

Targeted Principal Forgiveness Is Effective: Mortgage Modification and Financial Crisis

Philip Lewis Kalikman^{*‡} Joelle Scally^{†‡}

2023 September 29

Abstract

Research into the Global Financial Crisis finds forgiving mortgage principal ineffective at stemming defaults; authors argue that borrowers default because of illiquidity, not strategically. We argue the opposite. Targeted forgiveness is effective. And default is better explained by quantifying how illiquidity and strategy interact. We embed their interactions in a computational heterogeneous structural model. We introduce idiosyncratic default penalties: differing penalties underpin borrowers' differing deviations from pure-financial optimality. We run the model with heterogeneous microdata, estimating penalties from credit scores and payment histories. Forgiving low-score, deep-underwater borrowers would have eliminated nearly all their defaults, with net gain for lenders. JEL Codes: G01, G21, R00.

*Corresponding author: University of Cambridge, 19 Silver Street, Cambridge CB3 9EL United Kingdom; philip@kalikman.com (preferred), or pk563@cam.ac.uk. I am grateful to my advisors John Geanakoplos and Andrew Metrick, and my unofficial advisor Andrew Haughwout. I thank Joshua Barratt, Archishman Chakraborty, Tess Scharlemann, and Walter Torous for thorough discussion. For helpful comments, I thank Andrew Ellul, Ravi Jagadeesan, Wilbert van der Klaauw, Donghoon Lee, Lira Mota Mertens, Nagpurnanand Prabhala, Gary Pivo, Albert Saiz, Matt Sekerke, Tray Spilke, Gokhan Torna, and seminar participants at the American Real Estate and Urban Economics Association, Cambridge University, the Federal Reserve Bank of New York, the Financial Management Association, Institut Minés-Telecom, Massachusetts Institute of Technology, the Office of Financial Research, Oxford University, Stony Brook University, Texas A&M University, the University of Arizona, the University of Hawaii, Yale University, the Yale School of Management, the Yale SOM Financial Crisis Conference, and Yeshiva University. I thank Frank Innocenti, Kevin Kelliher, Dmitry Krivitsky, Ernest Miller, and Na'im Tyson for assistance with the FRBNY computing cluster. I thank the FRBNY for generously providing the data, computing resources, and espresso necessary to complete this project.

[†]Federal Reserve Bank of New York.

[‡]The views expressed herein are those of the authors, and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

Defaults on residential mortgages caused the 2007-9 Global Financial Crisis, erasing \$9 trillion in housing value and \$10 trillion in other asset values. Borrowers all over the country, with widely varying mortgage terms, financial situations, and regional house price dynamics, defaulted in record numbers.

Economists studying the Crisis and its aftermath are converging on consensus regarding two longstanding questions in mortgage finance: Why do borrowers default? And when defaults threaten the entire economy, how should policymakers intervene? The emerging consensus argues that borrowers default because of external shocks that reduce borrowers' ability to afford payments—the *double-trigger* model of default—and that, accordingly, the best intervention modifies the borrower's mortgage to keep the borrower owing the full amount but paying it off in lower payments over a longer time. Losing favor are the contrary notions that borrowers default based on a rational calculation of what's in their financial best interest—the *strategic* model of default—and accordingly, that forgiving principal, i.e. simply reducing or “writing down” the amount borrowers owe lenders, would efficiently ameliorate defaults.¹

We argue for conclusions that challenge the emerging consensus. Principal forgiveness would have cost-effectively reduced defaults and halted the Crisis. It should remain in policymakers' toolkits for use whenever price drops leave large numbers of borrowers with negative equity. Forgiveness would have worked because borrowers *are* strategic about their mortgages: for most people, mortgages are the largest and most consequential financial instruments with which they will ever interact. Their strategic thinking interacts with their level of liquidity, interactions neither the double-trigger model nor the strategic model adequately captures.

We capture these interactions in a computational structural model. Our model extends the existing state of the art in mortgage default modeling with innovations necessary to tease out how forgiveness can be an effective policy, as well as why studies of past principal forgiveness find it ineffective. In particular, in what we believe is a first in the literature, we introduce an idiosyncratic non-pecuniary penalty for default, and we estimate how this unob-

¹While the terms “double-trigger” default and “strategic” default, sometimes referred to as “can't-pay” default and “won't-pay” or “ruthless” default, are standard in the literature, they have acquired moral baggage that we wish to avoid. When we write that a borrower or her behavior is “strategic,” or “more strategic,” we mean only that she or it more closely resembles the prediction of a rational economic model. The “double-trigger” model—which refers to two circumstances “triggering” default, namely, negative equity and illiquidity—is sometimes also referred to as the liquidity-constrained model; some authors distinguish the terms to focus on positive-equity defaults, which are not the focus of our study.

servable characteristic relates to borrowers' observable credit scores. The penalty governs the strength of illiquidity and strategy interacting in determining the borrower's default decision: borrowers with higher penalties endure further deviations from the pure-strategic financial optimum before they default. Our model quantifies the size of this deviation borrower-by-borrower, and thus issues realistic predictions of behavior borrower-by-borrower. We can therefore derive a defensible answer to a counterfactual question regarding policymakers' responses to the Crisis: what if policymakers had implemented a targeted forgiveness program?

Principal forgiveness was rarely employed in the Crisis. Policymakers worried about offering aid to undeserving borrowers. When policymakers did implement forgiveness, it often targeted borrowers unlikely to respond to it and forgave so little of their debt as to leave them still with negative equity—i.e., with no reason to shift their strategic calculus over default. It is not surprising that an inadequate dose of treatment issued to an unreceptive population was ineffective. But it is not dispositive of the treatment's potential efficacy. We use our model to show that a simpler principal forgiveness program, targeting a larger dose of forgiveness to a broad group of borrowers, would have succeeded.

We select borrowers in the lower half of the distribution of credit score at origination and who live in ZIP codes where prices dropped enough to render their *scheduled* first liens—i.e., what they would have owed, regardless of whether they made payments on time—130% underwater.² We run the model against thousands of distinct loans, varying not only penalties but all economic variables, including borrower income and wealth, mortgage size, mortgage interest rate, loan-to-value ratio, and house price appreciation dynamics, avoiding representativeness assumptions that would muddy the measurement of the effectiveness of forgiveness.

Forgiving these borrowers' debt to 95% combined loan-to-value ratio would prevent 84% of their defaults. Additionally, the intervention would present no cost to taxpayers. Lenders writing down the balances owed by such borrowers would have earned on average approximately \$14,000 more than they would earn in the baseline, even taking into account forgiveness distributed to borrowers who would never have defaulted, as lenders would avoid losses they would have incurred selling borrowers' foreclosed homes. Borrowers in our study had “vanilla” mortgages—thirty-year, fixed-rate, fully-amortizing, non-government-insured

²These selection criteria are *ex ante* identifiable but are immune to the considerations of perverse incentives that drove policymakers to eschew forgiveness in the Crisis: no decision made by a borrower after such a program's announcement could affect her eligibility for it, as she cannot affect her original credit score by stopping payments in the present, and her individual default would not meaningfully affect prices in her ZIP code, let alone the history thereof.

mortgages—that were far safer than many that abounded in the crisis, implying that our figures are likely lower bounds on the number of averted defaults and the benefit to lenders implementing such a program.

Why is the current consensus developing against the effectiveness of principal forgiveness? In large part the reaction derives from the shortcomings of HAMP, the Treasury Department’s “Home Affordable Modification Program,” which was the largest mortgage modification program implemented to stem the tide of defaults in the Crisis. Starting in 2009, it spent approximately \$50 billion to incentivize mortgage servicers—the third-party businesses who communicate with and collect payments from borrowers, in exchange for a fee skimmed from the pass-through payments to the investors in the mortgage note—to offer modifications to mortgages of select borrowers.³ To be eligible for a HAMP modification, borrowers had to submit three-page affidavits affirming, under penalty of perjury, that they faced unexpected financial hardship such as divorce, illness, or job loss. The hardship requirement was designed to make sure aid only went to those borrowers who were behind on payments through “no fault of their own,” rather than to strategic defaulters who could afford to pay but chose not to. Borrowers who qualified received a series of algorithmically-determined changes to their mortgages, designed to reduce their monthly payments without reducing the amount they owed to lenders. At first, HAMP did not feature any principal forgiveness.

At the time, many commentators had called for forgiving principal. Because house prices had declined so drastically, many borrowers were “underwater” or in negative equity, i.e. they owed more on their mortgages than their homes were worth. Those borrowers faced a strong financial incentive to walk away from their mortgages and rent a different home for a lower monthly payment. When a borrower defaults, for any reason, lenders foreclose on the property and sell it to recover some of, but usually not all of, their lost principal. Borrowers who default nearly always already owe more principal than the home is worth—otherwise, they would sell the home to cover the owed principal and keep the positive difference—meaning lenders selling a home would rarely cover the cost of their outstanding principal even if selling the home were frictionless for the lender. On top of that shortfall, lenders face legal and administrative delays, neglect from tenants who know they are losing the property, and stigma associated with foreclosure that reduces the price a foreclosed home commands on the market. These forces typically lead the foreclosed home to sell at a discount to market

³The nature of the relationship among mortgage originators, servicers, and investors typically required that servicers, not original lenders or ultimate investors in the mortgage note, were the entities responsible for effecting modifications. This structure sometimes imposed barriers to modifications even when lenders wished to offer them, as servicers did not stand to gain from modifying mortgage terms, even if lenders did. For further discussion see Appendix D and McCoy (2013).

prices, leaving the lender well short of the full principal owed.

Policymakers in the crisis discovered quickly that foreclosure was not just costly to lenders but also had detrimental spillover effects on neighboring properties and regional jobs and housing, and, intermediated by fragile financial markets, the wider economy.⁴ As borrowers with positive equity can sell the home themselves if they cannot or do not want to keep making payments, writing down principal enough to restore borrowers to positive equity would have immediately stopped defaults. And it would have done so without the hassle imposed on lenders by foreclosure, and without foreclosure's negative externalities. Since no one knew when house prices would recover, some felt that principal forgiveness was a surer way to re-align borrowers' owed balances and payment sizes with the market in correction in house prices and halt the crisis in its path.

Some commentators suggested that writedowns would pay for themselves. Geanakoplos and Koniak (2008) and Geanakoplos and Koniak (2009) argued that deep underwater subprime borrowers were defaulting in such high numbers that lenders would earn more by forgiving debt to a level that incentivized borrowers to keep making payments than the lenders would recover by selling the defaulting borrowers' homes in foreclosure, and that the only need for government involvement was to coordinate the misalignment of incentives between the servicers who administered the loans and the lenders who would benefit but lacked the means to act directly.⁵ The intervention would be cost-free to taxpayers, would benefit both borrowers and lenders, and would achieve the social policy goal of significantly reducing defaults.

Other commentators called for even broader writedown policies. Zingales (2008) proposed writing down mortgage principal in proportion to the decline in home prices for *any* borrower who purchased in ZIP codes where prices fell by more than 20%, not just subprime borrowers, borrowers who defaulted, or borrowers in negative equity. Goetzmann and Koppell (2008) called for a no-questions-asked government refinance of any delinquent loan, which would cancel all past-due debt regardless of the reason the borrower had stopped making payments

⁴Whether default is itself an efficient outcome, or similarly whether mortgages should self-modify, has been studied by other authors. See Eberly and Krishnamurthy (2014), Shiller et al. (2017), Campbell, Clara, and Cocco (2020), Guren, Krishnamurthy, and Mcquade (2021), as well as Dubey, Geanakoplos, and Shubik (2005) and Zame (1993). Default may be socially optimal in some contexts. But we do not believe the level of default was optimal in the Global Financial Crisis, when spillovers into the entire economy destroyed jobs and livelihoods. Regardless, policymaker actions reveal that they did not believe coordinated defaults socially optimal in either the Global Financial Crisis or the Covid-19 crisis. Policymakers implementing a response to coordinate defaults need not take a view on whether *isolated* defaults are efficient. Furthermore, while self-modifying mortgages structures are interesting and potentially efficiency-enhancing, borrowers did not and do not have these mortgages.

⁵See again Appendix D.

or even local housing market conditions.

All of these interventions were proposed before HAMP was designed and implemented. But policymakers worried that forgiving debt was morally and politically fraught. What if forgiveness went to borrowers who it seemed could “afford” their payments, but stopped payments out of self-interested economic rationality, rather than because of some external hardship that was “not their fault”? What if a broad forgiveness program skewed incentives so that more borrowers would default, potentially exacerbating the crisis? In part for these reasons, a large-scale forgiveness program never got off the ground.⁶

But a sub-program of HAMP announced in 2010, years after the housing market entered its tumultuous decline, did test out a flavor of principal forgiveness: writing down a small portion of the debt of a small number of borrowers. This program, called HAMP Principal Reduction Alternative (HAMP PRA), provided for principal reductions that were in most cases so modest they were failed to offset the shortfalls between borrowers’ mortgage debt and their homes’ value. 85% of the participants of this program were left with first-lien loan-to-value ratios of 115% or higher, and many had second liens and so were even deeper underwater. Whether a borrower received a writedown was partially exogenous, providing economists an opportunity to study the effectiveness of principal forgiveness.

Such studies found principal forgiveness in PRA ineffective. Scharlemann and Shore (2016) found that principal reduction in PRA cost \$320,000 per foreclosure it averted, such a steep price tag that a policymaker would need an unrealistically large estimate of the social cost of foreclosure to justify it. Ganong and Noel (2020) found that PRA’s writedowns cost on average \$70,000 but with virtually zero impact on borrowers’ propensity to default versus modifications that achieved comparable payment reduction without forgiving principal.

These results confirmed that selecting borrowers who persisted in making payments for two years after the crisis, but stopped due to a documented liquidity shock—i.e., exactly those borrowers least likely to have ceased payments earlier due to strategic consideration of their financial benefit—and treating them with a dose of principal forgiveness too small to cure their negative equity, would not cost-effectively alter their behavior. Generalizing the lessons of principal forgiveness beyond the narrow sample and circumstances in HAMP PRA, let alone concluding that all principal forgiveness would be ineffective at any dose, would be a mistake.

In the onset of the 2020 Covid-19 crisis, Fannie Mae and Freddie Mac—the two government-sponsored enterprises which, by securitizing nearly all newly originated mortgages in the

⁶Nonetheless, by requiring borrowers to be behind on payments to qualify, HAMP itself did provide precisely the perverse incentive opponents of principal forgiveness argued would damage the economy.

United States, provide liquidity for, and set policy for, nearly the entire residential housing market—announced mortgage *forbearance* programs, i.e. pauses that delay payments but keep borrowers owing the same amount on their mortgages and paying over a longer time. A crisis where house prices climb while incomes face a temporary pause is tailor-made for a liquidity-responsive modification such as forbearance.

But what if a future crisis finds borrowers once again in deep negative equity? If commercial mortgage defaults lead to more bank failures and detrimental effects to the wider economy, should policymakers entertain forgiving mortgage debt? Is it ever sensible to bail out individual borrowers, or should we only bail out financial institutions?

We seek to answer these questions using a structural model. Our approach is motivated by the conflicting observations that the literature offers some evidence for the notion that principal forgiveness might have succeeded if it had been delivered differently, yet plenty of evidence that certain designs failed. The experiential record lacks enough natural variation to reconcile these contradictory findings. But the questions remain pressing, and will continue to bear on policymakers' responses to future crises.

When did principle forgiveness work? Haughwout, Okah, and Tracy (2016) showed that for securitized subprime loans that received private modifications before HAMP's implementation, reducing payments by forgiving principal reduced defaults nearly *fourfold* versus reducing payments without forgiving principal. These loans were not a representative sample, and are in many regards dissimilar from loans active today. Kaplan, Mitman, and Violante (2020) build an equilibrium model, calibrated to limited data, in which a counterfactual debt forgiveness program likewise succeeds in stemming defaults. Both studies suggest the possibility that principal forgiveness can be an effective tool for preventing default.

How should we reconcile the seemingly contradictory results that principal forgiveness performed poorly in HAMP PRA, yet performed well in a different context? It stands to reason that principal forgiveness would only be effective when it restores borrowers to positive equity. Might it have worked for HAMP borrowers if only it had been more generous? Or might it be that principal forgiveness is only effective on certain populations of borrowers? Since policymakers never implemented a sensible large-scale principal forgiveness program, the question cannot be answered by data alone.

We thus approach the question using a structural model. As we indicate above, the key innovation in our model is an idiosyncratic penalty for default. Scholars have long noted a wedge between borrowers' actual behavior and their optimal rational behavior. Guiso, Sapienza, and Zingales (2013) and others show that borrowers delay defaulting even when it is financially optimal, forgoing tens of thousands of dollars worth of financial value, and

sometimes over a hundred thousand dollars of financial value. This is a behavioral deviation from rationality too large to sweep under the rug, too large not to model explicitly.

The literature also establishes that different borrowers will endure different degrees of financial shortfall before succumbing to the incentive to default, a phenomenon as far as we are aware not captured by any previous structural model. We believe it is essential to pin down these differences: in order to know how different borrowers would respond to different modifications, we must also account for their significantly different tolerance for financial distress.

We thus explicitly model idiosyncratic borrower default penalties. A borrower’s penalty governs the interactions between her liquidity-derived and strategic motives. High-penalty borrowers more closely resemble the liquidity-constrained style of defaulter HAMP screened for, borrowers who continued making payments even when facing strong financial incentives to default. Low-penalty borrowers more closely resemble the strategic style of borrower.

