ_η - Unobserved standard deviation of delinquency cost

5. p_i - Net present value of future stream of payments

6. c_j - Lender cost of modification

The grid is constructed by maximizing the servicer's objective function on several points of support in each dimension. My estimation routine employs a grid with 30 points of support of $x_i'\beta + \xi_i$, 26 points in the supports of p_i and c_j, 17 points in the support of σ_ε, and 15 points in the supports of $w_i'\lambda$ and σ_η. Jointly these supports form a grid with $30 \times 26^2 \times 17 \times 15^2 = 77,571,000$ points. For any set of inputs, I use 1,000 simulation draws to integrate out the borrower private information ε_i within the servicer's cashflow maximization problem, while private information about the delinquency cost, η_i, is integrated out analytically. The grid placement and dimensions were carefully chosen through extensive exploration and experimentation. $x_i'\beta + \xi$, p_i and c_j dimensions have the greatest number of points in the support because these values will vary across borrowers for any guess of model parameters. The unobserved distribution parameters, namely σ_ε and σ_η, will apply to all borrowers in the same way because the servicer only observes the distribution parameters rather than the specific draws from that distribution. In theory, it would be possible to select only one point for these two dimensions if the true values were known. Nevertheless, counterfactual analysis in which σ_ε is shifted will require more points to evaluate new equilibrium. The $w_i'\lambda$ dimension does vary with borrower characteristics but plays a less critical role for model fit than the dimensions associated with the borrower home utility, H_i.

Parallel computing dramatically speeds up solving for model equilibrium at the nearly 80 million points on the grid. This parallel computing task is "embarrassingly parallel" in the sense that the solution for a given set of inputs can be computed completely independently from some other set of inputs. By leveraging the Texas Advanced Computing Center's Stampede2 supercomputer, I am able to use 78 compute nodes simultaneously to achieve an almost 75× improvement in computation speed of the full grid relative to using a single node.[52] The number of hardware threads (workers) available per node also plays a crucial role in speeding up this computation. By using Python's *multiprocessing* library I am able to take advantage of either 96 hardware threads on Stampede2's SKX Compute Nodes, or 160 hardware threads on the ICX Compute Nodes. Using these resources, I am able to form the grid in around 20 minutes.[53]

F Further estimation detail

This section provides further detail on the simulated maximum likelihood approach discussed in section 5 of the paper. The full estimation routine incorporates the servicer's cashflow maximization problem (previous Appendix section), simulates borrower unobservable draws and searches across potential parameter values to maximize the log-likelihood function. The log-likelihood function is also conditional on observed outcomes in the data. The estimation was done using Python. The complete process can be summarized in the following four steps:

1. **Load in data and the pre-solved servicer's equilibrium grid:** After data cleaning and forming the equilibrium grid, both must be loaded into the maximization routine. Since the grid has been computed in parts to facilitate parallel computing, it must be rejoined into a single Python object.

2. **Take S simulation draws for ξ_i and ε_i:** I take S standard normal draws for ξ_i and S standard uniform draws for ε_i. For any given set of model parameters, I then scale these draws up or down. To illustrate this: with $\xi_i \sim N(\mu_\xi, \sigma_\xi^2)$ the simulation draws $s_\xi \sim N(0,1)$ are re-scaled as $\hat{s}_\xi = \mu_\xi + s_\xi \sigma_\xi$.[54] I take standard uniform draws for ε_i instead of standard normal draws because the servicer observes the ξ_i when making loan modification decisions and must then take truncated draws of ε_i conditional on a borrower's type and the observed ξ_i draw. This procedure for taking truncated univariate densitites follows section 9.2.4 of Train [2009]. Halton draws or another form of systematic sampling could also be used to take draws from both the standard normal and the standard uniform distributions. (Train

[52] Interested readers can learn more about Stampede2 here: `https://portal.tacc.utexas.edu/user-guides/stampede2`
[53] It took about 65 hours to form an early iteration of this grid using a single node with eight threads.
[54] I set $\mu_\xi = 0$ and $\mu_\varepsilon = 0$.

[2009]).[55] In practice, I use $S = 110$.

