Liquidity vs. Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession

Peter Ganong and Pascal Noel*

January 7, 2020

Abstract

We exploit variation in mortgage modifications to disentangle the impact of reducing long-term obligations with no change in short-term payments ("wealth"), and reducing short-term payments with no change in long-term obligations ("liquidity"). Using regression discontinuity and difference-in-differences research designs with administrative data measuring default and consumption, we find that principal reductions that increase wealth without affecting liquidity have no effect, while maturity extensions that increase only liquidity have large effects. This suggests that liquidity drives default and consumption decisions for borrowers in our sample and that distressed debt restructurings can be redesigned with substantial gains to borrowers, lenders, and taxpayers.

*ganong@uchicago.edu, pascal.noel@chicagobooth.edu. This paper subsumes and extends a paper previously circulated as "The Effect of Debt on Default and Consumption: Evidence from Housing Policy in the Great Recession." We thank Sumit Agarwal, David Berger, John Campbell, Raj Chetty, Gabriel Chodorow-Reich, Joao Cocco, John Coglianese, Marco Di Maggio, Will Dobbie, Jan Eberly, Avi Feller, Xavier Gabaix, John Geanakoplos, Edward Glaeser, Paul Goldsmith-Pinkham, Gita Gopinath, Brett Green, Adam Guren, Sam Hanson, Nathan Hendren, Kyle Herkenhoff, Larry Katz, Rohan Kekre, Ben Keys, Arvind Krishnamurthy, David Laibson, Jens Ludwig, Yueran Ma, Laurie Maggiano, Neale Mahoney, Atif Mian, Kurt Mitman, Bill Murphy, Charles Nathanson, Elizabeth Noel, Christopher Palmer, Jonathan Parker, David Scharfstein, Therese Scharlemann, Antoinette Schoar, Amit Seru, Andrei Shleifer, Jon Spader, Jeremy Stein, Johannes Stroebel, Amir Sufi, Larry Summers, Adi Sunderam, Stijn Van Nieuwerburgh, Joe Vavra, Rob Vishny, Paul Willen, Owen Zidar, Eric Zwick, and three anonymous referees for helpful comments. We thank Ari Anisfeld, Therese Bonomo, Guillermo Carranza Jordan, Chanwool Kim, Lei Ma, Jing Xian Ng, and Peter Robertson for outstanding research assistance. Technical support was provided by the Research Technology Consulting team at Harvard's Institute for Quantitative Social Science. This research uses outcomes calculated based on depersonalized credit data provided by TransUnion, a global information solutions company, through relationships with Harvard University and the University of Chicago Booth School of Business. This research was made possible by a data-use agreement between the authors and the JPMorgan Chase Institute (JPMCI), which has created de-identified data assets that are selectively available to be used for academic research. All statistics from JPMCI data, including medians, reflect cells with at least 10 observations. The opinions expressed are those of the authors alone and do not represent the views of JPMorgan Chase & Co. While working on this paper, the authors were compensated for providing research advice on public reports produced by the JPMCI research team. We gratefully acknowledge funding from the Joint Center for Housing Studies, the Washington Center for Equitable Growth, the Hirtle Callaghan Fund, the Charles E. Merrill and Fujimori/Mou Faculty Research Funds at the University of Chicago Booth School of Business, and the National Bureau of Economic Research through the Alfred P. Sloan Foundation Grant No. G-2011-6-22 and the National Institute on Aging Grant No. T32-AG000186.

1 Introduction

Record foreclosure rates and reduced aggregate demand during the Great Recession sparked a vigorous policy debate about how to decrease defaults and increase consumption of struggling borrowers. Former Treasury Secretary Timothy Geithner explained that the government's "biggest debate was whether to try to reduce overall mortgage loans or just monthly payments" (Geithner 2014). Although it was generally believed that debt restructurings would affect both margins, the debate focused on the ideal mix of short-term liquidity provision and long-term debt reduction. A wide range of economists argued that failing to address long-term debt levels by permanently forgiving mortgage principal was a missed opportunity to increase housing wealth and one of the biggest policy mistakes of the Great Recession.[1] Others argued instead that if borrowers are liquidity constrained, focusing on short-term payment reductions is more cost effective (Eberly and Krishnamurthy 2014).

This policy debate hinges on underlying economic questions about the relative effect of short-term liquidity and long-term wealth. A broad literature evaluates changes in mortgage debt that *simultaneously* reduce both short-term payments and long-term obligations.[2] A parallel literature evaluates changes in house prices that *simultaneously* affect short-term borrowing capacity (through collateral effects) and long-term housing wealth.[3] Both literatures consistently find that the combined treatment of short-term liquidity and long-term wealth affect default and consumption. However, to investigate the underlying mechanisms driving default and consumption decisions and to inform the debate about liquidity- versus wealth-focused policy interventions, it is essential to *separately* estimate the effect of short-term liquidity and long-term wealth.