In what we believe is also a first in the literature, we estimate the relationship between borrowers’ observable credit scores and their unobservable default penalties. Lenders offer lower interest rates to borrowers with higher credit scores—even though the lenders have full knowledge of borrowers’ income and employment, the characteristics of relevant collateral, and every other relevant financial variable—because even after taking these characteristics into account, higher-score borrowers are less likely to default. This phenomenon is widely known, yet has not been explained in a quantitative framework. Our model provides a structural, quantitative underpinning for this variation in the behavior of borrowers with different credit scores: higher-score borrowers default less to the degree determined by their greater default penalties. As far as we are aware, ours is the first structural model to provide a borrower-level quantitative underpinning for borrowers’ distinct degrees of reluctance to default. It is also the first to model a rich, high-fidelity state space for each of thousands of different borrowers’ decisions across hundreds of thousands of joint realizations of income, savings, house prices, house price expectations, mortgage terms, and more; and the first to use borrowers’ actual mortgages, borrower-by-borrower and mortgage-by-mortgage, to represent borrower heterogeneity.

We believe its support for complete heterogeneity in borrower characteristics also makes it the first to feature realistic enough predictions of distinct borrowers’ behavior that a lender or policymaker can use it to predict how different borrowers will vary their default propensity in different financial situations and house price climates. We demonstrate the model’s in- and out-of-sample predictive accuracy further in the companion paper, Kalikman and Scally (2023). In this essay, we use the model to study how different borrowers would vary their

default propensity in the different financial situation created by principal forgiveness.

We study a policy which selects borrowers in the lower half of the distribution of credit score at origination and who live in ZIP codes where exogenous price drops render their scheduled first-lien loan-to-value ratio at 130% or higher, i.e. owing over 30% more on their loans than their homes are worth. Taking into account the unique loan terms, financial characteristics, and ZIP-level house price dynamics that these borrowers would have faced, we model each distinct borrower’s optimal default decision after receiving a principal-forgiving modification that reduces her combined loan-to-value ratio to 95%. We find principal forgiveness very effective at stopping defaults, reducing the rate of defaults from 62% without modification to merely 10% with modification. But such an effective intervention comes at no cost to taxpayers: lenders would earn on average \$14,000 from forgiving principal, as lenders receive both principal and interest from borrowers who choose to keep making payments after receiving forgiveness, rather than the proceeds from a distressed home sale at foreclosure.

Across the United States, borrowers who would have been eligible for this modification program constituted between 3% and 5% of the population of approximately 100 million mortgages active during the Global Financial Crisis. Thus, our model predicts that such a targeted principal forgiveness program would have reached as many borrowers as were modified under HAMP, with far greater effectiveness, at no cost to taxpayers, and with a net benefit to lenders on the order of tens of billions of dollars, comparable to HAMP’s total cost to taxpayers of \$50 billion.

2 Literature

While consensus is emerging against the effectiveness of principal forgiveness, the literature provides evidence both for and against. We characterize the evidence on both sides and propose an explanation that reconciles the divergent findings within a unified framework. We argue that principal forgiveness failed when it failed because it was applied too little, too late, to too few—and to the wrong few. Our structural model of mortgage default extends existing models in the literature with support for heterogeneous borrower characteristics, including credit scores that correlate to idiosyncratic private penalties for default. Using this model, we show both why certain borrowers with similar financial incentives exhibit different behavior as well as how to target principal forgiveness, both in terms of which borrowers should receive it and how much should they receive, for it to be effective.

The debate over mortgage modification hinges upon the deeper and longstanding debate over why borrowers with negative equity default on their mortgages. If they default because exogenous shocks reduce their ability to afford payments, i.e. the liquidity-constrained model also frequently termed the “double-trigger” model of default, then the most effective modification will focus on reducing payments. If they default because they rationally calculate that continuing to pay a mortgage in deep negative equity is throwing good money after bad, the “strategic” model of default, then the most effective modification will focus on increasing their equity. Thus, the consensus in favor of payment-reducing over principal-forgiving modification is developing hand-in-hand with the consensus in favor of the double-trigger model of default.

We consider it more accurate to identify borrower behavior as falling along a spectrum. In our model, every borrower acts according to a blend of double-trigger and strategic behavior. Liquidity constraints reduce the amount a borrower can consume in the present if she continues making payments, while low expectations of future house price appreciation reduce expected future consumption. Both effects reduce the borrower’s willingness to make payments. But borrowers vary in *how* willing they are to endure hardship to make payments; equivalently, they vary in how strategically their behavior appears. One borrower will default whenever it is financially optimal; another will endure a significant shortfall in utility before succumbing to the incentive to default. Explicitly modeling this heterogeneity not only better explains why borrowers defaulted when they did; it is also necessary to simulate counterfactual policy experiments that exploit it.

In the Introduction, we remarked that we wish to avoid the moral associations that follow labels of borrowers as “strategic” defaulters. Beyond our desire to avoid a moral debate, we also consider the labeling exercise to be limiting from the standpoint of economic precision. Excellent work has been done refining classifications of defaults: from imprecise and widely-varying estimates of the percentage of defaulters who are strategic or double-trigger, as summarized in Table 8 of Gerardi, Herkenhoff, et al. (2013), to the more recent and robust work of Ganong and Noel (2023). Using their novel and compelling method, based on the work of Pearl (2022)—whose argument that a structural model provides the sine qua non of causal and counterfactual analysis also motivated our own structural approach—the authors find that a liquidity shock is a necessary condition for default for 94% of defaulters, and refine the distribution of strategic, double-trigger, and cash-flow (essentially, double-trigger without the trigger of negative equity) defaulters.

But why stop at three categories, or at better precision in delineating who belongs to which of those categories? Even if nearly all borrowers experience a liquidity shock before

they default, Gerardi, Herkenhoff, et al. (2013) point out that most borrowers who experience liquidity shocks still do not default. Labels lack the power to distinguish those who do from those who do not, and discard dimensionally-rich information.

Our view is that the natural evolution in modeling default is to move beyond refining the boundaries of generic labels to discarding the labels entirely and working directly with the behavior and motivations of individual borrowers. If all borrowers decide to default making the same financial choices, using the same financial variables, but with different values of those variables as inputs, then we do not need to categorize the borrowers into types once they do default. In particular, introducing a categorization constrains the ability of a theory to issue predictions. We receive rich, high-dimensional input regarding borrower financial characteristics, and can immediately issue rich, borrower-specific predictions of behavior. Categorizing borrowers, as an intermediate step in the pipeline from observation to prediction, would constrain the precision of those predictions; whereas retaining the rich information about each borrower enables studying how each borrower would behave in other circumstances.

Ganong and Noel (2023) point out that their findings are rationalized by a model which incorporates a high private cost of mortgage default. Ours is such a model. It explains the decision to default by providing an explicit quantitative representation of how each borrower weighs considerations that represent illiquidity and strategy, weighing her willingness to tolerate reduced consumption just for the sake of not defaulting against her expectations for future consumption if she defaults and saves money.⁷

2.1 Principal Forgiveness in the Literature

As we discussed in the introduction, the literature has found HAMP PRA ineffective, but we believe PRA was too narrow a program on which to draw large-scale conclusions about principal forgiveness. Scharlemann and Shore (2016) and Ganong and Noel (2020) exploit pseudo-random variation in assignment to principal forgiveness in HAMP PRA in order to study its effectiveness. Scharlemann and Shore (2016) use a regression kink approach. They find that mortgages that received principal forgiveness defaulted with a quarterly hazard of 3.1%, versus comparable mortgages that defaulted with a quarterly hazard of 3.8% without principal forgiveness, suggesting that principal writedown was modestly effective at stemming default. But they find it costly, estimating that PRA spent approximately \$320,000 per foreclosure it prevented. Ganong and Noel (2020) use a regression discontinuity

⁷We explore this avenue further in Kalikman and Scally (2023).

approach. They find writedowns cost an average of \$70,000 more than modifications featuring similar levels of payment reduction without principal forgiveness. But they find virtually no additional impact on default from principal forgiving writedowns. Clearly, the type of principal forgiveness employed in HAMP PRA, on its target population, was not effective.

But we cannot conclude that principal forgiveness could not have worked at all, or should be entirely eschewed now and in the future. First, the sample of borrowers who received PRA were selected from the larger sample of borrowers who were screened by HAMP's penalty-of-perjury hardship affidavit, a screen designed to ensure that they were liquidity-constrained. These are borrowers whom one should expect, at any level of payment reduction, to be least responsive to the amount of principal forgiveness, since forgiveness acts through the borrower's expectations of *future* returns to keeping the house. Second, PRA was implemented in the last quarter of 2010. But prices started to drop in 2007; many borrowers had already defaulted by PRA's implementation. Thus the pool of borrowers surviving long enough to enter the sample in PRA omitted the most strategic borrowers, the borrowers with the worst expectations of future house prices, and the borrowers with the worst financial circumstances. Third, over 85% of PRA borrowers remained deep in negative equity, with loan-to-value ratios of 115% or higher, even before taking into account second liens. Such borrowers would still be so deep underwater after PRA's forgiveness as not to have any reason to change their strategic calculus over default.

This last point is crucial and to some extent counterintuitive. The function mapping writedown size into lender value is discontinuous at the point where the writedown becomes large enough to change the borrower's decision to default. Borrower value is continuous in writedown size, but lender value is filtered through the borrower's discrete decision to default, and is thus discontinuous at the point of borrower indifference between continuing to pay and defaulting. Consequently, lenders will lose more and more money as they forgive more and more of a borrower's debt without incentivizing that borrower not to default, yet will suddenly increase their value by writing principal down an even greater amount. This jump will be to a larger value than any nearby writedown of less principal, and for some borrowers the value to the lender from the smallest writedown that keeps the borrower in the home constitutes the lender's global maximum value. Writedowns even more generous than the borrower's point of indifference between continuing and defaulting resume a downward slope in lender value. It thus stands to reason that the right-size writedown might have benefited lenders and borrowers, even if the data have only examples of wrong-size writedowns.

The literature offers evidence precisely for this claim. Haughwout, Okah, and Tracy (2016) found that private modifications, i.e. implemented without government involvement

and thus including borrowers not subject to HAMP’s double-trigger screen, were nearly four times as effective when they forgave principal as when not. All the mortgages in the Haughwout, Okah, and Tracy (2016)’s sample had been securitized in subprime deals and thus, as with HAMP PRA borrowers, do not constitute a nationally representative sample but a selected sub-population.

The literature thus establishes that limited forgiveness failed for PRA borrowers, but greater forgiveness succeeded for subprime borrowers. We argue that these results, which would be contradictory if generalized to the entire population, can be reconciled by a model in which every borrower is partially strategic and partially double-trigger, but in which borrowers vary person-to-person in the degree to which they resemble the highly stylized extremes. To see this requires taking into account more than just which population a borrower comes from and which writedown size the borrower received, but *all* of the financial incentives facing each borrower in each pool. We design our model to do just that.

2.2 Models of Default

Our model extends the literature on what Foote and Willen (2018) term *hybrid* models. Hybrid models are gaining traction both because they better fit the data and because they encode established phenomena in the literature. Borrowers are known to behave consistently with features of both the double-trigger and strategic models. In Elul et al. (2010)’s study of mortgage default, both negative equity and illiquidity are associated with default, and the authors find interactions between negative equity and illiquidity: borrowers in deep negative equity are more likely to default given a liquidity shock than borrowers with moderate negative equity. Hybrid models offer a means to quantifying the magnitudes of such effects and their interactions.

The model we use is introduced in Kalikman and Scally (2023). It builds upon existing mortgage default models in several dimensions crucial to disentangling policy and counterfactual questions pertaining to heterogeneous borrower populations. One of the most important is the dimension of borrowers’ non-pecuniary penalties for default. The literature provides strong evidence that non-pecuniary default penalties are both real and heterogeneous. Bhutta, Dokko, and Shan (2017) study subprime negative-equity borrowers and find that they need to be far more underwater before they default than predicted by theory, arguing that borrowers weigh “emotional and behavioral factors,” before defaulting—thus arguing for a non-pecuniary default penalty. White (2010) argues that the financial industry deliberately and successfully cultivates non-pecuniary default penalties in borrowers through

a campaign of proactive moral suasion, including by misrepresenting the severity of the consequences that borrowers would face for defaulting. And, as discussed in the introduction, Guiso, Sapienza, and Zingales (2013) provide evidence not only that non-pecuniary penalties are common to many borrowers, and considerable in magnitude, but also that they are heterogeneous. 77% of borrowers they surveyed reported that they would not default on a mortgage, provided they could afford their payments, even if the value of the loan exceeded the value of the home by \$100,000. Of the 23% of borrowers who *would* default at that shortfall, 61% of those would *not* default when the shortfall was only \$50,000.

We thus model heterogeneous penalties explicitly, in what we believe is a first for the mortgage default literature. We build on prior structural models of default and prior forays into the introduction of a default penalty, extending to a framework which embeds all of them as special cases. That framework is a structural optimal option exercise model in the style of Black-Scholes-Merton. Rational, forward-looking borrowers compute the optimal mortgage decision in each possible future, under uncertainty in exogenous variables such as home prices and labor income, and recurse backwards to the present.

Campbell and Cocco (2015) introduced the benchmark model, which featured many borrower characteristics but not credit score or a penalty for default. Several authors advanced the literature in the direction of including non-pecuniary default penalties, a modeling feature supported by prior research in the data. Schelkle (2018) modeled a default-like penalty as a utility benefit for owning a home, but which thus introduces the same disincentive to sell a house as to default on it. In our model, borrowers may choose to build wealth by selling their homes without penalty. Laufer (2018) used a hybrid model featuring a default penalty to study equity extraction, without mortgage age as a state variable. In our model, borrowers pay down principal as mortgages age, changing their equity-driven incentives to default.

Schelkle (2018) and Laufer (2018) find widely varying estimates of the *size* of the borrower’s default penalty, from 1.5% of permanent income at one extreme to 29% at the other. It is possible that the discrepancy is an artifact of too-strong representativeness assumptions: Schelkle (2018) considered only a single interest rate and loan-to-value ratio; Laufer (2018) calibrated his model to mortgages in one county. Widely varying house price dynamics and mortgage terms not captured by these representativeness assumptions could account for different estimates of borrowers’ default penalties. We run our model with heterogeneous data from thousands of distinct borrowers in distinct geographies with distinct loan terms.

We believe the discrepancy in other authors’ estimates of penalty size might not be only a modeling artifact but in fact a faithful signal of the underlying reality that borrowers have

heterogeneous penalties. As we have already argued, the literature provides several points of evidence for such heterogeneity. This heterogeneity is also a bedrock principle in the mortgage finance industry. Mortgage lenders set rates by forecasting borrowers’ propensity to default using observable characteristics, such as employment and income history, to get at unobservable characteristics, such as future willingness to make payments. The most prominent non-pecuniary observable characteristic every borrower has—in fact must submit to lenders, by law—is a credit score.

Thus, in what we believe is likewise novel in the default literature, we treat the credit score as partially revealing a borrower’s private and idiosyncratic default penalty, and therefore estimate the relationship between credit score and penalty. Pinning down this relationship enables us to use the model to perform counterfactual policy experiments with realistically interpretable outcomes, while restricting the exercise to using information that would have been observable to policymakers.

We argue that for a model of default to fit the observed variation in the behavior of borrowers, it must account for why different borrowers exhibit different degrees of tolerance to being underwater before they default. Further, it must account for each borrower’s distinct financial circumstances and local house price dynamics. We argue in the companion paper Kalikman and Scally (2023) that this modeling approach better fits the data on default than a model without default penalties, without idiosyncratic default penalties, or without idiosyncratic default penalties correlated to credit score. As far as we are aware, this is the first effort in the academic literature to estimate a structural model of default separately, loan-by-loan, for each loan in a heterogeneous, nationally representative sample with variation in loan terms, borrower characteristics, and regional price dynamics, and with idiosyncratic penalties for default.

In the remainder of this essay, we first discuss the mathematical features of the model and introduce the data, and then show how the model reconciles the literature’s divergent outcomes on principal forgiveness within a single framework, and finally apply the model to study optimal mortgage modification policy.

3 Model

Our model is a finite-period backwards-recursive optimal option exercise model, in the vein of Black-Scholes-Merton and Campbell and Cocco (2015). Risk-averse, fully rational borrowers with CRRA utility for real consumption, constrained by income, savings, and mortgage

payments, maximize utility on each node in a recombining tree of forecast house prices, labor shocks, and mortgage statuses. In each state in each period, the borrower decides not only whether to default, but also whether to prepay the mortgage, and if so, whether to keep or sell the home.

Because we wish to study policies that are not one-size-fits-all but rather targeted to different borrowers, we also need to use different inputs to the model for each different borrower. We require of our model that it accurately reflect how each different borrower would behave under the particular realization of house price paths and employment circumstances that befell the borrower while also representing the borrower decision in light of any changes in mortgage terms that a policymaker might wish to offer. Thus each borrower in the model imagines a *different* tree of future possibilities with its own distinct house price dynamics, interest rate path, and interactions with her own present and future income, wealth, and mortgage terms.