3. **Calibrate the lender's loss on the $\tilde{p}_i$ and f_i as a share of p_i:** To simplify the computational burden, my model assumes that $\tilde{p}_i$ and f_i are both functions of p_i and rescaling factors ρ_m and ρ_f. Explicitly that $\tilde{p}_i = \rho_m \cdot p_i$ and $f_i = \rho_f \cdot p_i$. My estimation routine takes the appropriate re-scale factors ρ_m, ρ_f as arguments. I test my model with various different re-scale factors as a robustness test.

4. **Maximize the likelihood function using global search algorithm followed by a local search:**

   My log-likelihood function is not globally concave and features flat sections that create challenges for computational maximization routines. To increase the probability of finding the parameters that reach the global maximum of my log-likelihood function, I conduct a global search algorithm that emphasizes shifting away from potential local extrema. At the completion of the global parameter search I allow for a more extensive non-linear local maximization of the function based on the maximizing value found in the global search.

   (a) *Global search:* I set the global maximizing routine to take 500 random "steps" or iterations in the initial parameter guesses before conducting 200 iterations of a local Nelder-Mead, non-linear parameter search. The Nelder-Mead iterations are intentionally constrained in this step to limit estimation run-time. This procedure is conducted using Python's *basinhopping* function within the *scipy.optimize* library.[56] Basin-Hopping is a more general version of Simulated Annealing.

   (b) *Local search:* The local search takes the results of the global search as the initial guess for parameters and then again conducts a non-linear parameter search using the Nelder-Mead simplex method. This search is allowed to continue until convergence. This procedure is conducted using Python's *minimize* function within the *scipy.optimize* library.[57]

[55] For example, Bhat [2001] shows that Halton draws can greatly improve simulation precision in mixed logit setting over conventional sampling.

[56] See: `https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.basinhopping.html`

[57] See: `https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.minimize.html`

G Estimated model parameters

Table 17: Borrower utility coefficients, unscaled

	Q2 2004 Sample		Q2 2005 Sample		Q2 2006 Sample		Q2 2007 Sample	
	Ests	Std. Err.	Ests	Std. Err.	Ests	Std. Err.	Ests	Std. Err.
β_0	11.185	0.0483	8.913	0.0644	-11.127	0.1573	12.880	0.0753
$\beta_{Home\ Value}$	0.516	0.0004	0.919	0.0010	0.879	0.0013	1.068	0.0009
β_{LTV}	-12.022	0.0752	-23.952	1.9664	-18.455	0.1577	-27.221	0.0806
β_{DTI}	-10.877	0.0975	-5.675	0.3226	-10.783	0.3206	0.429	0.0691
$\beta_{Credit\ Score}$	0.194	0.0003	0.150	0.0002	0.193	0.0003	0.134	0.0002
$\beta_{\Delta UE}$	-24.730	0.0488	-25.958	0.0373	-24.490	0.0953	-33.999	0.0358
$\beta_{Log\text{-}income}$	-0.270	0.0043	0.076	0.0023	0.188	0.0118	0.622	0.0104
σ_ξ	48.316	0.0331	54.704	0.0611	60.615	0.1022	57.282	0.0807
σ_ε	4.489	0.0957	3.639	0.1514	4.007	0.0624	3.891	0.2759
N	21,310		10,864		8,928		10,782	
Log LL	6,661		5,875		6,649		7,664	

Note:
This table presents the estimated parameters for the borrower utility function before they have been re-scaled back to a $ equivalent. Standard errors are reported in parentheses.

Table 18: Borrower delinquency coefficients, unscaled

	Q2 2004 Sample		Q2 2005 Sample		Q2 2006 Sample		Q2 2007 Sample	
	Ests	Std. Err.	Ests	Std. Err.	Ests	Std. Err.	Ests	Std. Err.
λ_0	-0.641	1.7564	-0.045	0.2586	-0.512	0.0378	-0.348	0.0112
λ_{HV}	0.002	0.0028	-0.001	0.0042	0.002	0.0003	0.004	0.0001
λ_{LTV}	0.004	0.0078	-0.0	0.0004	0.0	0.0533	0.0	0.0094
λ_{DTI}	0.005	0.0024	0.002	0.5946	0.0	0.1194	-0.0	0.0315
$\lambda_{Credit\ Score}$	0.0	0.0	0.0	0.0003	-0.0	0.0001	0.0	0.0001
λ_{UE}	0.066	7.9203	-0.006	0.1383	-0.001	0.0211	-0.005	0.0064
$\lambda_{Log\ income}$	-0.0	0.0	0.0	0.0017	0.0	0.0001	-0.0	0.0002
σ_η	1.486	0.4669	2.500	0.0810	1.284	0.1716	2.467	0.1165
N	21,310		10,864		8,928		10,782	
Log LL	6,661		5,875		6,649		7,664	

Note:
This table presents the estimated parameters for the borrower delinquency cost before they have been re-scaled back to a $ equivalent.