We make progress on this question by exploiting two natural experiments to separately identify the impact of two distinct scenarios: reducing long-term obligations without changing short-term payments ("wealth") and reducing short-term payments without changing long-term obligations ("liquidity"). We find that mortgage principal reduction that increases housing wealth without affecting liquidity has no significant impact on default or consumption for underwater borrowers. In contrast, we show that maturity extension, which reduces payments in the short-term but leaves long-term obligations approximately unchanged, does significantly reduce default rates. Taken together, these results suggest that short-term liquidity drives default and consumption decisions for borrowers in our sample. This lesson

[1] See Goldfarb (2012) for a review of the academic support for principal reductions. For example, he reports that at a meeting to solicit ideas for fixing the ailing economy, President Obama "invited seven of the world's top economists... Nearly all said Obama should introduce a much bigger plan to forgive part of the mortgage debt owed by millions of homeowners who are underwater on their properties." As another example, Geanakoplos and Koniak (2009) argued that a plan to reduce payments and leave principal unchanged "wastes taxpayer money and won't fix the problem."

[2] See e.g. Agarwal et al. (2017*a*), Agarwal et al. (2017*b*), Abel and Fuster (2018), DiMaggio et al. (2017), Ehrlich and Perry (2015), Fuster and Willen (2017), and Tracy and Wright (2016).

[3] See e.g Aladangady (2017), Campbell and Cocco (2007), Carroll, Otsuka and Slacalek (2011), Guren et al. (2018), Mian and Sufi (2011), Mian, Rao and Sufi (2013), and Palmer (2015).

suggests that the collateral channel drives housing wealth effects. Furthermore, it can be used to inform the efficient design of distressed debt restructurings, with the potential for substantial gains to borrowers, lenders, and taxpayers.

Our first natural experiment isolates the effect of long-term wealth by comparing underwater borrowers who receive two types of modifications in the federal government's Home Affordable Modification Program (HAMP). Both modification types result in identical payment reductions for the first five years. However, one group also receives an average of $67,000 in mortgage principal forgiveness, which translates into long-term payment relief. Because borrowers remain slightly underwater even after substantial principal forgiveness, their short-term access to liquidity is unchanged. By exploiting quasi-experimental assignment of borrowers to each of these modification types, we capture the effects of long-term debt levels holding fixed short-term liquidity.

Our second natural experiment generates the opposite treatment: an increase in short-term liquidity with approximately no change in long-term wealth. We compare a set of HAMP borrowers who receive a small payment reduction to borrowers who receive a large payment reduction through alternative private sector modifications. The private sector finances this deeper payment reduction by first extending mortgage maturity prior to additional modification steps, such that the larger immediate payment reduction is offset by continued payments in the long term. This restructuring leaves the net present value (NPV) of total mortgage payments owed approximately unchanged. By exploiting a cutoff rule in assignment to these two modification types we isolate the effect of short-term liquidity provision holding fixed long-term wealth.

To study these natural experiments we build two new datasets with information on program participation and borrower outcomes. Our first dataset matches administrative data on HAMP participants to credit bureau records. We exploit detailed account-level information to construct a novel measure of consumer spending based on monthly credit card expenditures. Our second dataset uses de-identified mortgage and credit card data from the JPMorgan Chase Institute (JPMCI). It includes monthly information on all borrowers whose mortgages are serviced by Chase and who receive either a government-subsidized modification through HAMP or an alternative private modification. Our samples from both datasets are similar on observable borrower characteristics.

Using our first natural experiment, we estimate the causal impact of principal reduction on default by exploiting a cutoff rule in borrower assignment to the two HAMP modification types. Mortgage servicers evaluated underwater applicants for both modification types by calculating the expected gain to investors under each type using a standardized government-supplied formula. When the calculation shows that principal reduction is marginally more beneficial to investors, there is a sharp jump—41 percentage points—in the probability that a borrower receives principal reduction. We exploit this jump with a regression discontinuity estimator that compares borrowers on either side of this cutoff.

We find that principal reduction has no effect on default. Despite a $31,000 increase in principal forgiveness in the treatment group at the cutoff (translating to an 11 percentage point reduction in a borrower's loan-to-value ratio), default rates are unchanged. This implies a very large or possibly infinite cost to the government per avoided foreclosure. Even at the most optimistic point in our confidence interval, the government spent at least $365,000 per avoided foreclosure. This cost is almost an order of magnitude greater than estimates of the social cost of foreclosures (U.S. Department of Housing and Urban Development 2010).