Additionally, each borrower knows her own default penalty, i.e. the size of her psychological reluctance to default even when it is financially optimal to stop making payments. As we explained in the literature review, significant research establishes the existence of the default penalty; in our utility function it is the mechanism which causes borrowers to forgo financial gain on the order of a hundred thousand dollars. The literature also establishes that the amount borrowers would forgo varies borrower to borrower; we design our model to be faithful to this reality. In the model, in any state of the world when the borrower defaults, she pays a penalty in utils, distinct from the financial consequences of default such as having to search and move to a rental home with volatile rent and losing the upside in future home price appreciation. Different borrowers suffer different penalties, corresponding to their varying degrees of reluctance to default even when financially incentivized to do so. But we suppose the policymaker knows only the borrower's observable characteristics: her credit score and her history of payments. The policymaker estimates the borrower's default penalty from these observable variables, and constructs a forgiveness policy accordingly.

Heterogeneous default penalties enable our model to embed the existing structural models of mortgage default as particular cases. The double-trigger, strategic, and hybrid model are reproduced by particular parametrizations of our model assuming certain loan terms, property characteristics, income dynamics, and values for penalties. Heterogeneous penalties are also the crux of accurately reflecting the effects of targeted policy design. In particular we believe modeling heterogeneous penalties helps resolve a conundrum in modification policy: because different borrowers actually vary in how likely they are to continue payments without modification, a policy that pre-emptively reduces payments needs to account ex ante

for the likelihood that it modifies a borrower who would have continued making payments without the need for a modification. Prior work studying modifications has been afflicted by the “curse of averages”—the need to extrapolate the behavior of the average borrower to the average of borrower behaviors. We sidestep the curse by directly modeling every distinct borrower’s distinct decision using distinct data on that borrower’s distinct financial circumstances *and* embedding that borrower’s distinct degree of deviation from the purely financial-maximizing behavior.

Our model identifies the unnecessarily modified borrowers as, all else equal, those with the highest penalties, as those borrowers would have been least likely to default even without modification. By the same token, our model reproduces distributional features of borrower mortgage behavior micro-founded at the individual borrower level, rather than by appeal to coarse buckets or groups that would not permit finer grains of modification.

We feed into the model data from McDash, Equifax, and CoreLogic, linked at the loan, borrower, and property level. The linked datasets enable seeing the nearly complete portrait of a borrower’s financial conditions that the model requires to generate its realistic distinct state space for each borrower, with that borrower’s particular mortgage terms, financial situation, property characteristics, and credit score. While these large-scale, high-resolution data have been more widely available in recent years, as far as we are aware ours is the first paper to calibrate a structural loan-level model of default to a broad collection of complete borrower financial circumstances, rather than to representative samples with heterogeneity only in limited dimensions. We also believe this work to be the first to embed idiosyncratic default penalties estimated to borrowers’ credit scores, providing a structural, micro-level underpinning for the aggregate relationship between credit scores and borrowers’ payment behavior.

In Kalikman and Scally (2023), we argue that the model better fits the data on default than benchmark models, and further that it predicts defaults out-of-sample precisely enough both in aggregate and across the distribution of borrowers to glean realistic policy conclusions from its predictions of borrower behavior. In the remaining sections of this essay, we focus on those policy conclusions. In the remainder of this section, we describe the mathematical structure of the model.

3.1 Time

Time is discrete; $t = 0$ in the reference period, $1, \dots, T$ in the loan’s maturity before or T' after modification. The reference period $t = 0$ varies in historical time based on when the

borrower is eligible for a modification.

3.2 Mortgages

Mortgages mature in T remaining periods. They have remaining principal M_0 in period 0, scheduled remaining principal M_t in period t , coupon payments $\{m_t\}_{t=1,\dots,T}$, and interest rates $\{r_t^m\}_{t=0,\dots,T-1}$. For fixed-rate loans, the constant per-period mortgage rate is $r_t^m = r^m$, and the mortgage coupon payment $m_t = m$ therefore satisfies the usual full amortization schedule

$$M_0 = \frac{m}{r^m} \left(1 - \frac{1}{(1 + r^m)^T} \right) \qquad M_t = (1 + r^m)M_{t-1} - m.$$

Our model differs from some others in featuring exact mortgage amortization schedules rather than an approximation via a steady-state mortgage with a geometric approximation to amortization. It supports loans of arbitrary product structure, including adjustable-rate mortgages, interest-only mortgages, and mortgages with balloon payments. HAMP modifications with principal forbearance, for example, yield mortgages with large balloon payments.

3.2.1 Modifications and Equity Share Agreements

Policymakers considering modifications worried that writing down principal enough that borrowers had positive equity, or near enough to positive equity that a fast price rebound would restore positive equity, would enable such borrowers simply to sell their homes and walk away from their mortgages, without any benefit accruing to the lender or taxpayer who financed the modification. The equity share agreement mitigates that possibility and provides upside to the financier of the modification in the form of an option. We consider agreements characterized by three parameters: a **price floor schedule** Q_f^t for the sale price, below which the option has no value to the lender, a **maturity** Q_t , after which, if the borrower has not sold the home, the option expires with no value to the lender, and a **sharing percentage schedule** Q_p^t , the percent of the sale price in excess of the price floor which the borrower must reimburse the lender if she sells the home in t (before the option expires). In the simple case, $Q_f^t = Q_f$ and $Q_p^t = Q_p$ are constant. We typically take the

option maturity to be the entire term of the loan after modification $Q_t = T$, the price floor to be the house price as of the date of modification $Q_f = P_0$, and the sharing percentage to be a constant $Q_p = 50\%$.

3.3 Utility

Borrowers have constant relative risk aversion (CRRA) utility over real consumption C_t^s for each feasible state s in each time t , with relative risk aversion denoted by γ and interperiod discount rates, which may vary, denoted by $\{\beta_\tau\}_{\tau=0,\dots,T-1}$, so that the gross period-0 discount factor in period t is $B_t = \prod_{\tau=0}^t \beta_\tau$:

$$U^i \left(\{C_t^s(i)\}_{t \in T}^{s \in S(t)} \right) = E_0 \sum_{t=0}^T B_t \left(\frac{(C_t^s(i))^{1-\gamma(i)}}{1-\gamma(i)} - \mathbf{1}_{i \text{ defaults in } s} \lambda(i) \right).$$

The expectation is taken in period 0 over all realizable states s . The penalty λ^i or $\lambda(i)$ directly reduces the borrower’s utility in and only in the state and period when she defaults. It is not, as in comparable models, a multiplier on utility in states when the borrower owns the home, as this would introduce the same disincentive to sell the house free-and-clear as to default on the mortgage. As is standard, the amount of housing or housing services consumed is fixed regardless of whether the borrower is an owner or renter, so we suppress a separate housing term from the utility function.

3.4 House Prices

House prices are uncertain, exogenous, and specific to each distinct ZIP code of each borrower in the data. They follow an approximate geometric Brownian motion on a recombining trinomial tree. Trinomial trees support a discrete space of house prices and transition probabilities that model a wide range of possible series in expected mean and variance of house price changes while also enabling paths of prices which “jump” without requiring borrowers to recalculate revised expectations. Borrowers in the baseline model have time-revising expectations which are short-term trend-following and long-term mean-reverting. We also consider specifications in which they have constant expectations or certain-but-not-constant expectations of a particular series of price changes and price volatilities. We provide further mathematical details in Appendix [A.1](#)

By contrast with many other models of default, we do *not* assume a representative house price process. Consider Figure 1. We show the overall path for house prices for the United

States, as well as for each of the twenty cities composing the S&P/Case-Shiller House Price Index. We highlight the three metropolitan areas with the largest peak-to-trough decline. The variation in price paths across the twenty metro areas, including both overall rise and fall, pace of recovery, and size of peak-to-trough decline, is striking. Rather than paper over these differences for computational tractability, we consider it of paramount importance to model how different borrowers in these different regions would respond to the distinct dynamics in their own areas. Borrowers in our model experience considerably different paths of prices, form different expectations, and thus face different default decisions.

[Figure 1 about here.]

3.5 Labor Income

Labor income follows a two-state Markov process reflecting “employment” or “unemployment,” where unemployment is understood to encompass any shock to liquidity—not only job loss, but also unexpected expenses from medical costs or divorce. The borrower’s income in the employed state is L_t^h , a multiple of her initial income L_0^h determined by her age and multipliers estimated from BLS data. In the unemployed state, the borrower receives replacement income L_t^l as a fraction τ_{unemp} of normal income. The replacement fraction may be thought of as unemployment insurance income or as the borrower’s income net of the shock. The constant probability of unemployment is π_u and constant probability of re-employment π_e .

3.6 States and Uncertainty

States are vectors encompassing the exogenous state variables house prices P_t^s and labor income L_t^s as well as the endogenous control variables liquid assets A_t^s , remaining mortgage terms \vec{M}_t^s , whether the borrower still owns or has sold the home H_t^s , the years left, if any, the borrower will live rent-free after defaulting and before the lender completes foreclosure proceedings D_t^s , and the years, if any, since the borrower accepted a modification with an embedded lender option / equity share agreement Q_t^s :

$$s = \langle t; A_t^s, \vec{M}_t^s, H_t^s, D_t^s, Q_t^s; L_t^s, P_t^s \rangle .$$

A typical log-scale discretization of borrower wealth yields an overall state space on the order of 300,000 nodes for one borrower with one mortgage.

3.7 Savings, Borrowing, Unsecured Debt

The borrower chooses a consumption plan $\{C_t^s\}_{t \in T}^{s \in \mathcal{S}(t)}$, specifying the consumption level at any feasible state s in any time period t , to maximize U subject to a budget constraint on liquid assets, or cash-on-hand, A_t^s .

The borrower invests cash savings at the risk-free interest rate r_t^f . In addition to their mortgages, borrowers may borrow in unsecured debt at a (higher) borrowing rate r_t^b . Savings and borrowing rates may vary over time.

Unsecured borrowing is rare in structural models of mortgage default because it increases the computational complexity of the state space and borrower optimization, but it is an important and realistic aspect of borrower behavior. In the Equifax CRIS data, 92% percent of borrowers have observed unsecured debt when they default on their mortgages; 55% of the population have debt in excess of \$15,000. We do not require the borrower to stay above an arbitrary lower bound in unsecured debt, but we also do not allow borrowers to default on unsecured debt.

The borrower thus begins state s' in the next period with liquid assets $A_{t+1}^{s'}$ equal to

$$A_{t+1}^{s'} = \begin{cases} (1 + r_t^f)a_t^s, & a_t^s \geq 0 \\ (1 + r_t^b)a_t^s, & a_t^s < 0, \end{cases}$$

where a_t^s is the end-of-period cash-on-hand as defined in the budget constraint below.

3.8 Budget Constraint

See [A.2](#) for a detailed description of the budget constraint.

4 Data and Methods

4.1 Data Sources

Microdata were generously provided by the Federal Reserve Bank of New York and consist of Black Knight Financial Services' McDash Data, Equifax Credit Risk Insight Servicing (CRIS) Data, and CoreLogic Home Price Indices linked at the loan, property, or ZIP level. Black Knight Financial Services provides the industry-standard mortgage servicing dataset known typically as "the McDash Data" (hereinafter). These are loan origination and per-

formance data as reported monthly by loan servicers. Coverage in McDash is estimated at approximately 80% of first liens originated in the United States, with detailed loan performance histories since 2005. McDash is considered a representative sample of first liens and includes not only prime but also subprime and Alt-A loans in the crisis era. The data cover both origination characteristics and performance characteristics. We use McDash origination data for original mortgage size, purchase house price, mortgage interest rate, mortgage product structure, mortgage purpose and occupancy—we restrict to purchases of owner-occupied homes—and borrower characteristics front-end payment-to-income ratio and credit score. We use McDash performance data for contemporaneous interest rate, payment history, history of principal outstanding, whether the loan has received a modification, and modified terms of modified loans.

Because McDash data do not provide adequate second-lien coverage, we use the Equifax CRIS data to identify combined loan-to-value ratios. The CRIS database consists of monthly observations of credit characteristics by borrower and tradeline type. Coverage is virtually 100% of the universe of borrowers as virtually all formal lenders report borrowers' payments to Equifax and the other major credit bureaus. The database reports several credit scores each month for each borrower. The results herein are not sensitive to the choice of credit score.

Data are broken down by month, borrower, and tradeline type, and include number of accounts, total balance, total balance past due, and total credit limit or high credit. Because we observe separate tradelines for first and second mortgages, we can compute combined loan-to-value ratios. Equifax CRIS also observes credit card debt, which we use as a measure of borrowers' unsecured borrowing.

We use CoreLogic Home Price Indices, joined as of the contemporaneous date and at the zip code level to the McDash-reported property zip code, to estimate both the path of house prices for each property and the parameters of each borrower's expectations for her house price process. We take the savings rate as the 1-year treasury rate and use average credit-card APRs from the Federal Reserve Economic Data portal for borrowing rates. We use data from the 2007-9 Survey of Consumer Finances to estimate borrower's asset levels, and income data from the Bureau of Labor Statistics to estimate the expected path of labor income over the lifecycle.

4.2 Methods

We employ standard methods for defining mortgage-related variables of interest and estimating quantities not directly observable. We define default and modification using our total view into payment history, which enables us to disentangle temporary payment lapses from long-term defaults. We estimate borrower income from payments and the reported front-end payment-to-income ratios. We estimate liquid assets using SCF data regressed against borrower income, mortgage debt, and non-mortgage debt. We estimate combined loan-to-value ratios by scaling McDash first-lien loan-to-value by total mortgage debt reported in Equifax CRIS. And we estimate baseline borrower expectations of house prices and volatility of house prices and the related parameters of the house price trinomial tree from the history of ZIP-level house price indices. We describe the methods in greater detail in Appendix A.

4.3 Sample Selection

Our initial sample of mortgages are those active as of January 1, 2007 and with combined loan-to-value ratio of 80% or higher. We restrict to mortgages used to purchase a home that the owner would occupy, with price at least \$20,000 and no more than \$1,000,000. We further restrict to 30-year, fixed-rate, non-interest-only mortgages. We restrict to conventional mortgages, i.e. excluding those backed by Federal Housing Authority or Department of Veterans' Affairs insurance. We require the borrower's credit score and front-end payment-to-income ratio (PTI) to be observed at origination. Finally, we restrict to borrowers who have a single first mortgage. This restriction excludes borrowers who are also financing vacation or investment properties. The financial incentives facing such borrowers are different, and they are unlikely to be the focus for government relief policies.

These sampling restrictions select for the population of mortgages used in most comparable studies, but they are a relatively safe pool versus mortgages in the crisis era. As Li and Goodman (2014) and others discuss, mortgage origination in the crisis era concentrated in riskier product structures, such as mortgages with low teaser rates that made initial payments small. Such mortgages appear affordable to borrowers until rates reset to a floating index after the teaser period. Similarly, interest-only loans feature low, affordable-feeling payments but without amortization, meaning that borrowers do not develop equity in the property and are therefore more likely to experience negative equity and thus to default.

We describe further sampling restrictions in Appendix A.3.1. For the most part, these and our other selection criteria are adverse to our hypothesis: selecting borrowers who choose the least risky mortgages selects for those borrowers who are also likely to have the highest

personal reluctance to default, i.e., default penalties. Our selection thus likely leads us to underestimate the overall effectiveness we estimate for principal forgiveness. The one exception is selecting against ARM loans, which may cause us to underestimate the role of liquidity constraints, as interest rate resets to higher levels would exacerbate borrowers' liquidity shocks. But by the same token, selecting for fully-amortizing loans rather than including interest-only loans likely causes us to underestimate the equity effects of borrowers with partially-, non-, or negatively-amortizing loans, who would have less equity and thus an even greater strategic default incentive. Incorporating such mortgages would be a fruitful direction for further research, although origination in such exotic product structures is rare today and may well remain uncommon as long as lenders remain wary of the risks Li and Goodman (2014) characterize as stemming from those product structures.

The middle column in Table 1 summarizes standard statistics on this population.

[Table 1 about here.]

4.4 Calibration of Penalties

As we described in Section 2, prior literature establishes that non-pecuniary penalties for default exist, are significant in magnitude, and vary significantly across individuals. But if they are unobservable, how can we feed them into our model? We take advantage of two features of the data and our model. First, we know that borrower credit scores, which are fully observable to lenders and policymakers, correlate to their reluctance to default. That correlation underlies why lenders charge higher interest rates to borrowers with higher credit scores, even though the lenders know the borrowers' income, assets, employment, and so on. Second, we know the history of borrowers' actual continuation and default decisions. Given their financial incentives, that behavior is only consistent with certain sizes of penalties.

We combine these two observations into an estimation procedure that we believe sits at the intersection of simplicity, accuracy, and usefulness to policymakers. In particular, we select a method that would have been implementable just after the crisis. Our estimation procedure has two steps. First, we assume a linear relationship between borrowers' observable credit scores and their penalties. Any parametrization of that relationship maps each borrower to her penalty via her credit score. Second, we estimate the parameters of that relationship by running the model with the implied penalties, selecting the parameters that maximize the fit of borrowers' actual default rates to the default rates implied by the model when it uses the parameter-implied penalties.