Table 19: Servicer cost coefficients, unscaled

	Q2 2004 Sample		Q2 2005 Sample		Q2 2006 Sample		Q2 2007 Sample	
	Ests	Std. Err.	Ests	Std. Err.	Ests	Std. Err.	Ests	Std. Err.
c_{BoA}	1.027	0.0468	0.999	0.0454	0.779	0.0590	0.443	0.0253
$c_{Wells\ Fargo}$	0.223	0.0133	0.222	0.0110	0.235	0.0060	0.240	0.0014
$c_{Citimortgage}$	1.063	0.0745	1.026	0.0537	0.776	0.0958	0.684	0.1223
c_{GMAC}	1.090	0.8941	2.058	0.1608	2.384	0.1698	3.250	0.4315
$c_{JP\ Morgan}$	0.619	0.1405	0.226	0.0092	2.160	0.2508	1.685	0.2431
$c_{Fannie/Seterus}$	0.404	0.1158	0.215	0.0065	0.198	0.0102	0.174	0.0120
$c_{Green\ Tree}$	0.191	0.0296	0.205	0.0154	0.171	0.0218	0.179	0.0190
N	21,310		10,864		8,928		10,782	
Log LL	6,661		5,875		6,649		7,664	

Note:
This table presents the estimated parameters for servicer modification costs before they have been re-scaled back to a $ equivalent. Standard errors are reported in parentheses.

H Distributions of servicer losses

Figure 18: Relative distributions of servicer losses from incorrect modification compared to gains from correct modifications, by origination cohort