We next examine the causal impact of principal reduction on consumption using the same government modification program. Our preferred empirical strategy for analyzing consumption is a panel difference-in-differences estimator, which is more precise than our regression discontinuity estimator. We find that an average of $67,000 in principal reduction has no significant impact on underwater borrowers' credit card or auto expenditure. Translating our results into an annual marginal propensity to consume (MPC) for total consumption, our point estimate is that borrowers increased consumption by a statistically insignificant three-tenths of one cent per dollar of principal reduction, with an upper bound of less than one cent.

Using our second natural experiment, we estimate the causal impact of short-term payment reductions on default. This analysis exploits a cutoff rule that determines eligibility for HAMP using a regression discontinuity design. There is a sharp jump in the amount of payment reduction received by borrowers with private modifications just below the cutoff. Although there is a large change in short-term liquidity at the cutoff, because this deeper payment reduction is largely financed by extending mortgage maturities, there is no change in the NPV of total long-term payments owed.

In contrast to our results on the ineffectiveness of principal reduction, we find that short-term payment reduction significantly reduces default rates. Default rates fall sharply by 7 percentage points at the cutoff from a control mean of 32 percentage points, implying that a one percent payment reduction reduces default rates by about one percent. While our data and available research designs are unsuited for credibly estimating the causal effect of short-term payment reduction on consumption, we provide suggestive evidence from the time-series pattern of spending around modification that spending also rises when monthly payments fall.

Combining our empirical results, this paper's central contribution is to disentangle the effects of short-term liquidity and long-term wealth on borrower outcomes. We find that liquidity—and not wealth—drives consumption and default decisions for borrowers in our sample. This allows us to draw two types of lessons.

First, payment reduction can be structured to benefit borrowers, lenders, and taxpayers—a sharp contrast with principal reduction, which is both costly and ineffective for underwater borrowers. In particular, our default results show an inefficient allocation at the HAMP eligibility cutoff. The government spent substantial resources subsidizing HAMP modifications

above the cutoff with small payment reductions and high default rates. In contrast, borrowers below the cutoff received private modifications emphasizing maturity extension that required no government assistance, which had large payment reductions and low default rates. In fact, there is likely a Pareto improvement for borrowers, lenders, and taxpayers from shifting the cutoff to reallocate borrowers from HAMP to private modifications. Such a reallocation was prohibited by government rules requiring that HAMP be offered first to any eligible borrower above the cutoff. This requirement crowded out more effective modifications for up to 40 percent of HAMP borrowers.

This lesson can be used to characterize the default-minimizing modification structure for all borrowers. Since short-term liquidity reduces default rates but long-term wealth does not, the efficient modification structure maximizes liquidity provision.[4] We find the potential for substantial gains to borrowers, lenders, and taxpayers relative to existing public and private modifications. One way to quantify the potential gains is to analyze a hypothetical modification that maximizes the amount of payment reduction offered to borrowers while holding fixed the costs to lenders and taxpayers. If our discontinuity-based treatment effects extrapolate to other HAMP borrowers, it would have been possible to cut default rates by one-third, avoiding 267,000 defaults at *no additional cost* to lenders or taxpayers.

Second, our consumption results help distinguish between the liquidity- and wealth-based explanations for the robust relationship between housing wealth and consumption. Consumption responses to home equity gains could reflect an increase in long-term wealth or a relaxation of collateral constraints. Because house price changes typically affect both wealth *and* collateral, it has been difficult to separate these effects (Cloyne et al. 2019). However, a reduction in mortgage principal that leaves a borrower underwater increases that borrower's NPV of wealth (by reducing their long-term debt obligations), but does not relax their immediate collateral constraint. Hence, our setting isolates the wealth channel holding the collateral channel fixed. Our estimated MPC from principal reduction is an order of magnitude smaller than prior estimates of the MPC out of housing wealth that combine both channels.[5] Thus, our results suggest that the wealth channel alone is weak and that relaxing collateral constraints is a *necessary* condition for housing wealth to stimulate consumption. This finding complements prior work that isolates the effect of collateral holding wealth fixed, which finds that relaxing collateral constraints is a *sufficient* condition for housing wealth to stimulate consumption.[6]

Because liquidity drives housing wealth effects, we find that the tight link between housing

[4]This is consistent with the conclusions in Eberly and Krishnamurthy (2014). The lessons about ex-post renegotiation also help inform a growing theoretical literature about optimal ex-ante mortgage design and its macroeconomic implications. See Campbell, Clara and Cocco (2018), Eberly and Krishnamurthy (2014), Greenwald, Landvoigt and Van Nieuwerburgh (2018), Guren et al. (2018), Gorea and Midrigan (2017), Hedlund (2015), and Piskorski and Tchistyi (2010).

[5]A large literature examines the consumption response to house price changes and typically estimates an MPC of around around five cents per dollar. See e.g. Aladangady (2017), Campbell and Cocco (2007), Carroll, Otsuka and Slacalek (2011), Guren et al. (2018), and Mian, Rao and Sufi (2013).