Mathematically, we begin with a pool of borrowers $\mathcal{L} = 1, \dots, i, \dots, N$ of loans active and current on payments at some time t . Letting $\lambda(i) = \lambda^i$ be borrower i 's penalty, we assume

$$\lambda^i = \alpha + \beta \cdot (\text{Credit Score})^i,$$

where we observe $(\text{Credit Score})^i$ for each borrower.⁸ Then, for any particular choice of α, β , we define two quantities: the empirical cumulative default rate in \mathcal{L}

$$CDR_{\mathcal{L}}^T = \frac{\# \text{ of loans in } \mathcal{L} \text{ that defaulted by } T}{\# \text{ of loans in } \mathcal{L}},$$

and the model-implied cumulative default rate, which is indirectly a function of α and β as they determine each model input λ^i :

$$\widehat{CDR}_{\mathcal{L}}^T(\alpha, \beta) = \frac{\# \text{ of loans in } \mathcal{L} \text{ model predicts default by } T \text{ given } \{\lambda^i\}_{i=1, \dots, N}}{\# \text{ of loans in } \mathcal{L}}.$$

(In the definitions above, both cumulative default rates condition also on the observed series of house prices between t and T , suppressed in the notation.) Finally, we estimate $\hat{\alpha}$ and $\hat{\beta}$ to minimize

$$\left| CDR_{\mathcal{L}}^T - \widehat{CDR}_{\mathcal{L}}^T(\alpha, \beta) \right|.$$

We test our procedure by estimating parameters on one random sample of loans but then selecting a new, entirely distinct sample and verifying that the model's out-of-sample performance does not degrade. We first take a randomly selected sample of loans active as of January 2007 and with combined loan-to-value ratio of 80% or higher. We select loans with high LTVs because they are more likely to experience negative equity and therefore to default. Characteristics of this sample are presented in the middle column of Table 1, with those of the defaulting borrowers presented in the rightmost column. Certain salient features of these statistics are that the characteristics of defaulting borrowers tilt consistently in the expected directions: they have larger loans, a higher percentage of second liens, higher second-lien balances, higher LTVs, higher payment-to-income ratios, lower wealth, lower income, and lower credit scores. Origination vintage skews heavily towards more recent loans (2005 and 2006 vintage) because these loans are those less likely to have amortized below 80% LTV.

We then draw a new sample of loans with the same sampling restrictions, and run the

⁸While the equation relating credit score and penalty suggests estimation by linear regression, this is not our approach and in fact is not possible as we do not know the entries of the vector $\{\lambda^i\}$ ex ante. In a different direction, we also recognize that more variables besides credit score may correlate to penalties. We consider this and related possibilities in Appendix E.

model again using the estimated coefficients to determine each borrower’s penalty. We compare the predicted default rate generated by the model to the default rate of these distinct loans in actuality. The in-sample and out-of-sample fit are displayed in Figure 2.

[Figure 2 about here.]

We find very strong out-of-sample performance: the model fits the level of defaults out-of-sample even though none of those loans were used to calibrate the model. In this graph, the model calibration targets matching the year-5 cumulative default rate. Thus the shape of each curve, as well as the terminal values at year 7, are all out-of-sample estimates that we take to validate the calibration’s performance. The model not only fits well on the out-of-sample pool; it also fits the shape of defaults out-of-sample on years 0 through 4 and 6 and 7. These results give us confidence that the model fits defaults accurately enough to use for drawing realistic policy conclusions. The model also fits the cross-sectional characteristics of defaulting borrowers; in Appendix E, we show that this is a genuine feature of idiosyncratic penalties and not merely an artifact of correlation between borrowers’ credit scores and other risk factors. The model’s out-of-sample performance is robust to selection criteria and calibration method, e.g. calibrating in-sample and out-of-sample on distinct vintages.⁹

4.5 Summary of Model Parametrization

The full set of model parameters, data sources, and baseline values are summarized in Table 2. The bulk of parameters, those listed below without a numerical value, are input into the model as different values for each borrower i , depending on data in the indicated sources according to the inference methods described in the preceding sections. We assume fixed values primarily for transaction costs and for variables other than those that pertain to mortgages and house prices. In the baseline specification, we take the borrower’s interperiod discount rate β to be constant at 0.985. This value is consistent with other authors’ selected values, as well as with risk-free rates of approximately 1.5%, approximately those that prevailed throughout the period of interest. We set the coefficient of relative risk aversion to 3.5 again for consistency with other authors’ parameter selection. Our results are robust to parameter selection within ranges considered reasonable by other authors.

[Table 2 about here.]

⁹See Kalikman and Scally (2023).

5 Results

We test the effectiveness of targeted principal forgiveness by selecting borrowers from our sample who experience deep negative equity through “no fault of their own.” This is to say, we deem a borrower eligible once her home would be 30% underwater even if she had paid her mortgage in full. This way, borrowers cannot increase the odds of their eligibility by missing payments.

Mathematically, we extend the idea of scheduled principal remaining to a new concept, the *Scheduled Loan-to-Value Ratio* (SLTV). Scheduled loan-to-value divides the principal balance scheduled to be remaining on the first mortgage at any time t by the *actual* house price that obtains at that time:

$$\text{Scheduled LTV} = \text{SLTV}_t = \frac{M_t^0}{P_t},$$

where M^0 is the initial principal balance, M_t^0 is the scheduled principal remaining at t , which does not depend on the borrower’s payment history, and

$$P_t = \frac{\text{HPI}_t^z}{\text{HPI}_0^z} P_0$$

is the house-price-index-adjusted house price at time t for the home (given its zip code z). Scheduled LTV at a given date cannot be known in advance of that date by the borrower or the policymaker, because it depends on the realized house price. But policymakers in 2009 could have observed the 2009-scheduled LTV for all mortgages active then (and thereafter observed contemporaneous SLTV for all mortgages), so this policy both would have been implementable and would not have introduced a perverse incentive.

Approximately 8% of borrowers in our overall sample are included in the sub-sample of borrowers who experience 130% SLTV. Borrowers experience these events primarily starting in 2008 and as late as 2012. We restrict our attention to those whose first 130-SLTV event occurs from January 2008 through December 2011. As we showed in Figure 1, the timing of price drops varied considerably in different regions, again emphasizing the importance of selecting a heterogeneous sample and running the model separately for each loan with house price dynamics particular to that loan.

Borrowers who experience 130% SLTV are overall similar to the population while exhibit ex ante signals of slightly greater risk: higher initial loan-to-value ratio, higher initial payment-to-income ratio, higher mortgage interest rates, and lower credit scores than the

average borrower, as summarized in the rightmost column of Table 3.

[Table 3 about here.]

The specific principal forgiving modification we study is as follows. On the first date that a borrower experiences 130% SLTV, a policymaker offers her a writedown to 95% combined loan-to-value ratio, i.e. forgiving enough debt on both first and subordinate liens so that the borrower has some positive equity net of all liens. (Recall that the inclusion criterion applies only to the first lien, but if a borrower had a second lien, we assume that the modification would take both liens into account in order to be sure to restore the borrower to positive equity.) Borrowers have the option of refusing modification if they calculate the present value of the modification to be less than the present value of continuing or defaulting with their pre-existing mortgage. Borrowers might refuse because if they accept, they must agree to an equity share agreement (see Section 3.2.1) should prices suddenly rebound and incentivize them to sell the home.

We find writing down principal for all these deep underwater borrowers prevents many defaults but does not pay for itself. This result is broadly consistent with other findings in the literature, which claim positive costs for forgiveness. Under the baseline scenario, the model predicts that approximately 55% of the 130% SLTV borrowers would default within five years. After writedown, about 8% of borrowers will still eventually default.

[Figure 3 about here.]

A principal writedown to 95%-CLTV for these borrowers costs on average \$68,500, but prevents about 85% of defaults that would occur under the baseline scenario.¹⁰ The writedown gives away 49 cents upfront of each dollar to borrowers who would not have defaulted—less than Gerardi, Herkenhoff, et al. (2013) estimated, because this is a more targeted policy, but still a seemingly high amount of waste.¹¹

However, the measurement of upfront expenditure is not the best measure of waste, as the forgiven dollars earn lenders more returns from future interest payments from borrowers who remain in the home as well as savings on expected losses from foreclosed properties. Taking into account the improvement in receipts from borrowers who would have defaulted but

¹⁰Note that those few borrowers who do default do not do so immediately after receiving the modification: the prices in their ZIP codes fall further, rendering them back underwater, after which point they default.

¹¹This figure differs slightly from the percent of borrowers who would not have defaulted without modification because borrowers have different loan balances.

continue paying because of forgiveness, we calculate that the average cost of a modification—the difference in total lender receipts after all modifications versus in the baseline, divided by the number of loans modified—is only \$5,600 in this sample.¹² The cost per prevented default likewise totals only \$12,900, significantly less than the cost estimated in other work for HAMP modifications. Depending on a policymaker’s estimate of the social cost of a foreclosure, this may be a small cost to pay.¹³

5.1 Narrowing the Target

As this experiment is only one of many tractable ways to estimate the cost of principal forgiveness for a deep negative-equity pool of borrowers, it produces a lower bound on the effectiveness of principal forgiveness.¹⁴ Could policymakers do better by targeting even more narrowly? A common feature of principal forgiveness proposals was an inclusion criterion meant to select for borrowers who were likeliest to default and likely to have the highest losses given default. Geanakoplos and Koniak (2009) proposed modifying only subprime borrowers, as it was those borrowers on whom lenders lost the most money foreclosing. It was also these borrowers whose defaults were cascading into systemic failure in the financial system, making them an appropriate target for policymakers. Zingales (2008) proposed including only borrowers in ZIP codes where prices fell by a certain threshold.

¹²We find a consistent result when examining a rudimentary implementation of the intervention proposal in Zingales (2008). Zingales’s proposal would reach more borrowers because the inclusion criterion is less restrictive: borrowers who were scheduled 30% underwater on amortizing mortgages necessarily live in ZIP codes where prices fell by over 20%. But the cost per prevented default would be an order of magnitude higher than that for targeted forgiveness.

¹³This figure also relies on an estimate of the recovery to lenders of selling a foreclosed home of 65% of the home’s value, which is standard in the literature but would still earn the lender more than the 25% recovery on lost principal. Geanakoplos and Koniak (2009) estimated subprime foreclosures yielded on average.

¹⁴The lower bound arises for two reasons. First, for any one borrower, the best amount of forgiveness, from the lender’s point of view, may not be that which restores the borrower to 95% CLTV, or even one that restores the borrower to positive equity at all. Even in a model with frictions or transaction costs and penalty all set to zero, borrowers should not default precisely at the negative equity threshold if their expectations for future price increases are high enough and the carrying costs of continuing on the mortgage low enough. Accordingly, it may be sufficient to restore borrowers to near-, but not actually, positive equity in order to dissuade them from defaulting, yielding a cheaper upfront expenditure. Second, borrowers differ. The indifference-threshold forgiveness required to induce one particular borrower not to default may be to a higher or lower CLTV than that required to induce a different borrower not to default. The forgiveness policy we studied offered different borrowers different levels of forgiveness: some borrowers received none; others received a lot. The most effective forgiveness policy would almost certainly distribute a much more diverse menu of offers to different borrowers. Even if political considerations required policymakers to offer a substantively equal modification to all borrowers—whether by forgiving the same dollar-denominated amount of debt, or by forgiving all borrowers to the same CLTV—then again 95% might not be the optimal CLTV level net of forgiveness. Forgiving less principal might spend less money, while forgiving more principal might prevent more re-defaults.

We design a refinement of our targeted forgiveness policy to test the ideas in these two suggestions, targeting borrowers' credit scores and their homes' locations. Borrowers with lower credit scores are more likely to default, a fact the model reproduces as these borrowers have lower penalties. Borrowers in neighborhoods worse-hit by house priced declines are also more likely to default, not only because price drops reduce their equity, but also because those price drops lower their expectations, reducing their estimated present value of future house price gains if they keep their homes. These inclusion criteria would have been ex ante available and observable to policymakers. And targeting such borrowers, any modification, not just principal forgiveness, should be more cost-effective.¹⁵ But can simply selecting borrowers more restrictively be enough to shift a forgiveness policy all the way from costly to cost-free? Or, as some authors have argued, must any forgiveness program incur positive costs, which thus can only be justified in light of greater social costs of foreclosure?

Two implementability concerns might arise with offering forgiveness only to low credit score borrowers. First, might it persuade borrowers to take actions that harm their credit score? Second, might it be seen as rewarding irresponsible behavior? We mitigate the first concern by using borrowers' credit scores as of loan origination, not their contemporaneous credit scores. Thus borrowers still could not take a decision which would affect their eligibility for the modification program after its announcement.¹⁶ We recognize that the second concern may have presented policymakers with a politically fragile program, but we consider the program worth studying regardless of how commentators might perceive the morality of helping those some label irresponsible. The concern of *political* viability might also have been mitigated by the broader observation that providing enough help to stem foreclosures would, in light of foreclosures' strong negative externalities and the continually unwinding crisis, make the entire economy better off.

In Table 4, we summarize the results of the writedown modification based on applying it only to each of four buckets: lower-half or upper-half of the sample of neighborhoods based on which would have the highest expectations of Home Price Appreciation, and lower-half or upper-half of the sample of credit score. Despite some correlation between price expectations

¹⁵Three main factors drive the net cost or benefit of a forgiveness policy: first, the degree of effectiveness of targeting forgiveness only to those borrowers who would have defaulted without modification, second, the degree of effectiveness of persuading those borrowers who do receive forgiveness not to default—that is, offering a writedown generous enough to change behavior—and finally, the degree of forgiveness that exceeds that necessary to change behavior, i.e., offering forgiveness not much more generous than that needed. In our refinement, we isolate the effect of the first—targeting—while holding constant the size of forgiveness and, therefore, the propensity of each individual borrower admitted to the policy to re-default.

¹⁶Would such a program cause riskier behavior in equilibrium, as borrowers anticipated that they might be “bailed out” of unaffordable loans? In short, no: lenders, not policymakers, price borrower risk, and lenders would not lend to borrowers who were too risky. For further discussion see Section 5.2.

and credit score, buckets end up relatively uniformly distributed, with between 23.5% and 26.5% of the sample population in each. The table shows that principal forgiveness loses money only when it forgives principal for the best-off or least likely to default borrowers: those with high home price appreciation expectations and high credit scores or penalties. Only a quarter of these borrowers would default. Forgiving principal of low-expectations borrowers is self-financing, and forgiving low-expectations low-penalty borrowers earns \$29,800 per such borrower.

[Table 4 about here.]

We conclude that if policymakers focused a mortgage modification intervention only on subprime borrowers, they would almost certainly find the principal writedown program self-financing. There is not a standard definition used to identify subprime loans in our data, but mortgages in the lower portion of the distribution of credit scores are more likely subprime. As shown in Table 5, at any reasonable threshold for including such borrowers in a modification program, modifying such borrowers would still significantly reduce their defaults, and would earn lenders on the order of \$14,000 to \$21,000 per modified loan.

[Table 5 about here.]

The population of borrowers with low credit score and who experience 130% SLTV constituted between 3% and 5% of the national population of approximately 100 million mortgages. Thus, a forgiveness policy for subprime borrowers might have reached as many borrowers as were modified under HAMP, with far greater effectiveness, and at no cost to taxpayers—in fact with a net benefit to lenders on the order of HAMP’s total cost of \$50 billion.

5.1.1 Highly Targeted Refinements

We are excited by the idea that we could use our model to determine the best distinct forgiveness amount to offer to each borrower individually (and independent of offers to other borrowers), the best blanket policy to offer uniformly to all borrowers, or any optimal policy under any other set of distributional constraints. But we do not wish to gild the lily. Such refinements will eventually exhibit declining marginal returns to complexity, just as the precision of the estimates of value of such improvements may degrade.

The same applies to questions not of whom to include or how much to forgive, but of how precisely to structure a modification, e.g. by favoring forbearance over forgiveness. We

do not feel this question is the relevant one for our study:. We seek to settle a concrete counterfactual question regarding a particular modification program that was proposed but not implemented, and that could have been implemented the benefit of a computationally intense model. What would have happened?¹⁷

We welcome the possibility that policymakers might wish to use our model to study optimal modification policy design in the future, and trust them to find the right balance of simplicity versus precision-targeting We believe we have made an important contribution by establishing that principal forgiveness is a powerful tool that can and should remain in policymakers' toolkit. We believe these results already sufficient to show that principal forgiveness can be effective, that when it was not effective, it was ineffective because it targeted the wrong populations in the wrong dose, and that one can make it even more effective by refining the degree to which one targets who receives it and how.

5.2 Discussion

We have attempted to show that a modification that writes down principal for borrowers whose mortgages would be underwater solely because of price declines introduces no perverse incentives, costs little on average to implement, and would likely self-finance for a broad swath of those borrowers. We report results only for loans with complete relevant documentation and only for vanilla 30-year fixed-rate purchase mortgages. Including only these loans likely understates the effectiveness of the modification regimen studied: the effectiveness of the selection criterion is largely achieved by the higher propensity of selected borrowers to default absent a modification; borrowers with exotic loan terms, particularly those without amortization, would have been more likely to default, and therefore may have been even more likely to push our estimates of the cost-effectiveness of principal forgiveness upward.

In this specification of the model, we do not allow for correlation between labor income shocks and house price shocks. But there is evidence that income shocks were worse in areas that suffered greater house price declines. Accounting for such correlation would likely further amplify the effectiveness of the principal forgiveness studied, as borrowers facing both negative equity and reduced income would be even more likely to default, all else equal, and therefore fewer forgiven dollars would flow to borrowers who did not need them.

Our model does not consider general equilibrium effects. While we are excited about the potential to expand the model to incorporate general equilibrium effects in future research, herein we are inclined to eschew the question of whether modification is efficient in general

¹⁷A more detailed discussion is provided in Appendix [B.3](#).

equilibrium. Policymakers acted in the financial crisis; their actions reveal that they considered the crisis a disequilibrium systemic event in need of intervention. How should they have intervened?