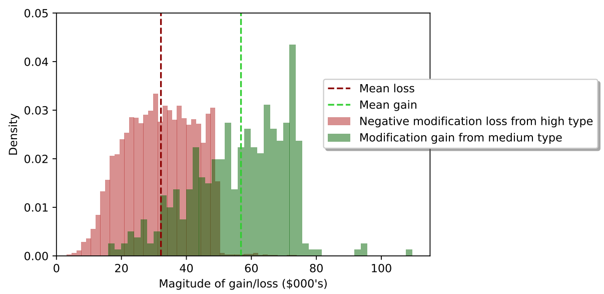

(a) Q2 2004

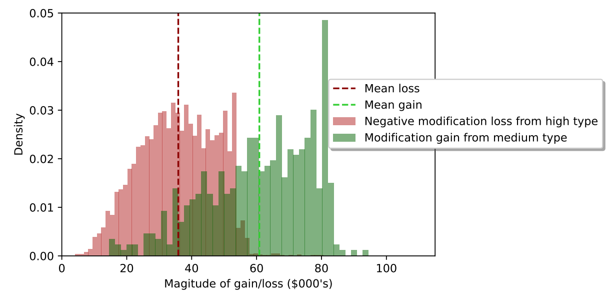

(b) Q2 2005

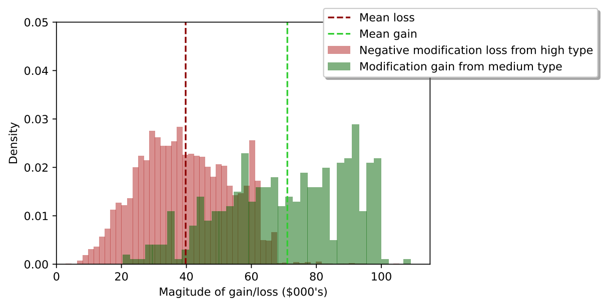

(c) Q2 2006

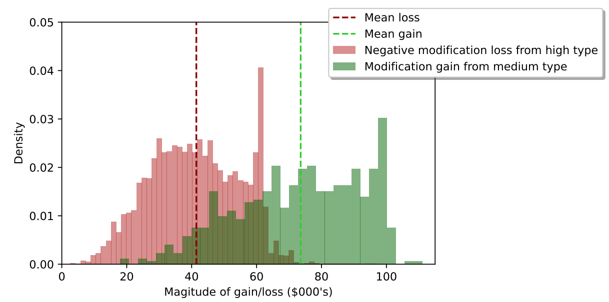

(d) Q2 2007

I Servicer simulated accuracy across origination samples

Figure 19: Servicer modification accuracy, by origination cohort

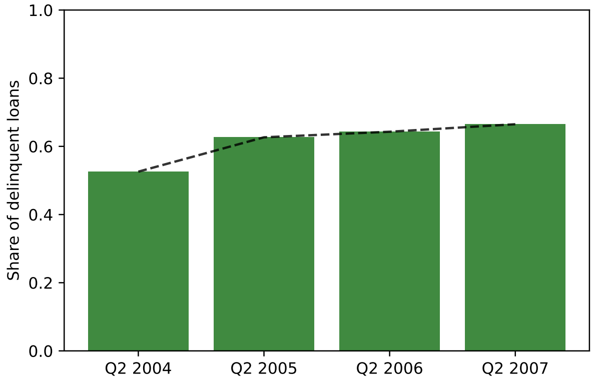

(a) Servicer correct: No Modification

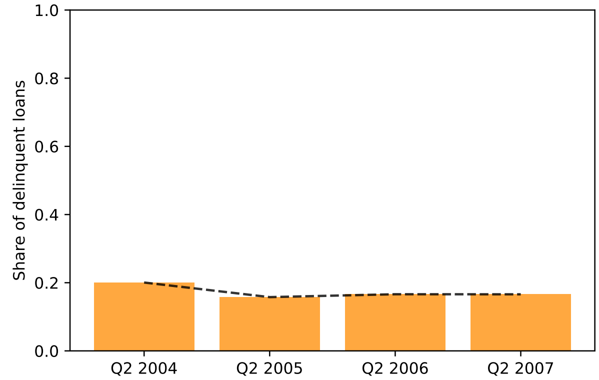

(b) Servicer Mistake: Modification

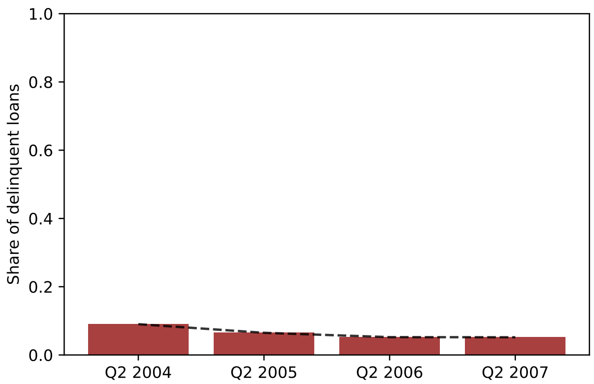

(c) Servicer mistake: No Modification

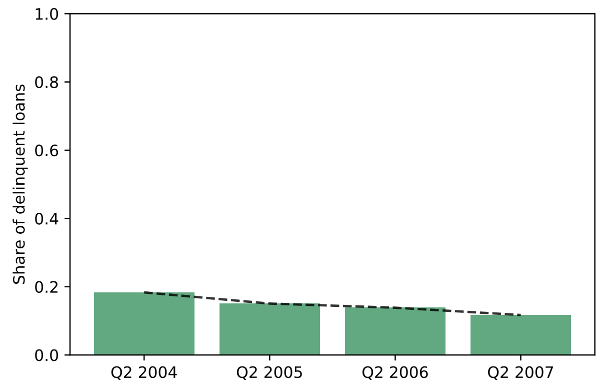

(d) Servicer correct: Modification

J Additional notes on information varying counterfactual

The focus of the information-shifting counterfactual is to study the impact of changing loan servicers' information about borrowers. To accomplish this, I vary the share $\frac{\sigma_\xi^2}{\sigma_\xi^2+\sigma_\varepsilon^2}$ explained by σ_ξ^2 to preserve the overall variance of H_i. Under my approach I shift this share in the following way:

- Take some share $s \in [0,1]$ and compute new σ_ξ^2 and σ_ε^2 as follows:

$$\hat{\sigma}_\xi^2 = s \cdot (\sigma_\xi^2 + \sigma_\varepsilon^2)$$
$$\hat{\sigma}_\varepsilon^2 = (1-s) \cdot (\sigma_\xi^2 + \sigma_\varepsilon^2)$$