[6]See Agarwal and Qian (2017), Cloyne et al. (2019), Defusco (2017), and Leth-Petersen (2010).

wealth and consumption breaks down when borrowers are underwater. Home equity gains do not relax collateral constraints for underwater borrowers and therefore do not affect consumption because households cannot increase borrowing to monetize these gains. Indeed, we show that collateral constraints drive a wedge between the MPC out of cash and the MPC out of housing wealth for underwater borrowers. Thus, policies such as principal reduction are unable to stimulate demand when borrowers are so far underwater that home equity gains fail to relax binding collateral constraints. This highlights the general principle that when borrowing constraints matter for real outcomes, programs can be ineffective if they fail to target these constraints.

The ineffectiveness of long-term principal reduction at boosting short-term consumption has two further implications for models. First, our results provide evidence that the timing of liquidity matters, consistent with the predictions of models with incomplete markets. A substantial literature has implemented tests for incomplete markets by showing that current consumption *responds* to *current* liquidity (e.g. Johnson, Parker and Souleles 2006, Zeldes 1989). We provide complementary evidence by showing that current consumption is *unresponsive* to changes in *future* liquidity. Second, our findings contribute to a new literature, which finds little direct linkage between debt levels and consumption when debt is modeled as a long-term contract (Kaplan, Mitman and Violante (2017), Justiniano, Primiceri and Tambalotti 2015). This contrasts with debt overhang models in which forced deleveraging of short-term debt leads to depressed consumption during a credit crunch (Eggertsson and Krugman 2012, Guerrieri and Lorenzoni 2017). We show in a simple model with long-term debt that if nothing forces borrowers to immediately delever when they are far underwater, the mechanical link between debt levels and consumption is removed and principal reduction becomes less effective.

The remainder of the paper is organized as follows. Section 2 describes the data. Sections 3 and 4 analyze the effect of principal reduction on default and consumption, respectively. Section 5 analyzes the effect of payment reduction on default. Section 6 provides discussion and interpretation of the empirical results. The final section concludes.

2 Data

We use two datasets. Our first dataset matches administrative HAMP participation data to consumer credit bureau records. This dataset allows us to analyze the mechanisms assigning borrowers to each modification type in HAMP, which we exploit to estimate the impact of principal reduction. Our second dataset comes from a bank that is also a servicer that offers both government-subsidized HAMP modifications as well as private modifications. This allows us to analyze variation in short-term payment reduction between public and private modifications and to examine administrative spending data.

2.1 Matched HAMP Credit Bureau File

The U.S. Treasury releases a public data file on the universe of HAMP applicants. This

loan-level dataset includes information on borrower characteristics and mortgage terms before and after modification. Crucially, it also includes the expected gain calculation run by servicers when evaluating borrowers for each modification type.

In order to observe consumption for borrowers in the HAMP public file, we use de-identified consumer credit bureau records from TransUnion. HAMP program rules require servicers to report borrower participation to credit bureaus. We use the universe of records for borrowers flagged as having received HAMP. We have monthly account-level information between January 2010 and December 2014 for each borrower.

We develop proxies for both durable and nondurable consumption based on the credit bureau records. For durable consumption, we follow DiMaggio et al. (2017) by using changes in auto loan balances as a measure of car purchases. DiMaggio et al. (2017) document that leveraged car purchases account for 80 percent of new car sales. While prior work relied on observing jumps in total auto loan balances to infer new loans, our product account-level data allows us to observe new loans directly.

The detailed nature of our credit bureau data also allows us to construct a new measure of consumption based on credit card expenditures. In particular, we calculate monthly expenditures using end of month balances and payments made in a given month.[7] We are able to construct this measure for 83 percent of all credit and charge card accounts (not all servicers report monthly payments). We find average credit card spending of $452 per month in our sample, which is 84 percent of the average credit card spending per adult in 2012 (Federal Reserve System 2014), commensurate with the 83 percent of cards for which we observe expenditures.

We match borrowers in the HAMP dataset to their credit bureau records using loan and borrower attributes present in both files: metro area, modification month, origination year, loan balance, and monthly payment before and after modification. When two borrowers are listed on a mortgage, we measure consumption using the credit bureau records of both borrowers. We are able to match half of the records in our sample window, resulting in a panel dataset of about 106,000 underwater households eligible for both HAMP modification types.[8]

The imperfect match rate does not bias our sample in terms of any observed borrower characteristics. Online Appendix Table 1 reports summary statistics for our sample before

[7] Let b_t denote the balance at the end of month t, and p_t be the payment made in month t. We calculate expenditure in month t as $e_t = b_t - b_{t-1} + p_t$. See online Appendix B.1.1 for details on construction of the expenditure variable. Because interest rates and fees are not reported, we do not distinguish between new purchases, interest charges, and fees in this dataset. In the bank dataset described in Section 2.2, we can isolate purchases and confirm that our results are unchanged.