Nonetheless, we believe that extending the model to general equilibrium would likely enhance the measured effectiveness of principal forgiveness. Forgiving debt early enough could have stemmed the tide of foreclosure, and led to an earlier and faster recovery in house prices, enabling the economy to avoid the negative consequences of the collapse that ensued.¹⁸ The positive knock-on effects to the rest of the economy are missed in our present estimates.

What about the concern that borrowers in general equilibrium might internalize the expectation of a future “borrower bailout,” be thus incentivized to take out riskier loans in the future, and thereby cause even greater economic damage later on? Three considerations mitigate the possibility.

First, the historical probability of a coordinated, national drop in house prices on the order of that experienced in the financial crisis is extraordinarily low. Equilibrium borrowers would estimate the probability of bailout based on the likelihood of a coordinated crisis, not based solely on the likelihood of their own mortgage experiencing negative equity.

Second, in our program, taxpayers do not bear the cost of modification, and lenders are not forced to modify—they do so to improve their own bottom lines, with the knock-on effect that both they and borrowers avoid the destructive externalities of foreclosure. Since the policy entails no government bailout, borrowers would not have any reason to form an expectation of a government bailout.

Third and finally, there is no risk of incentivizing borrowers to *engage* in riskier behavior in the future, even if they did anticipate a bailout. Again, lenders bear the cost of modification in our study, and agree to the renegotiation because it is in their own financial interest. Borrowers might therefore *wish* to take out riskier loans, but they cannot receive them simply because they demand them. If lenders did not calculate such new, riskier lending to be to their benefit, they would not lend. Rather, equilibrium lenders would anticipate the behavior of their borrowers, and would not agree to issue loans to borrowers inclined to reckless future behavior—or would price such risk correctly in the interest rates on and quantity of loans supplied.

¹⁸Not all authors agree with this intuition; Kaplan, Mitman, and Violante (2020) provide an analysis in which forgiveness has a still positive but very modest impact on economic recovery as measured by consumption, and no impact on house prices.

6 Conclusion

We leverage a borrower-level structural model to show that principal forgiveness can work and can be cost-effective when targeted in the right dose, to the right populations, and in the right circumstance. We show that forgiving the principal of borrowers with low original credit scores living in neighborhoods where exogenous price drops rendered their mortgages deep underwater could have helped millions of people who suffered in the Global Financial Crisis, and perhaps even thereby averted further economic disaster from that crisis, without cost to lenders. We show that lenders may have stood to gain as much from such a program as taxpayers paid for a modification program that proved far less effective. We show that those borrowers whom policymakers could efficiently modify were identifiable *ex ante*, that the modification policy was implementable at the time of the crisis, and that it would have avoided perverse incentives.

Our work suggests a resolution to a developing conundrum in the mortgage literature: it explains disparate results regarding the effectiveness of principal forgiveness as arising from comparing incomparable treatments on incomparable groups. We rationalize the disparate outcomes on both groups as consistent with a single model with individual-level differences in willingness to suffer financial disutility before defaulting—the default penalty. We thus argue that principal forgiveness may continue to be an effective tool for foreclosure mitigation, particularly when applied to borrowers with a low reluctance to default, high negative equity, high losses given default, and low expectations of future price increases.

Our results contribute to a growing literature which recognizes the importance of exploiting heterogeneity in borrower characteristics to create opportunities for designing effective policies that assumptions of representativeness would miss. They also thereby demonstrate the policy relevance of borrower-specific, highly realistic computational structural models, models which Pearl and Mackenzie (2018) argue are the only way to fully disentangle the web of interactions that underlie rational individuals’ economic behavior and generate both robust explanations for observed behavior and realistic predictions of future behavior in various circumstances.

We believe our contributions also meet an increasingly recognized need for economic models that take distributional variation into account. In 2016, when she was Chair of the Board of Governors of the Federal Reserve System, Janet Yellen remarked that

Even though the tools of monetary policy are generally not well suited to achieve distributional objectives, it is important for policymakers to understand and monitor the effects of macroeconomic developments on different groups within

society.¹⁹

More recently, current Fed Chair Jerome Powell said in 2021 that

disparities in economic outcomes along various lines have been a persistent and growing feature of our economy for several decades. [The Fed] can contribute to addressing those through our tools. It does, of course, take a much broader set of policies from the legislature to really make significant progress on this.²⁰

By leveraging a structural model to show the efficacy of a targeted policy intervention, we have attempted to make a contribution to fill the gap identified by Chair Yellen’s and Chair Powell’s remarks. Economists and policymakers can understand and address distributional consequences of individual-level behavior and enter a new era of policy that recognizes individual differences by using high-fidelity models of individual behavior, models which are realistic about idiosyncratic differences not only in their financial terms but also in how they deviate from idealized behavior. Further research into such models will enable economists to derive a distributionally-accurate understanding of the economy. And it will help policymakers implement fairer, more effective policy.

References

- Adelino, Manuel, Kristopher Gerardi, and Paul S. Willen (2013). “Why Don’t Lenders Renegotiate More Home Mortgages? Redefaults, Self-Cures and Securitization”. In: *Journal of Monetary Economics* 60.7, pp. 835–853. ISSN: 0304-3932.
- Agarwal, Sumit, Gene Amromin, et al. (Dec. 1, 2011). “The Role of Securitization in Mortgage Renegotiation”. In: *Journal of Financial Economics* 102.3, pp. 559–578. ISSN: 0304-405X.
- Agarwal, Sumit, Kristopher Gerardi, and Vincent Yao (May 18, 2018). *The Political Economy of Loan Modification*. SSRN scholarly paper ID 3180850. Rochester, NY: Social Science Research Network.
- ATTOM Data Solutions (Jan. 15, 2019). *U.S. Foreclosure Activity Drops to 13-Year Low in 2018*. URL: <https://www.attomdata.com/news/most-recent/2018-year-end-foreclosure-market-report/> (visited on 12/09/2019).
- Bezanson, Jeff et al. (July 19, 2015). *Julia: A Fresh Approach to Numerical Computing*. arXiv: 1411.1607 [cs]. URL: <http://arxiv.org/abs/1411.1607>. preprint.

¹⁹Yellen, 2016.

²⁰Powell, 2021.

- Bhutta, Neil, Jane Dokko, and Hui Shan (Dec. 2017). “Consumer Ruthlessness and Mortgage Default during the 2007 to 2009 Housing Bust: Consumer Ruthlessness and Mortgage Default”. In: *The Journal of Finance* 72.6, pp. 2433–2466. ISSN: 00221082.
- Board of Governors of the Federal Reserve System (US) (Nov. 1, 1994). *Commercial Bank Interest Rate on Credit Card Plans, All Accounts*. URL: <https://fred.stlouisfed.org/series/TERMCBCCALLNS> (visited on 10/25/2021).
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (Feb. 2018). “Diagnostic Expectations and Credit Cycles”. In: *The Journal of Finance* 73.1, pp. 199–227. ISSN: 00221082.
- Bradley, Michael G., Amy Crews Cutts, and Wei Liu (June 2015). “Strategic Mortgage Default: The Effect of Neighborhood Factors:” in: *Real Estate Economics* 43.2, pp. 271–299. ISSN: 10808620.
- Braun, Lauren (Jan. 25, 2018). *U.S. Housing Market Has Gained Back All \$9 Trillion in Value Lost During Recession*. Zillow.
- Bricker, Jesse et al. (2011). “Surveying the Aftermath of the Storm: Changes in Family Finances from 2007 to 2009”. In: p. 38.
- Campbell, John Y., Nuno Clara, and João F. Cocco (Aug. 2020). *Structuring Mortgages for Macroeconomic Stability*. Working paper 27676. National Bureau of Economic Research.
- Campbell, John Y. and João F. Cocco (Aug. 2015). “A Model of Mortgage Default”. In: *The Journal of Finance* 70.4, pp. 1495–1554. ISSN: 00221082.
- (Nov. 1, 2003). “Household Risk Management and Optimal Mortgage Choice”. In: *The Quarterly Journal of Economics* 118.4, pp. 1449–1494. ISSN: 0033-5533.
- Case, Karl E., Robert J. Shiller, and Anne K. Thompson (2012). “What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets”. In: *Brookings Papers on Economic Activity* 43 (2 (Fall)), pp. 265–315.
- Corbae, Dean and Erwan Quintin (2015). “Leverage and the Foreclosure Crisis”. In: *Journal of Political Economy* 123.1, pp. 1–65.
- Corradin, Stefano (June 1, 2014). “Household Leverage”. In: *Journal of Money, Credit and Banking* 46.4, pp. 567–613. ISSN: 1538-4616.
- Crews Cutts, Amy and William A Merrill (Mar. 2008). *Interventions in Mortgage Default: Policies and Practices to Prevent Home Loss and Lower Costs*. Working paper. McLean, VA: Freddie Mac, p. 78.
- Dubey, Pradeep, John Geanakoplos, and Martin Shubik (2005). “Default and Punishment in General Equilibrium¹”. In: *Econometrica* 73.1, pp. 1–37. ISSN: 1468-0262.
- Eberly, Janice and Arvind Krishnamurthy (2014). “Efficient Credit Policies in a Housing Debt Crisis”. In: *Brookings Papers on Economic Activity* 2014.2, pp. 73–136. ISSN: 1533-4465.

- Elul, Ronel (2016). “Securitization and Mortgage Default”. In: *Journal of Financial Services Research* 49.2, pp. 281–309. ISSN: 1573-0735.
- Elul, Ronel et al. (2010). “What “Triggers” Mortgage Default?” In: *The American Economic Review* 100.2, pp. 490–494. ISSN: 0002-8282. JSTOR: [27805045](https://www.jstor.org/stable/27805045).
- Fannie Mae (Oct. 2, 2019). *Loan-Level Price Adjustment (LLPA) Matrix*.
- Federal Reserve Bank of New York (May 2019). *Quarterly Report on Household Debt and Credit*.
- Federal Reserve Board (2019). *Delinquency Rate on Single-Family Residential Mortgages, Booked in Domestic Offices, All Commercial Banks*. URL: <https://fred.stlouisfed.org/series/DRSFRMACBS> (visited on 12/09/2019).
- Ferreira, Fernando and Joseph Gyourko (June 2015). *A New Look at the U.S. Foreclosure Crisis: Panel Data Evidence of Prime and Subprime Borrowers from 1997 to 2012*. w21261. Cambridge, MA: National Bureau of Economic Research.
- Foote, Christopher, Kristopher Gerardi, Lorenz Goette, et al. (Apr. 2010). “Reducing Foreclosures: No Easy Answers”. In: *NBER Macroeconomics Annual 2009, Volume 24*. Vol. 24. University of Chicago Press, pp. 89–138.
- Foote, Christopher L, Kristopher Gerardi, and Paul S Willen (May 30, 2008a). *Subprime Facts: What (We Think) We Know about the Subprime Crisis and What We Don’t*. Working paper. Federal Reserve Bank of Boston, p. 57.
- (Sept. 2008b). “Negative Equity and Foreclosure: Theory and Evidence”. In: *Journal of Urban Economics* 64.2, pp. 234–245. ISSN: 00941190.
- Foote, Christopher L. and Paul S. Willen (Nov. 2018). “Mortgage-Default Research and the Recent Foreclosure Crisis”. In: *Annual Review of Financial Economics* 10.1, pp. 59–100. ISSN: 1941-1367, 1941-1375.
- Foster, Chester and Robert Van Order (1984). “An Option-Based Model of Mortgage Default”. In: *Housing Finance Review* 3.4, pp. 351–372.
- Fuster, Andreas and Paul S. Willen (Nov. 2017). “Payment Size, Negative Equity, and Mortgage Default”. In: *American Economic Journal: Economic Policy* 9.4, pp. 167–191. ISSN: 1945-7731, 1945-774X.
- Ganong, Peter and Pascal Noel (Oct. 2020). “Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession”. In: *American Economic Review*. NBER Working Paper Series 110.10, pp. 3100–3138. ISSN: 0002-8282.
- (May 1, 2023). “Why Do Borrowers Default on Mortgages?” In: *The Quarterly Journal of Economics* 138.2, pp. 1001–1065. ISSN: 0033-5533.

- Geanakoplos, John (Mar. 2014). “Leverage, Default, and Forgiveness: Lessons from the American and European Crises”. In: *Journal of Macroeconomics* 39, pp. 313–333. ISSN: 01640704.
- (Aug. 2010a). “Solving the Present Crisis and Managing the Leverage Cycle”. In: *FRBNY Economic Policy Review*.
- (Apr. 2010b). “The Leverage Cycle”. In: *NBER Macroeconomics Annual 2009, Volume 24*. Vol. 24. University of Chicago Press, pp. 1–65.
- Geanakoplos, John D. and Susan P. Koniak (Mar. 4, 2009). “Opinion — Matters of Principal”. In: *The New York Times. Opinion*. ISSN: 0362-4331.
- (Oct. 29, 2008). “Opinion — Mortgage Justice Is Blind”. In: *The New York Times. Opinion*. ISSN: 0362-4331.
- Gerardi, Kristopher, Adam Hale Shapiro, and Paul Willen (Sept. 1, 2009). *Decomposing the Foreclosure Crisis: House Price Depreciation versus Bad Underwriting*. SSRN scholarly paper. Rochester, NY: Social Science Research Network.
- (Dec. 3, 2007). *Subprime Outcomes: Risky Mortgages, Homeownership Experiences, and Foreclosures*. SSRN scholarly paper. Rochester, NY: Social Science Research Network.
- Gerardi, Kristopher S., Kyle Herkenhoff, et al. (2013). “Unemployment, Negative Equity, and Strategic Default”. In: *SSRN Electronic Journal*. ISSN: 1556-5068.
- Gerardi, Kristopher S., Andreas Lehnert, et al. (2011). “Making Sense of the Subprime Crisis”. In: *Lessons from the Financial Crisis*. Wiley-Blackwell, pp. 109–117. ISBN: 978-1-118-26658-8.
- Ghent, Andra C. (June 1, 2011). “Securitization and Mortgage Renegotiation: Evidence from the Great Depression”. In: *The Review of Financial Studies* 24.6, pp. 1814–1847. ISSN: 0893-9454.
- Ghent, Andra C. and Marianna Kudlyak (Sept. 2011). “Recourse and Residential Mortgage Default: Evidence from US States”. In: *Review of Financial Studies* 24.9, pp. 3139–3186. ISSN: 0893-9454, 1465-7368.
- Glaeser, Edward L. and Todd M. Sinai, eds. (2013). *Housing and the Financial Crisis*. National Bureau of Economic Research Conference Report. Chicago ; London: The University of Chicago Press. 431 pp. ISBN: 978-0-226-03058-6.
- Goetzmann, William N. and Jonathan G. S. Koppell (Oct. 1, 2008). “The Trickle-Up Bailout”. In: *The Washington Post*. ISSN: 0190-8286.
- Goodman, Laurie, Walt Scott, and Jun Zhu (July 2018). *How Beneficial Are Streamlined Modifications?* Research report. Urban Institute, p. 46.

- Guiso, Luigi, Paola Sapienza, and Luigi Zingales (Aug. 1, 2013). “The Determinants of Attitudes toward Strategic Default on Mortgages”. In: *The Journal of Finance* 68.4, pp. 1473–1515. ISSN: 1540-6261.
- Guren, Adam M and Timothy J McQuade (2018). “How Do Foreclosures Exacerbate Housing Downturns?” In: p. 66.
- Guren, Adam M., Arvind Krishnamurthy, and Timothy J. Mcquade (2021). “Mortgage Design in an Equilibrium Model of the Housing Market”. In: *The Journal of Finance* 76.1, pp. 113–168. ISSN: 1540-6261.
- Gyourko, Joseph and Joseph Tracy (Mar. 1, 2014). “Reconciling Theory and Empirics on the Role of Unemployment in Mortgage Default”. In: *Journal of Urban Economics* 80, pp. 87–96. ISSN: 0094-1190.
- Haughwout, Andrew, Donghoon Lee, et al. (2011). *Real Estate Investors, the Leverage Cycle, and the Housing Market Crisis*. Staff Report, Federal Reserve Bank of New York.
- Haughwout, Andrew, Ebiere Okah, and Joseph Tracy (June 2016). “Second Chances: Subprime Mortgage Modification and Redefault”. In: *Journal of Money, Credit and Banking* 48.4, pp. 771–793. ISSN: 00222879.
- Haughwout, Andrew, Richard Peach, and Joseph Tracy (Sept. 2008). “Juvenile Delinquent Mortgages: Bad Credit or Bad Economy?” In: *Journal of Urban Economics* 64.2, pp. 246–257. ISSN: 00941190.
- Jagtiani, Julapa and William W. Lang (Dec. 9, 2010). *Strategic Default on First and Second Lien Mortgages During the Financial Crisis*. SSRN scholarly paper. Rochester, NY: Social Science Research Network.
- Kalikman, Philip Lewis and Joelle Scally (June 7, 2023). *Mortgage Default: A Heterogeneous-Agent Model*. preprint.
- Kaplan, Greg, Kurt Mitman, and Giovanni L Violante (Aug. 5, 2020). “The Housing Boom and Bust: Model Meets Evidence”. In: *Journal of Political Economy* 128.9.
- Kau, James, Donald Keenan, and Kim Taewon (1994). “Default Probabilities for Mortgages”. In: *Journal of Urban Economics* 35.3, pp. 278–296. ISSN: 0094-1190.
- Kau, James B and Donald C Keenan (1995). “An Overview of the Option-Theoretic Pricing of Mortgages”. In: *Journal of Housing Research*, p. 29.
- Kau, James B. (1992). “A Generalized Valuation Model for Fixed-Rate Residential Mortgages”. In: *Journal of Money, Credit and Banking* 24.3, pp. 279–99.
- Laufer, Steven (Apr. 2018). “Equity Extraction and Mortgage Default”. In: *Review of Economic Dynamics* 28, pp. 1–33. ISSN: 10942025.
- Li, Wei and Laurie Goodman (Nov. 2014). *Measuring Mortgage Credit Availability Using Ex-Ante Probability of Default*. Urban Institute, p. 35.