- To take simulation draws as in Train [2009] to evaluate borrower integrals:

$$\hat{\xi}_i \;\, draws = \underbrace{\mu_\xi}_{=0} + \sqrt{\hat{\sigma}_\xi^2} \times d \qquad \underbrace{d \sim N(0,1)}_{Standard\ normal\ simulation\ draws}$$
$$\hat{\varepsilon}_i \;\, draws = \underbrace{\mu_\varepsilon}_{=0} + \sqrt{\hat{\sigma}_\varepsilon^2} \times d$$

It is easy to see that variance remains consistent with initial variance because s terms cancel out:

$$\hat{\sigma}_\xi^2 + \hat{\sigma}_\varepsilon^2 = (\sigma_\xi^2 + \sigma_\varepsilon^2) \quad \forall s$$

K Effects of modification subsidies and servicer information on equilibrium outcomes, all cohorts

Figure 20: Effect of modification subsidies and servicer information on delinquency, by origination cohort

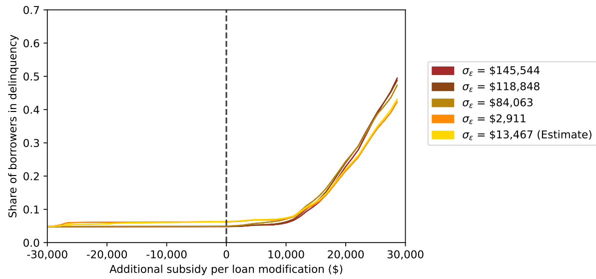

(a) Q2 2004

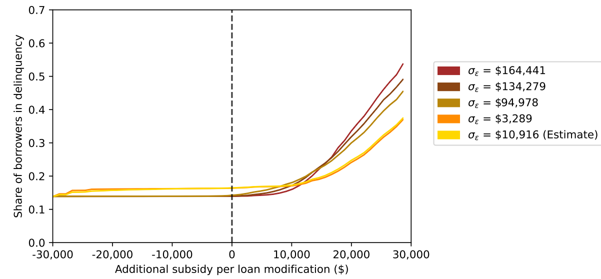

(b) Q2 2005

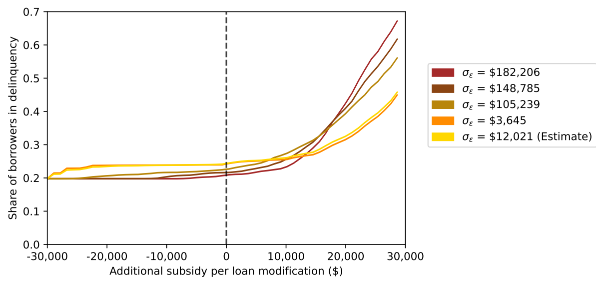

(c) Q2 2006

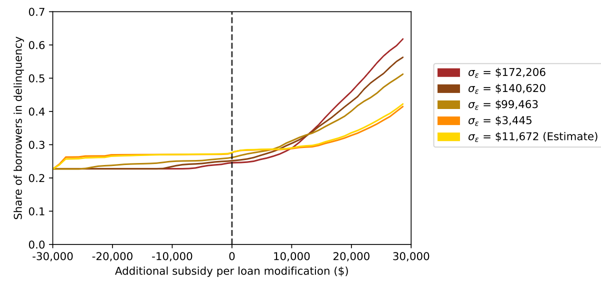

(d) Q2 2007

Figure 21: Effect of modification subsidies and servicer information on modifications, by origination cohort