[8] See online Appendix B.1.1 for details on the matching procedure. Our match rate is less than 100 percent due to rounding and changing reporting requirements. The main data limitation is that pre-modification principal balance and monthly payment fields are rounded in the Treasury HAMP file, which introduces a discrepancy between the same loans in both files. Another limitation is that construction of the Treasury file required new reporting processes for participating servicers, and the reporting requirements changed several times as the program developed. As a result, Treasury explains that there are occasional inaccuracies in the underlying data (U.S. Department of the Treasury 2014*a*).

and after the credit bureau match. This table shows that borrower characteristics are similar in the matched sample. The final column shows that the difference in means for any characteristic is less than one-fifth of a standard deviation. For our regression discontinuity design to identify the causal impact of principal reduction on default in the presence of incomplete matching, we need the match rate to be smooth at the cutoff. We show that this is the case in online Appendix Figure 1. In Section 4.2 we show that our consumption result is unchanged (though slightly less precise) when we estimate it using the borrowers in the bank dataset, which does not rely on matching and is described in the following section.

2.2 JPMCI Bank Dataset

Our second dataset includes de-identified account-level monthly information on all mortgages serviced by Chase bank and spending by mortgagors who also had a Chase credit card. The dataset covers 2009 to 2016. We focus on two subsamples of borrowers. The sample we use as a robustness check to study the effect of principal reduction on consumption includes all HAMP borrowers with both a mortgage and a credit card with Chase. We observe credit card spending for 10,741 borrowers one year before and after modification.

The sample we use to study the effect of payment reduction includes all borrowers who receive either a government-subsidized HAMP or private modification. This includes 59,726 mortgages owned or securitized by Fannie Mae and Freddie Mac (the government sponsored enterprises, or GSEs) and 86,580 mortgages which are owned or have been securitized by Chase. We limit the sample to modifications performed in the fourth quarter of 2011 or later, when the particular versions of the private programs we study were sufficiently established.[9] We analyze the impacts on GSE-backed and non-GSE-backed mortgages separately in Section 5.

3 Effect of Principal Reduction on Default

In this section we analyze the effect of principal reduction on borrower default. We compare borrowers who received two different types of government-subsidized modifications, with both types receiving identical short-term payment reductions but one type receiving additional principal reduction. Using a regression discontinuity (RD) empirical strategy we find that substantial principal reductions have no effect. We can rule out prior cross-sectional estimates that were used to justify the program.

3.1 Variation in Principal Reduction in the Home Affordable Modification Program

The government instituted the HAMP program in 2009 as a response to the foreclosure crisis. It provided government subsidies to help facilitate mortgage modifications for borrow-

[9] Both Chase and the GSEs had a variety of other private modification programs with different designs that preceded HAMP.

ers struggling to make their payments. In total, 1.8 million borrowers received modifications through the program.

The government designed HAMP's eligibility criteria to target the borrowers it perceived as most likely to benefit from modifications. Borrowers must have current payments greater than 31 percent of their income, be delinquent or in imminent default at the time of their application, attest that they are facing a financial hardship that makes it difficult to continue making mortgage payments, and report that they do not have enough liquid assets to maintain their current debt payments and living expenses. In almost all cases, borrowers must be owner-occupants and have loan balances of less than $730,000.[10]

The primary goal of HAMP modifications is to provide borrowers with more affordable mortgages. All borrowers who receive modifications have their payment reduced to reach a 31 percent Payment-to-Income (PTI) ratio for at least five years. This rule results in substantial modifications for many borrowers. The mean payment reduction is $680 per month, or 38 percent of the borrower's prior monthly payment.

Our research design relies on contrasting borrowers assigned to two distinct modification types. Both modification types result in the same payment reduction for the first five years, but each type achieves this payment reduction in a different way.

The first modification type provides what we call a "payment reduction" modification. Figure 1a shows the average annual payments for borrowers in this modification type relative to their payments under the status quo. This modification implements up to three steps to achieve the 31 percent PTI target. First, the interest rate is reduced down to a floor of two percent for a period of five years, after which it gradually increases to the market rate. Second, if the target is not reached after the interest rate reduction, the mortgage maturity is extended up to 40 years. Third, if the target still is not reached, a portion of the unpaid balance is converted into a non-interest-bearing balloon payment due at the end of the mortgage term.

The second modification type is what we call a "payment and principal reduction" modification (also known as the HAMP Principal Reduction Alternative). The first step in this modification is to forgive a borrower's unpaid principal balance until the new monthly payment achieves the 31 percent PTI target or their Loan-to-Value (LTV) ratio hits 115 percent, whichever comes first. If the borrower's monthly payment is still above the target, then the interest rate reduction, maturity extension, and principal forbearance steps described above are followed as needed. 245,000 borrowers received these modifications.