- Mae, Fannie (n.d.). “Fannie Mae Statistical Summary Tables”. In: (), p. 5.
- McCoy, Patricia A (2013). “Barriers to Foreclosure Prevention During The Financial Crisis”. In: *ARIZONA LAW REVIEW* 55, p. 51.
- Mian, Atif, Amir Sufi, and Francesco Trebbi (Dec. 2010). “The Political Economy of the US Mortgage Default Crisis”. In: *American Economic Review* 100.5, pp. 1967–1998. ISSN: 0002-8282.
- Pearl, Judea (Mar. 4, 2022). “Probabilities of Causation: Three Counterfactual Interpretations and Their Identification”. In: *Probabilistic and Causal Inference: The Works of Judea Pearl*. 1st ed. Vol. 36. New York, NY, USA: Association for Computing Machinery, pp. 317–372. ISBN: 978-1-4503-9586-1.
- Pearl, Judea and Dana Mackenzie (May 15, 2018). *The Book of Why: The New Science of Cause and Effect*. Basic Books. 470 pp. ISBN: 978-0-465-09761-6. Google Books: [9H0dDQAAQBAJ](#).
- Powell, Jerome, ed. (June 22, 2021). *Lessons Learned: The Federal Reserve’s Response to the Coronavirus Pandemic*. Washington, D.C.
- Scharlemann, Therese C. and Stephen H. Shore (Oct. 1, 2016). “The Effect of Negative Equity on Mortgage Default: Evidence From HAMP’s Principal Reduction Alternative”. In: *The Review of Financial Studies* 29.10, pp. 2850–2883. ISSN: 0893-9454.
- Schelkle, Thomas (Sept. 2018). “Mortgage Default during the U.S. Mortgage Crisis”. In: *Journal of Money, Credit and Banking* 50.6, pp. 1101–1137. ISSN: 00222879.
- Scott, Walter (Feb. 27, 2015). *Comparing the Performance of Home Affordable Modification Program (HAMP) Modifications and Non-HAMP Modifications: Early Results*. Fannie Mae.
- Shiller, Robert (Oct. 2007). *Understanding Recent Trends in House Prices and Home Ownership*. w13553. Cambridge, MA: National Bureau of Economic Research.
- Shiller, Robert J. et al. (2017). “Continuous Workout Mortgages: Efficient Pricing and Systemic Implications”. In: *Journal of Economic Behavior & Organization*.
- Table 3. *Median Usual Weekly Earnings by Age, Race, Hispanic or Latino Ethnicity, and Sex* (2021). URL: <https://www.bls.gov/webapps/legacy/cpswktab3.htm> (visited on 10/25/2021).
- U.S. Department of the Treasury (May 10, 2013). *Making Home Affordable Program Performance Report Through March 2013*. U.S. Department of the Treasury.
- (Apr. 2012). *The Financial Crisis Response In Charts*. U.S. Department of the Treasury.
- Vandell, Kerry D. (1995). “How Ruthless Is Mortgage Default? A Review and Synthesis of the Evidence”. In: *Journal of Housing Research* 6.2, pp. 245–264. ISSN: 1052-7001. JSTOR: [24832828](#).

- WatchBlog (Oct. 3, 2018). *As TARP Turns 10, Looking Back at the U.S. Government's Assistance to the Banking Sector*. URL: <https://blog.gao.gov/2018/10/03/as-tarp-turns-10-looking-back-at-the-u-s-governments-assistance-to-the-banking-sector/> (visited on 03/15/2019).
- White, Brent T. (2010). "Underwater and Not Walking Away: Shame, Fear and the Social Management of the Housing Crisis". In: *SSRN Electronic Journal*. ISSN: 1556-5068.
- Wickham, Hadley (2016). *Ggplot2: Elegant Graphics for Data Analysis (3e)*. Springer-Verlag. ISBN: 978-3-319-24277-4.
- Wickham, Hadley et al. (Nov. 21, 2019). "Welcome to the Tidyverse". In: *Journal of Open Source Software* 4.43, p. 1686. ISSN: 2475-9066.
- Yellen, Janet (Oct. 14, 2016). "Macroeconomic Research After the Crisis". *The Elusive 'Great' Recovery: Causes and Implications for Future Business Cycle Dynamics* (Boston, MA).
- Zame, William R (1993). "Efficiency and the Role of Default When Security Markets Are Incomplete". In: *American Economic Review* 83.5, pp. 1142–1164. JSTOR: [2117553](https://www.jstor.org/stable/2117553).
- Zingales, Luigi (Oct. 2008). *Plan B*. URL: https://faculty.chicagobooth.edu/luigi.zingales/papers/research/plan_b.pdf (visited on 12/05/2019).

Figures and Tables




Figure 1: House Prices, United States and 20 Metro Areas, 2000 to 2020




Figure 2: Cumulative Default Rates, Empirical versus Model In-sample and Model Out-of-sample




Figure 3: Cumulative Default Rates: Baseline versus Writedown to 95% CLTV

Attribute Date Observed	All Borrowers January 2007	Borrowers who Default January 2007
Original Purchase Price	\$221,000 (\$125,000)	\$226,000 (\$124,000)
Original Loan Size	\$186,000 (\$98,000)	\$196,000 (\$98,000)
Original First-lien LTV	86.2% (8.8%)	88.8% (8.9%)
Original First-lien PTI	33.16% (10.39%)	35.7% (10.0%)
Original First-lien Annual Payment	\$13,700 (\$7,200)	\$14,700 (\$7,300)
Original First-lien Interest Rate	6.21% (0.57%)	6.36% (0.64%)
Original Credit Score	717 (58)	686 (61)
Current Credit Score	714 (73)	662 (91)
Current First-lien LTV	82.8% (9.5%)	85.8% (9.8%)
Current Second Lien Balance (when positive)	\$37,000 (\$28,000)	\$42,000 (\$31,000)
Current Combined LTV	90.6% (8.3%)	92.4% (8.6%)
Current Total Annual Payment	\$15,100 (\$8,000)	\$15,900 (\$8,200)
Current Income	\$49,500 (\$65,800)	\$47,700 (\$62,800)
Current Net Liquid Assets	-\$18,000 (\$31,000)	-\$18,400 (\$30,000)
House Price Appreciation Log-Mean	0.056 (0.011)	0.057 (0.011)
House Price Appreciation Log-Volatility	0.058 (0.018)	0.059 (0.017)
Months until Default (Defaulters)		43.7 (26.8)
Distribution by Status as of 2021		
Active	76.3%	
Default	19.3%	100%
Paid	4.4%	
Distribution by Origination Vintage		
Before 2001	<1%	<1%
2001	<1%	<1%
2002	2%	2%
2003	5%	4%
2004	10%	8%
2005	36%	35%
2006	46%	51%

Table 1: Summary Statistics, 30-year Fixed-Rate Loans with 80%+ LTV in January 2007

Group	Parameter	Model Symbol	Value or Data Source
Prices and Rates for Borrower i:			
	House Prices	$\{P_t(i)\}_{t=0,\dots,T}$	McDash + CoreLogic HPI
	House Price Log-Drifts	$\{\mu_t(i)\}_{t=0,\dots,T-1}$	Zip-level CoreLogic HPI
	House Price Log-Volatilities	$\{\sigma_t(i)\}_{t=0,\dots,T-1}$	Zip-level CoreLogic HPI
	Risk-free Savings Rates	$\{R_t^f\}_{t=0,\dots,T-1}$	1-year Treasuries
	Unsecured Borrowing Rates	$\{R_t^b\}_{t=0,\dots,T-1}$	FRED Credit Card APRs
	Inflation Index	$\{I_t\}_{t=0,\dots,T-1}$	1.01^t
Mortgage of Borrower i:			
	Term	$T(i)$	McDash
	Combined Initial Mortgage Principal	$M_0(i)$	Equifax CRIS + McDash
	Mortgage Interest Rate Series	$\{r_t^m(i)\}_{t=0,\dots,T-1}$	McDash
	Mortgage Coupon Schedule	$\{m_t(i)\}_{t=1,\dots,T}$	McDash
	Mortgage Principal Schedule	$\{M_t(i)\}_{t=0,\dots,T}$	McDash
	Modified Term	$T'(i)$	Experiment-dependent
	Lender Option Price Floor	$\{Q_t^f(i)\}_{t=0,\dots,T'}$	Experiment-dependent
	Lender Option Time Limit	$\{Q_t(i)\}_{t=0,\dots,T'}$	Experiment-dependent
	Lender Option Sharing Percent	$\{Q_t^p(i)\}_{t=0,\dots,T'}$	Experiment-dependent
Other Characteristics of Borrower i:			
	Idiosyncratic Default Penalty	$\lambda(i)$	Estimated (see Sec. 4.4)
	Income	$L_t^h(i)$	McDash & BLS
	Starting Liquid Assets	$A_0(i)$	McDash/CRIS/SCF Estimate
	Probability of Unemployment	π_u	0.07
	Probability of Re-employment	π_e	0.35
	CRRA Risk-aversion parameter	γ	3.5
	Discount factor	β	0.985
Transaction Costs, Taxes, Fees, etc.:			
	Income tax rate	τ_{inc}	0.20
	Property tax rate	τ_{prop}	0.015
	Property maintenance/insurance/HOA	τ_{maint}	0.01
	Transaction cost of selling	τ_{sell}	0.05
	Transaction cost of prepaying	τ_{prepay}	0.015
	Transaction cost of moving	τ_{move}	0.005
	Rent-to-price ratio	τ_{rent}	0.04
	Distressed sale ratio	$\tau_{\text{distressed}}$	0.65
	Unemployment income replacement ratio	τ_{unemp}	0.60

Table 2: Model Parameters, Data Sources, and Baseline Values.

Attribute Date Observed	All Borrowers January 2007	Borrowers with 130% SLTV Event Date of First 130% SLTV Event
Original Purchase Price	\$221,000 (\$125,000)	\$251,000 (\$121,000)
Original Loan Size	\$186,000 (\$98,000)	\$217,000 (\$97,000)
Original First-lien LTV	86.2% (8.8%)	88.3% (9.71%)
Original First-lien PTI	33.16% (10.39%)	38.6% (12.44%)
Original First-lien Annual Payment	\$13,700 (\$7,200)	\$16,300 (\$7,300)
Original First-lien Interest Rate	6.21% (0.57%)	6.42% (0.59%)
Original Credit Score	717 (58)	710 (59)
Current Credit Score	714 (73)	695 (110)
Current First-lien LTV	82.8% (9.5%)	131.8% (1.68%)
Current Second Lien Balance (when positive)	\$37,000 (\$28,000)	\$39,000 (\$35,000)
Current Combined LTV	90.6% (8.3%)	136.6% (12.8%)
Current Total Annual Payment	\$15,100 (\$8,000)	\$16,800 (\$7,700)
Current Income	\$49,500 (\$65,800)	\$51,200 (\$64,800)
Current Net Liquid Assets	-\$18,000 (\$31,000)	-\$22,216 (\$37,245)
House Price Appreciation Log-Mean	0.056 (0.011)	0.056 (0.011)
House Price Appreciation Log-Volatility	0.058 (0.018)	0.060 (0.014)
Months until Default (Defaulters)	43.7 (26.8)	17.7 (19.2)
Distribution by Status as of 2021		
Active	76%	56%
Default	19%	42%
Paid	4%	2%
Distribution by Origination Vintage		
Before 2001	<1%	<1%
2001	<1%	<1%
2002	2%	<1%
2003	5%	1%
2004	10%	3%
2005	36%	18%
2006	46%	28%
2007		41%
2008		8%

Table 3: Borrower characteristics: overall and those who experience 130% Scheduled LTV

HPA	Credit Score	Default (%)	Post-Mod Def. (%)	Averted (%)	Waste (%)	Avg Cost (\$)
High	High	25	3	96	72	51,000
High	Low	48	5	89	49	4,600
Low	High	50	3	93	44	-2,800
Low	Low	74	14	81	22	-29,800

Table 4: Summary of Writedown Effectiveness and Cost by Target Population

Credit Score Group	Pre-Mod Default %	Post-Mod Default %	Average Cost (\$)
Bottom 50%	62	10	-14,000
Bottom 40%	65	12	-18,000
Bottom 30%	68	14	-20,000
Bottom 20%	71	17	-20,000
Bottom 10%	75	22	-21,000

Table 5: Summary of Writedown Effectiveness and Cost by Credit Score Criterion

Appendices

A Model and Methods Details

A.1 Trinomial Tree

There are $2t + 1$ possible exogenous values of the house price index in period t . Denoting the time by the superscript and the discrete realizable values by the subscript,

$$P^t = \{P_{-t}^t, \dots, P_{-1}^t, P_0^t, P_1^t, \dots, P_t^t\}.$$

The levels of the prices in the grid are determined by the parameters initial price $P^0(i)$, log-drifts $\{\mu_t(i)\}_{t=0, \dots, T}$, log-volatilities $\{\sigma_t(i)\}_{t=0, \dots, T}$, and grid parameters $\{\hat{\sigma}_t(i)\}_{t=0, \dots, T}$. Each loan (i) has its own instances of these parameters, though we suppress the notation hereafter.

These parameters determine both the level of prices at nodes in the grid and the transition probabilities between those nodes. The j^{th} level of the house price index P_j^t is

$$P_j^t = P_0 \cdot \exp \left(\sum_{\tau=0}^{t-1} \mu_\tau + \hat{\sigma}_t \cdot j \right). \quad (1)$$

The transition probabilities between prices in subsequent periods, $Pr(P_{j+1}^{t+1}|P_j^t)$, $Pr(P_j^{t+1}|P_j^t)$, and $Pr(P_{j-1}^{t+1}|P_j^t)$, are determined by the moment conditions

$$E_{j'=j-1, j, j+1} \left[\log(P_{j'}^{t+1}) - \log(P_j^t) \right] = \mu_t \quad (2)$$

$$Var_{j'=j-1, j, j+1} \left[\log(P_{j'}^{t+1}) - \log(P_j^t) \right] = \sigma_t^2 \quad (3)$$

as well as the usual constraints on probability distributions:

$$\begin{aligned} 0 &\leq Pr(P_{j+1}^{t+1}|P_j^t) \leq 1 \\ Pr(P_{j+1}^{t+1}|P_j^t) + Pr(P_j^{t+1}|P_j^t) + Pr(P_{j-1}^{t+1}|P_j^t) &= 1. \end{aligned}$$

The grid parameters $\{\hat{\sigma}_t\}_{t=0, \dots, T}$ are chosen as the minimum values feasible to satisfy Equations (1), (2), and (3).

House prices during the Global Financial Crisis followed a path that many would regard as

inconsistent with the expectations of borrowers before 2008. Our model thus supports time-varying expectations with “surprises”; borrowers are rational but have imperfect foresight. That is, a borrower with a mortgage at period 0 may have a forecast of log-drifts and log-volatilities as of period 0, $\{\mu_0^0, \mu_1^0, \mu_2^0, \dots\}$ and $\{\sigma_0^0, \sigma_1^0, \sigma_2^0, \dots\}$, that imply an expected house price path $E_0[P_t], E_0[P_{t+1}], \dots$, which does not lie within the grid of prices $\{P_{-1}^1, P_0^1, P_1^1\}$ that the borrower forecasts. We account for this possibility by computing additional “out-of-grid” price levels determined by the actual price history the borrower would experience and simulating Monte Carlo histories across price paths with jumps or surprises.

If the borrower sells the home, she receives the sale price P less transaction costs $\tau_{\text{sell}} \cdot P$ and $\tau_{\text{move}} \cdot P$ for searching for and moving to a new rental home. If she defaults, she pays the moving cost only but lives rent-free for a predetermined period reflecting the time it takes the lender to foreclose. Either way, she thereafter rents a comparable property. We take rent to be proportional to price ($\tau_{\text{rent}} P_t^s$), up to a ceiling and floor so that rental costs do not get unrealistically high or low.