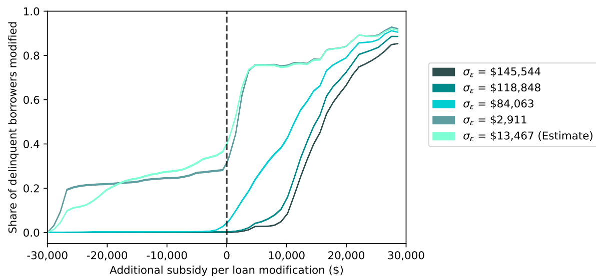

(a) Q2 2004

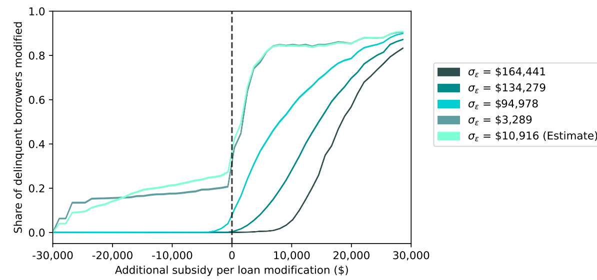

(b) Q2 2005

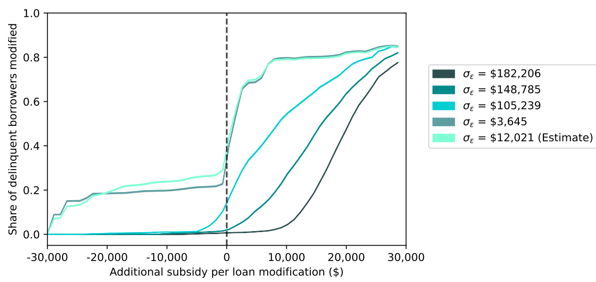

(c) Q2 2006

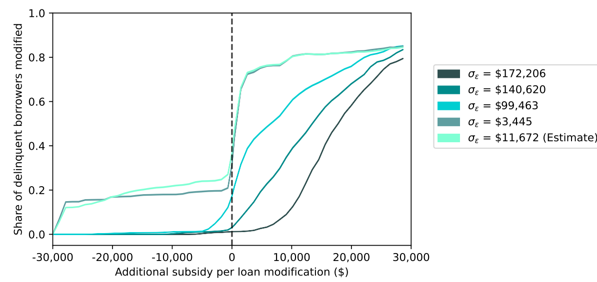

(d) Q2 2007

Figure 22: Effect of modification subsidies and servicer information on foreclosure, by origination cohort

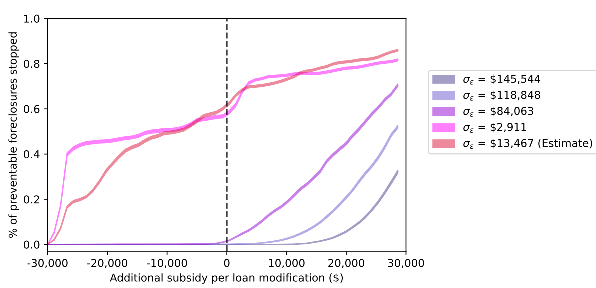

(a) Q2 2004

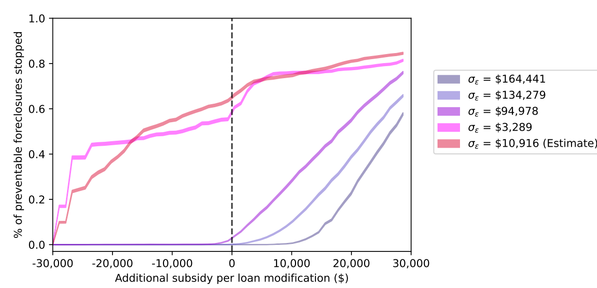

(b) Q2 2005

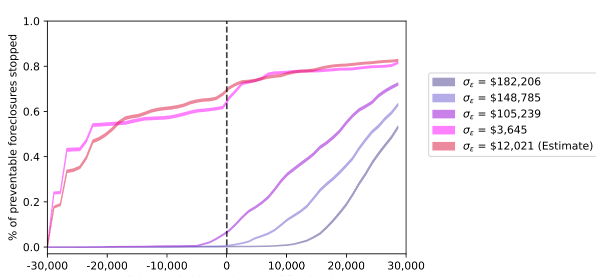

(c) Q2 2006

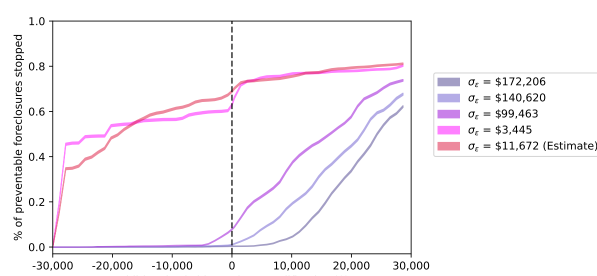

(d) Q2 2007