The government introduced these principal reduction modifications in October 2010 in response to growing concern that long-term debt levels, rather than just short-term debt payments, were responsible for high default rates and depressed consumption. The government devoted substantial resources towards supporting principal reduction modifications. On av-

[10]These two criteria rule out borrowers who might be particularly likely to strategically default. However, such ineligible borrowers are responsible for a small share of defaults. Eighty-six percent of defaults in 2009 were for borrowers who met the owner-occupancy and loan balance criteria (Agarwal et al. 2017*a* Table 1).

erage, the government paid an additional $20,000 per modification to support modifications with principal reduction (Scharlemann and Shore 2016).

By comparing borrowers who receive these two types of modifications, we can estimate the effect of long-term debt obligations holding short-term payments constant. The two types of modifications have identical effects on payments in the short term, but dramatically different effects on long-term payments and homeowner equity. Figure 1a shows that payment reductions are identical for the first five years, after which payments rise more sharply for borrowers with payment reduction modifications. Figure 1b summarizes the financial impacts of these modifications for borrowers in our sample. Borrowers with principal reduction modifications receive an average of $67,000 more principal reduction.[11]

The monetary value of the principal reduction depends on borrower behavior. To a borrower who prepays her mortgage the next day, principal reduction is worth $67,000, but it is worth nothing to a borrower who immediately defaults and never repays. We calculate the value to borrowers using two methods. First, we calculate that the incremental reduction in the NPV of payments owed under the mortgage contract if the borrower repays on schedule is $34,000. This calculation assumes borrowers discount future cash flows at the average market interest rate, consistent with the empirical findings in Busse, Knittel and Zettelmeyer (2013) for auto loans.[12]

Second, we calculate the NPV of expected payments using observed prepayment and default behavior of HAMP borrowers. Prepayment raises the NPV to the borrower and default lowers it. We provide details on our valuation method in online Appendix C.1. The default effect dominates and we calculate a change in NPV of $28,000.

Program administrators took steps to ensure that borrowers understood the new mortgage terms. The cover letter for the modification agreement prominently listed the new interest rate, mortgage term, and amount of principal reduction. Additionally, the modification agreement included a summary showing the new monthly payment each year, as shown in online Appendix Figure 2. Borrowers appear eager to take up modifications. Conditional on being offered a modification, 97 percent of borrowers accepted the offer.

3.2 Identification: Discontinuity in Principal Reduction at Treasury Model Cutoff

Borrower assignment to different modification types is determined in part by a cutoff rule, and in part by servicer and lender type. This assignment generates quasi-experimental variation, which we will exploit in our empirical strategies. We discuss the cutoff rule—which

[11] Some borrowers in the payment reduction modification type received small amounts of principal reduction. This is because some servicers wanted to provide principal forgiveness outside of the Treasury incentive program, which only paid incentives for forgiveness above 105 percent LTV and required the forgiveness to vest over three years.

[12] This is also consistent with a "market-based" conception of wealth where valuation does not differ across individuals. However, for an individual conception of wealth, the gains are still substantial even for a more impatient borrower. For example, if instead we assume the borrower's discount rate is twice the mortgage interest rate, principal forgiveness reduces the NPV of payments owed under the contract by $18,000.

we use in a regression discontinuity to estimate the impact of principal reduction on default—in this section. We use variation in servicer and lender type to estimate consumption impacts, and so we defer an explanation of that variation to Section 4.

Our quasi-experimental variation covers the period with the most severe delinquency rates in the recent crisis. Our sample of borrowers have their first delinquencies in the fourth quarter of 2009, just before the peak of the delinquency crisis, which did not begin abating until 2013. Online Appendix Figure 3 plots the delinquency rate for all U.S. borrowers over time.

Principal reduction is determined in part by a calculation examining which modification type is expected to be most beneficial for the lender. Using a model developed by the U.S. Treasury Department, servicers calculate the expected NPV of cash flows for lenders under the status quo and under each of the two modification types described in Section 3.1. The NPV model takes into consideration government-provided incentives as well as the expected impact that modifications will have on default and prepayment.

The Treasury NPV model is designed to encourage principal reduction modifications by reducing contracting frictions between lenders and servicers. The government could not force servicers to offer principal reduction to borrowers under the existing contracts between servicers and lenders. However, the government could compel servicers to run the Treasury NPV model. Servicers are bound by their fiduciary duty to the lenders to maximize repayment, and as a result are more likely to offer the modifications shown to be most beneficial to lenders.

Our empirical strategy exploits a large jump in the share of borrowers receiving modifications with principal reductions when the NPV model shows it will be marginally more beneficial to lenders than the alternative. This jump is shown in Figure 2a.