A.2 Budget Constraint in Detail

$a_t^s =$	End of period cash-on-hand	(4)
A_t^s	Starting assets	(5)
$+ L_t^s \cdot (1 - \tau_{\text{inc}})$	Labor income less income tax	(6)
$- \mathbb{1}_{\text{Has Mortgage in } s} \cdot m_t$	Mortgage coupon payment	(7)
$+ \mathbb{1}_{\text{Has Mortgage in } s} \cdot \tau_{\text{inc}} r_{t-1}^m M_{t-1}$	Mortgage interest deduction	(8)
$- \mathbb{1}_{\text{Owns Home in } s} \cdot ((\tau_{\text{maint}} + \tau_{\text{prop}}) P_t^s)$	Maintenance and property taxes	(9)
$+ \mathbb{1}_{\text{Owns Home in } s} \cdot (\tau_{\text{inc}} \tau_{\text{prop}} P_t^s)$	Property tax deduction	(10)
$- \mathbb{1}_{\text{Rents Home in } s} \cdot \mathbb{1}_{\text{Did not recently default}} \cdot (\tau_{\text{rent}} P_t^s)$	Rent	(11)
$- I_t C_t^s$	Inflated price of consumption	(12)
$- \mathbb{1}_{\text{Prepays Mortgage in } s} \cdot (1 + \tau_{\text{prepay}}) \cdot M_t$	Principal plus prepayment fees	(13)
$+ \mathbb{1}_{\text{Sells Home in } s} \cdot (1 - \tau_{\text{move}} - \tau_{\text{sell}}) \cdot P_t^s$	Price less moving costs, broker's fees	(14)
$- \mathbb{1}_{\text{Sells Home in } s} \cdot \mathbb{1}_{t < Q^t} \cdot Q^p \cdot \max(P_t^s - Q_f^t, 0)$	(See Sec. 3.2.1)	(15)
$- \mathbb{1}_{\text{Defaults in } s} \cdot \tau_{\text{move}} \cdot P_t^s$	Moving costs	(16)
$+ \mathbb{1}_{\text{Defaults in } s} \cdot \max(\tau_{\text{distressed}} \cdot P_t^s - M_t, 0)$	Distressed sale price less principal	(17)
$- \mathbb{1}_{\text{Defaults in } s} \cdot \mathbb{1}_{\text{Recourse}} \cdot (M_t - \tau_d P_t^s)$	Principal less distressed sale price	(18)

The agent will conclude the period with a level of liquid assets (4), that she will invest until the next period if positive, or borrow if negative, determined as follows. She begins the period with a given level of liquid assets (5), and receives labor income net of taxes (6). If she still has a mortgage, she pays the coupon (7), but receives a deduction from income

tax for her mortgage interest (8). Whether she has the mortgage or not, provided she owns the home, she pays maintenance and property taxes (9), but deducts property tax from income tax (10). If instead she rents her home, she pays rent unless she recently defaulted, in which case she lives rent-free for a fixed period (typically two years) (11). These costs are all determined by her initial state. Subsequently, she makes several choices: first, a level of consumption (12), and then decisions regarding her mortgage. If she prepays the mortgage, she pays the outstanding balance to the lender, plus a fee proportional to that balance (13). If she prepays or has previously prepaid, she may sell the home. She receives the home price less proportional broker’s fees and moving costs (14). If she sells the home after receiving a modification with an equity share agreement, she may have to pay the lender its share of equity. (See the more detailed discussion in Section 3.2.1.) Finally, if she defaults, she must pay the same moving costs she would after selling (16). The home is sold in a distressed sale. If the distressed sale price covers her outstanding principal, the lender receives the principal and she receives the difference (17); if not, and if the lender has recourse, the borrower pays the shortfall out of liquid assets (18).

A.2.1 Irreversibility

The model features standard assumptions on irreversibility or absorbing states for computational tractability:

- A borrower who leaves her house, whether through sale or default, moves into a house of the same “size” or “quality.” Borrowers do not adjust house size or quality. Accordingly, the term for housing in the utility function is suppressed.
- A borrower who terminates a mortgage, whether due to prepayment, default, or completion, does not procure a new mortgage. In particular, while there is prepayment, there is no refinance and no equity extraction or other unscheduled adjustment of equity other than complete prepayment. Mortgages either amortize according to the amortization schedule determined at origination or modification, or they are fully prepaid.
- A borrower who rents does so until the terminal period. Borrowers do not ever purchase new homes, even if they could afford to do so entirely out of cash-on-hand.

The plausibility of these assumptions has been widely discussed in the related literature. The irreversibility of these decisions is not a strictly accurate assumption, and it limits

the relevance of resulting predictions in general equilibrium. However, the inaccuracy is mitigated to some extent by the true frictions that befall defaulters in mortgage markets. In addition to reducing defaulters' credit scores, default events generally stay on a credit report for seven years; both the lowered score and the record of the default limit access to future credit that would be necessary for most such borrowers to purchase another home. This model does allow for prepayment, including prepayment without sale. Borrowers may therefore capture the value of future home price appreciation either by reducing liquid assets in one fell swoop, or by continuing to pay the mortgage coupon, though for most borrowers, prepaying without selling the home would require them to take out expensive unsecured debt.

A.3 Methods

A.3.1 Further Sampling Restrictions

We restrict to properties in ZIP codes in which the CoreLogic home price index is observed throughout the model period. The CoreLogic home price data cover approximately 7,500 ZIP codes. ZIP codes where there are few transactions do not have price indices computed. As transaction volume correlates with population, restricting to ZIP codes with defined home price indices likely does not break the generality of the results.

We restrict to mortgages that McDash indicates do not admit lender recourse. The literature is divided on the effects of recourse on a borrower's decision to default: on the one hand, Ghent and Kudlyak (2011) find that it reduces a borrower's default propensity; on the other, Guiso, Sapienza, and Zingales (2013) find that borrowers do not have a clear understanding of whether lenders have recourse to pursue their assets, let alone reliable estimates of lenders' likelihood of doing so or of succeeding. It is not clear how recourse would affect our estimates, as while borrowers who had expectations of lender recourse would be less likely to default *ex ante*, they would also be less likely to re-default after modification.

Equifax matches tradeline credit data to mortgage loans in McDash using a proprietary probabilistic matching algorithm. We restrict our analysis to loans with high match confidence; we further ensure correct matching by selecting borrowers whose first mortgage balance in the model reference period as reported in McDash is within 5% of the balance reported in Equifax. The analysis herein is not sensitive to including loans that are more or less confidently matched. We exclude loans that transferred servicers during the period of interest.

A.3.2 Definition of Default

Mortgage servicers do not adhere to a standard definition of the period in which a borrower defaults, and thus neither the Equifax CRIS nor the McDash data indicate a date of default. The major challenge to identifying defaults is that borrowers often fall behind on payments and then make up missed payments, or “cure,” and then may stay current until prepayment or completion of the loan—or may fall behind again and default. Additionally, servicers vary in when they initiate foreclosure proceedings, and foreclosure can take varying amounts of time. Fortunately, because this is a historical study, the typical challenge of identifying a true default is mitigated. We define defaults in the data according to the *final* status of the mortgage. Because we are using mortgage history data from 2021, we observe whether the mortgage prepays, defaults, or is modified at any point after a given period. Their terminal state may be one of the following: paid off or prepaid, foreclosed upon and liquidated, transferred to a servicer that did not continue report to McDash, modified, or still active in 2021.

In the model, borrowers do not miss payments and then cure. Whenever a borrower in the model misses a payment, she will miss all future payments. The borrower will remain in the home for two years while the lender completes foreclosure proceedings, but there is no opportunity for the borrower to change her mind or access additional liquidity to make up a past missed payment. This feature is more realistic than it may appear: it means the focus of our model is only terminal defaults, not delinquencies-with-cures, and we identify only terminal defaults in the data. The model restriction and the data are consistent.

Again excluding mortgages that transferred servicers, we define mortgages that were paid or prepaid by 2021, or that are still active in 2021, as not having defaulted, even if the borrower temporarily fell behind on payments. Excluding servicer transfers may bias the sample towards safer loans but is necessary in order not to mis-classify delinquent mortgages that eventually cure after a transfer as defaults. Mortgages that were foreclosed upon before 2021 did default. Additionally, if a mortgage was delinquent and then received a modification, We define it as having defaulted. The reason for this choice is that in the study of defaults, the most logical assumption is that a delinquent mortgage that received a modification received the modification precisely because the modifying party (typically the mortgage servicer) considered modifying the mortgage more profitable than the alternative of initiating foreclosure proceedings. But for this to be the case, the servicer must have expected the borrower not to cure and continue making payments on the original mortgage. We therefore treat such outcomes as defaults; they are equivalent to a default on the original mortgage and an origination of a new mortgage.

For mortgages that default, defining the *date* of default presents its own challenges. Because default in the model refers to the first point after which the borrower decides it is preferable to stop making payments and ceases making any payments, we take defaults in the data also to be the first point after which the borrower fully ceases making payments. Therefore, if a borrower becomes delinquent but cures, and subsequently becomes delinquent again but does not cure, her default date is the month before the second delinquency. The cure event is ignored.

Consider the following examples for three hypothetical loans, where e.g. 30DPD refers to a loan that is 30-59 days past due.

[Table 6 about here.]

Borrower 1 did not default. Borrower 2 defaulted in period 4: her earlier delinquency is ignored because she cured, whereas as of period 4 she began a series of missed payments that culminated in a liquidation. Borrower 3 also defaulted in period 4: she decided to miss her first payment in the following period, and never to make a subsequent payment. Borrower 4 is excluded from the sample. Because we do not observe her loan performance after the period in which her mortgage was transferred to a different servicer that did not report to McDash, we do not know whether she eventually cured her mortgage.

A.3.3 Definition of Modification

The McDash data include partial coverage of loss mitigation data provided by servicers, consisting primarily of flags indicating months in which a mortgage receives a modification, whether that modification was conducted as part of a program such as HAMP, and whether that modification involved various features such as principal forgiveness, interest rate reduction, or term extension. These flags sometimes do not coincide with actual observed changes in mortgage variables. To account for such discrepancies, we also follow the modification-identification procedure outlined in Goodman, Scott, and Zhu (2018). The authors identify a loan as receiving a modification when it transitions from 60 or more days delinquent to current while simultaneously increasing term, changing interest rate by more than 10 basis points, or changing principal balance or principal-plus-interest payment by more than 3 percent. We take the earliest date of such a coincident change (in both mortgage status and mortgage variables) after the reference period as the first date of modification of a loan.

A.3.4 Estimation of Combined Loan-to-Value Ratio

The McDash data do not have property-level identifiers and have low second-lien coverage. They thus do not provide a reliable source of information regarding the *combined* loan-to-value ratio on a mortgage, the total loan-to-value ratio after accounting for junior liens. Observing the correct combined LTV (CLTV) is crucial to investigating default behavior. A borrower who owes 90% of the value of a home on her first mortgage could sell the home if she found payments unaffordable—unless she had a second mortgage on which she owed 20% of the value of the home. Using only first-lien LTV to infer positive equity would therefore confuse the interpretation of the decision by such a borrower to default. The effect of missing junior liens is just as detrimental when the borrower is already in negative equity even just accounting for the first lien: the borrower’s own estimate of the probability of house price changes restoring her to positive equity may be substantially lower than what one would infer from the data without accounting for the junior liens she had.

Accordingly, following the method also used by Crews Cutts and Merrill (2008), we adjust for combined LTV on loans using the Equifax CRIS data. We have already restricted to borrowers who have only one first mortgage, and we further restrict to those who have only one or zero closed-end second (CES) mortgages for the nine months following the reference period, and one or zero Home Equity Lines of Credit (HELOCs) at the same time. This restriction increases the likelihood that any CES and HELOC loans in the borrower’s credit report are on the same property as the first mortgage, though it is not logically impossible that a borrower could have a junior lien on a property other than that backing her senior lien. We further restrict to those borrowers for whom the first mortgage balance reported in Equifax CRIS differs by not more than 5% from that reported in McDash. (Because of timing mismatches, the two sources may report different outstanding principal balances on what is actually the same mortgage.) For these borrowers, we add the total outstanding balance on the CES and HELOC loans as observed in CRIS to the first mortgage balance observed in CRIS in order to compute the CLTV as a multiple of the first-lien LTV on the property. We restrict to properties for which this multiple is less than 1.8, as junior liens are typically smaller in volume than senior liens, so multiples higher than 1.8 are likely due to data inaccuracies or unusual circumstances. We then compute total principal outstanding on the property as the CRIS-derived multiple times the McDash-reported first-lien principal outstanding (first-lien LTV). We use the McDash principal rather than the Equifax CRIS principal because the rest of the loan performance characteristics, including the interest rate, loan status and therefore inferred date of default, and so on, are reported in McDash but not Equifax. Accordingly, the modeled borrower makes her default decision based on the

timing of data reported in McDash.

Unlike first liens, second liens do not have origination or performance characteristics reported in McDash. As a result, while we observe the total outstanding balance on such liens via the Equifax CRIS data, we do not know the interest rate or product structure for these second liens. As is consistent with the literature, we approximate the amortization schedule of these loans by assuming they are capitalized into the first mortgage balance and amortize at the same rate and term, rather than treating them separately. Accordingly, the borrower who defaults in the data is assumed to default on all mortgages backing a property simultaneously; they are treated as one large mortgage in the model.

A.3.5 Estimation of Income

We use McDash-reported front-end payment-to-income ratio to calculate the borrower's annual income. Documentation of income is known to be limited, particularly before the financial crisis, so this may overstate income estimates in many cases. An investigation of an additional dataset with borrower-level income and employment data revealed that it matched too few borrowers who default during the sample period to offer a viable alternative to this method. This dataset has better coverage for recent years and could prove useful in extending the analysis to more recent and/or future defaults.

A.3.6 Estimation of Liquid Assets

We model cash on hand for the borrower as a predicted level of liquid assets net of the level of non-mortgage debt as observed in CRIS, where the prediction is formed from the 2007-2009 Survey of Consumer Finances data by regressing log liquid asset levels against log income, log mortgage debt, and log non-mortgage debt for respondents with values of each of these variables in the middle 90% of respondents. We define mortgage debt as the sum of mortgage and HELOC debt, non-mortgage debt as the sum of credit card debt, auto debt, education debt, and other debt, and assets as the sum of checking and savings accounts, mutual fund holdings, savings bonds, stocks, and other significant assets. As indicated in Table 7, log levels of income, mortgage debt, and non-mortgage debt are all statistically and economically significant predictors of log liquid assets, with an R^2 of 41%.

[Table 7 about here.]

A.3.7 Estimation of House Price Expectations and Realizations

Each borrower forms an expectation of the mean and volatility of house price surprises by taking the historical average in her ZIP code from 1975 until 2002. Her expectations are independent of other borrowers' expectations, are uncorrelated with labor income surprises and other exogenous factors. Expectations may be formed according to several specifications. In the simplest, they are fixed, and do not interact with the path of prices taken in computing scenario-specific expected default rates, lender recoveries, or other statistics of modification effectiveness. In this specification, a borrower who lives in a ZIP code with high home price appreciation through 2002, who originates a mortgage in 2007, and then experiences a drastic decline in home prices, will continue to expect high prices. We also study a specification in which borrowers revise expectations after price surprises (positive or negative).

The expected default rate and lender recoveries are computed conditional on the realization of prices for the given ZIP code from the reference period up until the last period of interest. In studying modifications, we take the actual realization of house price changes as the baseline: borrowers who receive a modification face the same home prices that they would have without a modification. In reality an effective modification program might reduce defaults enough to feed back into prices and reverse or accelerate the reverse of price declines; We do not account for that feedback effect. We consider alternate specifications for home price realizations in Appendix B.1.

A.3.8 Simulation

To compute default probabilities and expected lender receipts for a pool of loans \mathcal{L} , we first solve the model by backwards recursion for each loan in the pool. We then draw 25,000 forward paths, i.e. realizations of house price and labor income draws across time, for *each* loan. Typically, we condition on a given starting path of house prices, so that only labor income is uncertain along the beginning of each path. Conditional on a loan being active in a given period t , the borrower's state in $t + 1$ is determined by the borrower's decision regarding the mortgage and regarding assets in t : a borrower who chooses to continue paying the mortgage will have the mortgage in $t + 1$ with an asset level determined by the savings chosen at t according to the budget constraint in Section 3.8. Note that in forward paths, prices need not follow edges along the price tree. In this way we are able to model "surprises," whereby the price series conditioned upon in computing a statistic such as default probability or lender receipts is not anticipated by the borrower in the reference period. We thereby compute for each loan its expected probability of default by a given period in the future

as well as the expected dollars paid to the lender over the life of the loan as the sum of these figures (in the case of default probability, 1 for defaults and 0 otherwise) period-by-period divided by the number of paths. We then compute pool-level figures by summing over the individual figures computed for each loan in the pool. Loans in a pool are each subject to a given house price path determined by the relevant experimental specification, but there is no interaction between loans in the pool. That is, there is no reason that a given loan defaulting on a given path should make it more likely that another randomly selected loan on another randomly selected path will default. In reality, surprises that result in one mortgage’s default likely are at least partially correlated with others that would lead to other defaults. Omitting this correlation is therefore likely to bias estimates of cost-effectiveness of a mitigation program downward.

B Limitations and Alternate Scenarios

B.1 Alternate Prices

As we show in Figure 1, the recovery in house prices from the crisis has been robust since the waning years of decline, around 2013. But in 2009, there was no way to know when prices would rebound, if ever. Policy has to operate with uncertainty over the future.

Had prices continued to fall, the effectiveness of principal forgiveness would likely have been even greater, even ignoring general equilibrium effects. And the general equilibrium effects of principal forgiveness would further enhance estimates of its effectiveness, as it would have been likely to turn around the economic decline more quickly.

To get at these impacts, we model two scenarios for realized prices. In the baseline scenario, prices follow their observed path. We also study a “protracted crash” scenario. Principal forgiveness is likely to have been especially effective at encouraging borrower who would default when prices continued to drop to pay instead. We can study this possibility by simulating a worse price realization to identify whether principal forgiveness is indeed more effective in such a scenario. In the protracted crash scenario, we assume prices in each ZIP code plunged an additional two standard deviations, and recovered to their actual values after four years.