We identify the effect of principal reduction on default using the cutoff in the expected benefit to lenders with a regression discontinuity design. Let the receipt of principal reduction treatment be denoted by the binary variable $T \in \{0, 1\}$, where 0 represents receiving a payment-reduction-only modification, and let X capture the characteristics of the borrower. The Treasury NPV model calculates the expected NPV to lenders $ENPV(T, X)$ under either scenario. Our running variable V is the normalized predicted gain to lenders of providing principal reduction to borrowers, that is

$$V(X) = \frac{ENPV(1, X) - ENPV(0, X)}{ENPV(0, X)}. \tag{1}$$

A realization v reflects the anticipated percent gain to the lender from principal reduction relative to a standard modification. The cutoff that affects assignment to treatment or control is at $v = 0$.

Borrowers near this cutoff are those for whom the Treasury model predicts a large average reduction in default from principal reduction that is offset by reduced cash-flows from non-

defaulting borrowers. We normalize the predicted gain by $ENPV(0, X)$ to avoid a high concentration of low-balance mortgages near the cutoff. We describe the sample construction in more detail in online Appendix B.1.1, provide more details on what gives some borrowers high or low values of v in the Treasury NPV model in online Appendix B.1.2, and further discuss the normalization in online Appendix B.1.3.

The treatment effect of receiving principal reduction is determined by the jump in default divided by the jump in the share receiving principal reduction at the cutoff. Let Y be the outcome variable of interest (such as default). The fuzzy RD estimand is

$$\tau = \frac{lim_{v\downarrow 0}E[Y|V=v] - lim_{v\uparrow 0}E[Y|V=v]}{lim_{v\downarrow 0}E[T|V=v] - lim_{v\uparrow 0}E[T|V=v]}. \tag{2}$$

The parameter τ identifies the local average treatment effect of providing principal reduction to borrowers near the cutoff. We follow the standard advice for RD designs from Lee and Lemieux (2010) and Imbens and Kalyanaraman (2012) to estimate $\hat{\tau}$ using a local linear regression. Our analysis dataset is the matched HAMP credit bureau dataset, which includes the predicted gain to investors of providing principal reduction v.

In Table 1a we compare summary statistics for borrowers in our sample near the assignment cutoff to the characteristics of delinquent borrowers in the Panel Study of Income Dynamics between 2009 and 2011. Borrowers in our sample are broadly representative of delinquent underwater borrowers during the recent crisis. We provide more detail on this comparison in online Appendix Section B.1.4.

Predicted default rates based on pre-determined covariates trend smoothly through the cutoff, as shown in online Appendix Figure 4. Some servicers ran only one NPV calculation and reported this single number as the NPV calculation for both "payment reduction" and "payment & principal reduction" modifications, meaning that they reported $ENPV(1, X) = ENPV(0, X)$. Following the advice of U.S. Treasury staff, we assume that observations exactly at zero reflect misreporting and subsequently drop them from the analysis sample. Online Appendix Figure 5 shows that density in the analysis sample is smooth around the cutoff. We provide additional detail on both covariate balance and smoothness in online Appendix B.1.3.

3.3 Results: Effect of Principal Reduction on Default

Figure 2a shows that there is a discontinuous jump of 41 percentage points in the share of borrowers receiving principal reduction at the cutoff. Measured in terms of dollars of principal reduction, the treatment size at the cutoff is \$31,000.[13] This reduces borrower LTV

[13] This is smaller than the average of \$67,000 across all principal reduction recipients. Because the program targeted LTV of 115, borrowers with lower pre-modification LTV are eligible for less principal reduction. The running variable v in our design is the Treasury NPV model's estimate of the relative gain (or loss) to investors from principal reduction. The absolute value $|v|$ is larger when a borrower is eligible for more principal reduction. Because we study borrowers with $v \approx 0$, our research design identifies a treatment effect for borrowers who are eligible for a smaller, but still substantial, amount of principal reduction.

by 11 percentage points, which amounts to a $17,000 reduction in the NPV of borrower payments owed over the full mortgage term. Importantly, there is no jump in monthly payment reduction at the cutoff, highlighting that the treatment we are analyzing is a reduction in mortgage principal that leaves short-term payments unchanged. The relationship of the four aforementioned variables with respect to the running variable is shown in online Appendix Figure 6.

We find that principal reduction has no impact on default. Figure 2b shows the reduced form of the fuzzy RD specification, plotting the default rate against the running variable. We define default as being 90 days delinquent at any point between modification date and March 2015, when our HAMP dataset ends, which is an average of three years. This is the measure of default used to disqualify a borrower from the HAMP program and is the common measure used in the prior literature discussed in Section 3.4. There is no jump in default rates at the cutoff, and we can rule out a reduction of more than five percentage points using the 95 percent confidence interval. Online Appendix Figure 7 shows that our estimates are close to zero for a wide range of bandwidth choices, and these results are discussed in more detail in online Appendix B.1.3.