While it is not possible to know what else might have happened to house prices, it is partially possible to know how a modification program such as the one studied here would have fared conditional on prices following a different path. We hypothesize that had prices

fallen even further than they actually did, the effectiveness of principal forgiveness would have increased. The mechanism is straightforward: more borrowers would default if prices fell further, and principal forgiveness would provide the right cushion to these borrowers.

Indeed, the hypothesis is confirmed by the experiment. If prices had fallen by an additional two standard deviations in each zip code, and taken four years to recover to their observed levels, the writedown modification program would prevent a slightly smaller fraction of a larger number of defaults—84% instead of 85% avoided—out of about 59% rather than 50% of borrowers who default. But because there are more defaults, the modification wastes less money, giving away 40 instead of 49 cents on the dollar to borrowers who would not have defaulted. The net cost of the modification is negative. Lenders earn on average \$12,100 per modified loan. (The cost per avoided default is likewise negative, though not meaningful as increasing the number of avoided defaults would increase the cost per avoided default.) The comparison is summarized in Table 8.

[Table 8 about here.]

B.2 Income Uncertainty

If one analyzes modifications under a traditional default framework, in which borrowers who default are either double-trigger or strategic defaulters, then one might be tempted to conclude that forgiveness would only work for strategic defaulters, and in particular that only borrowers who maintained income or employment over the short term would continue making payments. In the companion paper Kalikman and Scally, 2023, we argue that this view of default is limited. In fact, this reasoning is just such an example of the limitation: borrowers means of paying are not limited to cash flows from income; what matters is their *willingness* to tap other means of payment.

Nonetheless, as we do not observe the path of borrowers' income, it is worth considering whether our results are sensitive to the assumption that borrower income evolves according to a Markov process with fixed transition probabilities. Because these transition probabilities are estimated unconditional on whether borrowers defaulted, they likely overstate the probability that a borrower maintained her pre-default income. Thus the results likely understate the effectiveness of the modification policy: if modified borrowers were less likely than we assume to be able to afford their payments, then they were more likely to default than we estimate. The dollars wasted on borrowers who would have paid anyway are reduced. It is, of course, also possible that some borrowers would have been so destitute that even after modification they still would have defaulted. A back-of-the-envelope calculation suggests

the pool of such borrowers would be an order of magnitude smaller than of the borrowers whom the model mis-classified as likely to continue paying even without modification.

It would be feasible to size, or at least bound, these effects by a counterfactual study in which borrowers who receive forgiveness then suffer some number of years of certain unemployment. We do not expect the results would meaningfully alter the conclusions herein. We also reiterate that, due to general equilibrium effects, implementing a policy such as the one we advocate herein would be more likely to restore borrowers to employment than happened in reality, another reason our estimates of effectiveness are likely lower bounds.

B.3 Comparison with forbearance-only modifications

Another way to investigate the importance of liquidity versus strategy is to ask: what if borrowers received forbearance modifications rather than forgiveness—i.e., modifications which only reduced their present payments, to the same level achieved by a comparable modification with forgiveness, but without reducing the total balance owed lenders? We did not set out to determine the *optimal* modification policy. It would likely be feasible to do so using our model, but our goal is not to pontificate with twenty-twenty hindsight. It is to ask a concrete counterfactual question about ideas that policymakers actually considered at the time, that were advocated broadly in academic and policy circles, and that could have been implemented *without* the benefit of a computationally intense model.

Nonetheless, it is reasonable to wonder how forgiveness would compare to forbearance. As we have already argued, *ex ante* we should expect such modifications to be less effective than principal-forgiving modifications. The reason remains the simple discontinuity between the presence of a default incentive for a borrower who has negative equity versus the absence of that very incentive when the borrower has positive equity. In forthcoming extensions, we may use the model to further investigate this question, particularly in the context of the more recent Covid crisis and the ensuing forbearance. Our hunch is that in this case, forbearance *was* the optimal policy.

But it bears observing that in the Global Financial Crisis, conversion of present principal to a balloon payment offers the lender little additional upside. We can conduct a back-of-the-envelope estimation: suppose a borrower purchased an originally \$200,000 house at 95% LTV and at 6% with a 30-year fixed-rate mortgage. After two years the borrower owes about \$186,500. Suppose the house price has fallen to \$143,000. Then it would take about \$50,500 to forgive the borrower to 95% LTV, and the lender would get 50% upside in future home price appreciation. By contrast, forbearing the same amount of principal would at most

(if the borrower completes payments) earn the lender \$50,500 in 28 years, discounted even at a lower rate of 5% worth a little shy of \$13,000 in present value. The crux of the issue is thus whether \$13,000 compensates the lender for both the borrower’s greater likelihood of default and the loss of potential home price appreciation. (Similarly, increasing the size of the balloon payment just increases the likelihood that the borrower re-defaults.) While the advantage of a given policy will depend on the planner’s particular utility function, we take that in the crisis, the planner placed a high emphasis on stopping the spiraling crisis in defaults and foreclosures. The trade-off does not appear stark enough definitely to favor forbearance in such a context; further research would be required.

B.4 Expectations Revisions

Shiller (2007) and Bordalo, Gennaioli, and Shleifer (2018) argue that borrowers revise expectations after price surprises, overweighting recent news. To reflect these significant house price surprises and the expectations revisions that borrowers likely undertake following such drastic surprises, for a particular realized path of house prices, we recompute the entire tree each period that the borrower retains her mortgage. If the borrower does not terminate the mortgage in period 0, her decision in period 0 determines her successor state in period 1. She forms a new forecast of future prices at this state in period 1, and therefore recomputes the entire remaining truncated tree, with the contemporaneous level of the house price determined exogenously but her asset level and other state variables determined endogenously. This modeling approach is computationally taxing but enables the model to capture two important realistic features: that the significant price drops in financial crises are not foreseen, and that such drops likely induce expectations revisions in borrowers who endure them.

With expectations revisions, borrowers default more quickly even when they have the same penalties, because negative surprises cause them to revise their expectations of future price appreciation downward. Cumulative default curves therefore exhibit steeper early rises. Other results are qualitatively similar.

C Who Benefits Lenders After Modification?

Under the baseline scenario, 44% of borrowers make lenders better off by being modified. Overall, borrowers for whom principal writedown is self-financing resemble the 130%-LTV population, but with some minor expected differences. Beneficial-to-target borrowers have lower credit scores, slightly higher original LTVs, slightly higher PTI ratios, and slightly

lower expectations of and volatilities for house prices. (Low volatility is bad for borrowers in the model because the mortgage embeds the default option, and volatility of the underlying increases option value. That is, the higher house price volatility, the higher the probability that prices eventually recover, improving upside without bound, while in the downside defaulting is always still possible and limits losses.)

On average, these borrowers are not particularly different in these characteristics from the borrowers who do not benefit lenders after modification. This analysis suggests that the only reliable way to reach a policy conclusion on modification is to base the conclusion on an analysis that makes no premature representativeness assumptions, but rather fully represents borrower heterogeneity up until the final relevant conclusion is computed, as under the veil of representativeness the two populations would incorrectly appear to perform similarly.

[Table 9 about here.]

D Barriers to modification

It is tempting to conclude that if principal forgiveness, or any modification, could really be Pareto-improving, then lenders would select the optimal modification themselves. Under this view, the low number of principal-forgiving modifications throughout the crisis would constitute evidence against any argument that modifications could be effective. While there is not consensus in the literature on the reasons for low modification rates, it is difficult to interpret the debate as validating that the market is efficient—but there were strong barriers to modification that provide good reason to doubt that lenders modified efficiently.

McCoy (2013) and Geanakoplos (2010a) describe several such barriers. The most important failure in the modification market is a misalignment between the incentives of the lenders—who own mortgages and therefore stand to benefit from modifying them profitably—and the servicers, the third party entities who collect payments from borrowers and were typically the only entities contractually enabled to implement a modification. Servicer incentives are structured to favor foreclosure, for which they can recoup costs, and against modification, which bears upfront costs in underwriting and negotiation (technical skills servicers also lacked) and then often only reduces their compensation even if effective.

Other incentive misalignment and coordination problems also limit the market’s ability to deliver effective modifications. Ownership of an individual mortgage was often divided among different investors, some of whom had different incentives from others, characterized by the popular term “tranche warfare.” Other groups of investors may have agreed on the

benefit of modification, but had no mechanism for coordinating and communicating that agreement to servicers. Agarwal, Amromin, et al. (2011), for example, find that securitized loans are significantly less likely to be modified than comparable loans held in securities, suggesting that the securitization itself impedes modifications that otherwise would be pursued. Other authors, such as Adelino, Gerardi, and Willen (2013) and Ghent (2011), argue that securitization was not a key impediment to renegotiation.

As servicers could prevent modifications that would have benefited both lenders and borrowers, we take the role for government in our study as only one of coordination the implementation of a modification that might not have been possible without government coordination, but which would not require any expenditure of taxpayer funds beyond the presumably small cost of communicating to achieve coordination.

E Penalties

Taking for granted that penalties exist, are significant, and vary across borrowers, but are unobservable, how should we match borrowers to their penalties?

Two approaches suggest themselves: first, using borrowers' complete payment histories to bound plausible penalties, independent of their other observable characteristics; and second, using borrowers' other observable characteristics to proxy for and thus estimate their penalties. Within the second approach, the modeler has the choice of which observable characteristics to select and which to omit.

In this essay, we prefer the simplest possible flavor of the second approach. We wish our method to be implementable by policymakers, and even policymakers in 2009 could have estimated penalties to credit scores by looking at payment histories through 2009, and also by looking at defaults in prior years. And since credit score is the only boldface borrower characteristic used by lenders for estimating their default propensity but which has no direct mathematical interpretation in a model of default, it is for both reasons the most logical characteristic to use for estimating penalties.

When we do expand the number of admissible predictors of default penalty, we find credit score remains the most robust predictor of a borrower's propensity to default. We consider a reduced form logistic regression in which income, wealth, age, and credit score together predict borrower penalty. Credit score remains the strongest explanatory factor predicting a borrower's likelihood of default. We do find that introducing other factors can explain a realistic non-monotonicity in the data, where borrowers with high credit score

are sometimes disproportionately more likely to default than their estimated penalty would suggest. We speculate their relatively greater default propensity is due to their greater financial sophistication, greater access to legal resources, and greater awareness of the limited financial consequences of default.

We could also take the opposite route: entirely ignore borrowers' observable characteristics and simply estimate penalties as a purely idiosyncratic and uncorrelated error term unique to each borrower. This approach exploits the discrepancy between the continuity in financial incentives facing borrowers who decide to continue or default and the discontinuity in her behavior if she does, using months when the borrower continues to lower-bound penalties and the month when she defaults to reveal an upper bound.

Our policy conclusions are robust to these alternate specifications. We consider alternate approaches further in Kalikman and Scally (2023).

Among other directions for future research, the model, and thus our estimation procedure, assumes that each borrower expects her own penalty not to change in the future. It is possible that shocks such as the Global Financial Crisis, in which borrowers observe mass defaults by neighbors, re-evaluate their estimate of the reputational cost they would suffer in the eyes of their neighbors and also learn information about the foreclosure process, its impact on reputation and credit access, its procedural inconveniences which may further lead them to re-evaluate their penalties; in the opposite direction, borrowers who plan to build credit over time may anticipate that their future penalties would grow with time. If these effects are of first order, then accounting for the possibility of changing penalties should be a fruitful direction for future research.

E.1 Are Penalties Just Proxying For Other Risk Factors?

Borrowers with low credit scores are more likely to default. Is this because credit score reveals information about the borrower's penalty, as we claim? Or is it only because borrowers with lower credit scores are more likely to have higher-interest-rate loans, negative shocks, and other risk factors which together fully account for their higher default propensities? That lenders continue to use credit scores despite full transparency into borrowers' financial circumstances suggests that the predictive power of credit score is not fully explained by these other risk factors.

We can also shed light on these possibilities using the model in a simple test. If credit score does not act through penalty, then the model-predicted default rates for pools of borrowers with different credit scores but a fixed penalty should not be meaningfully worse

than the model-predicted default rates for pools of borrowers with different credit scores and varying penalties: the same borrowers, same financial circumstances, and same correlated risk factors enter in each case, so if the non-penalty risk factors are sufficient to explain how credit score and default propensity correlate, then they should fully account for the variation in default behavior of borrowers with different credit scores.

We test this hypothesis by computing the cumulative default rates for borrowers across five quintiles of credit score. The results for the top four quintiles are shown in Figure 4. Empirical cumulative default rate curves are shown in solid lines. The predictions of the model with idiosyncratic penalties are displayed in the tight dotted lines, while the alternate hypothesis in which all borrowers have the same penalty but credit score alone explains default rates generates the predicted default curves in the long dashed lines.

[Figure 4 about here.]

The fixed-penalty model does generate some heterogeneity in borrower behavior across quintiles of credit score, indicating that credit score does, to some extent, predict defaults by correlation with other risk factors. But if that correlation fully explained the relationship between credit score and default rate, then we would see the dashed lines overlap or approximate the solid lines. Instead, the dashed lines exhibit an inferior fit; the model with idiosyncratic penalties more closely tracks the empirical variation in default rate across credit scores. This exercise suggests both that credit score's power to predict defaults is not fully explained by correlation to other risk characteristics, and that some part of the explanatory power of credit score on default rates is in its correlation to idiosyncratic penalties.

Appendix Figures and Tables

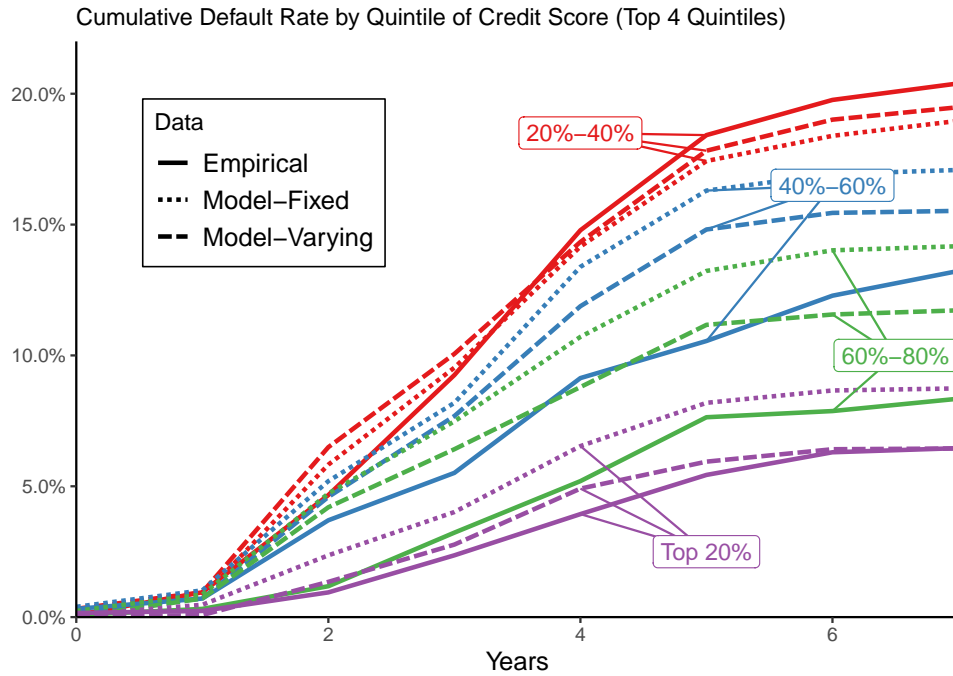


Figure 4: Cumulative Default Rates, Fixed vs. Varying Penalties, by Quintile of Credit Score

Period	Loan 1	Loan 2	Loan 3	Loan 4
0	Current	Current	Current	Current
1	30 DPD	30 DPD	Current	30 DPD
2	60 DPD	60 DPD	Current	60 DPD
3	90 DPD	30 DPD	Current	90 DPD
4	60 DPD	Current	Current	Transferred
5	30 DPD	30 DPD	30 DPD	
6	Current	60 DPD	60 DPD	
7	Current	90 DPD	90 DPD	
\vdots	\vdots	\vdots	\vdots	
T	Paid	Liquidated	Liquidated	

Table 6: Hypothetical Loan Performance Examples

	Log Liquid Assets
(Intercept)	-11.69*** (0.34)
Log Mortgage Debt	0.23*** (0.03)
Log Non-Mortgage Debt	-0.09*** (0.01)
Log Income	1.64*** (0.03)
R ²	0.41
Adj. R ²	0.41
Num. obs.	7994
RMSE	2.11

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 7: Regression Coefficients: SCF Estimation of Liquid Assets

Statistic	Baseline	Protracted Crash
Default Without Modification (%)	55	59
Default After Modification (%)	8	9
Defaults Averted (%)	85	84
Upfront Writedown Waste (%)	49	40
Average Cost of Modification (\$)	5,600	-12,100
Cost per avoided default (\$)	13,000	Negative / Undef.

Table 8: Summary of Writedown Effectiveness and Cost, Historical versus Protracted Crash

Statistic	Sample Loans	130% SLTV Event	Benefits Lender (44%)
Credit Score	717	710	700
First-lien LTV (%)	86	88	89
First-lien PTI (%)	33	39	40
First-lien Size (\$K)	186	217	212
House Price (\$K)	221	251	243
Interest Rate (%)	6.21	6.42	6.47
Long-term ZIP-level Price Drift	0.056	0.056	0.53
Long-term ZIP-level Price Vol	0.058	0.060	0.058

Table 9: Borrower characteristics: overall, experience 130% SLTV, modification benefits lender