Our results imply a large or possibly infinite government cost per avoided foreclosure. While we do not follow borrowers through to completed foreclosures within our data, government reports show that 45 percent of HAMP borrowers who default eventually end up with a foreclosure (U.S. Department of the Treasury 2017). Thus, even taking the most optimistic point in the confidence interval for the effect of principal reduction on default, this translates into at most a 2.3 percentage point reduction in foreclosure during the window we study.[14] The government spent about $8,000 per modification to support the additional principal reduction of the size we analyze in our treatment group. This translates into a cost of at least $365,000 per avoided foreclosure, almost an order of magnitude larger than common estimates of the social costs of foreclosure (U.S. Department of Housing and Urban Development 2010).

Principal reduction was also costly to lenders. Even when using the most optimistic point in the confidence interval, we estimate that lenders had to forgive at least $1.3 million in principal to prevent one foreclosure. However, this writedown was partially offset by two forces. First, government subsidies would have reimbursed a portion of the cost, as described above. Second, lenders would not have expected to recoup all of this principal because some borrowers would have defaulted under the status quo. Altogether, after accounting for these forces, we estimate that lenders would have lost at least $402,000 for each foreclosure prevented. Online Appendix C.2 contains additional detail on the calculations in the two

[14] Although we find little impact on foreclosures within our three-year analysis window, it is possible that once these borrowers regain positive equity several years in the future, foreclosures for the principal reduction group will be lower than for those who did not receive it. Unfortunately this is not something we can analyze with the data in this paper. Furthermore, to the extent that the policy goal was short-term housing market stabilization, the benefit of future foreclosure reduction is limited.

prior paragraphs.

Would more aggressive principal reduction have been a superior policy? Foreclosures could have been mechanically avoided by bringing borrowers all the way into positive equity, but this also would have been an expensive strategy after prices had fallen substantially. The average underwater borrower evaluated for principal reduction had approximately $100,000 in negative equity. Even if foreclosures were completely eliminated by forgiving 100 percent of this negative equity (since defaulting borrowers could then sell their home and avoid a foreclosure), this would require $1.3 million in writedowns to avoid a single foreclosure. The intuition behind this finding is that most underwater borrowers keep paying even in the absence of principal reduction, so negative equity needs to be eliminated for many borrowers in order to avoid one foreclosure. Eliminating borrowers' negative equity becomes more attractive as the baseline foreclosure rate without principal reduction rises. In the limit, if every home is going to be foreclosed on in the absence of principal reduction, then offering principal reduction is costless to the lender because they would never have received this principal initially. We calculate that, in the absence of any alternative modification steps, eliminating all negative equity is cost-effective from the investor's perspective only when the default rate exceeds 77 percent.

3.4 Comparison to Prior Evidence on Default

Our results are inconsistent with prior evidence based on the cross-sectional relationship between negative equity and default. For example, Haughwout, Okah and Tracy (2016) use data on modifications performed prior to HAMP and find, using cross-sectional variation, that borrowers who received principal reductions equivalent to ours saw an 18 percentage point reduction in default. Furthermore, there is a strong cross-sectional relationship between the amount of negative equity and mortgage default rates across all borrowers (Gerardi et al. 2018).[15]

Indeed, the U.S. Treasury Department developed a model based on this historical data that predicts a substantial reduction in default from principal reduction, which is inconsistent with our findings. The Treasury generated this estimate as part of its model to predict the benefits of modifications to lenders (Holden et al. 2012). We implement the Treasury redefault model (U.S. Department of the Treasury 2015) in the public HAMP data and calculate the predicted impact of principal reduction at the cutoff. The Treasury model expects a reduction in default of 7.3 percentage points at the NPV cutoff, which we can rule out using our 95 percent confidence interval.

Why is our causal estimate so much smaller than what is predicted by the cross-sectional

[15] Outside of mortgages, Dobbie and Song (2019) analyze future payment reductions for credit card borrowers. In contrast to our findings, they find that reducing future payments by 8 percent of the total debt owed leads to a reduction in short-term default of 1.6 percentage points. When scaled to an equivalent treatment size, this is larger than our point estimate but within our confidence interval. One possible explanation is that borrowers behave more strategically with respect to credit card debt because the consequences of default are less severe than defaulting on a mortgage, which often results in foreclosure.

relationship between borrower equity and default and models calibrated to this relationship? One possibility is that the cross-sectional evidence was misleading because borrowers with less equity were also borrowers who purchased homes near the height of the credit boom and who therefore might have been less credit-worthy on other dimensions. Palmer (2015) shows that changes in borrower and loan characteristics can explain 40 percent of the difference in default rates between the 2003-2004 and the 2006-2007 cohorts. Another possibility is that the large price reductions that left many borrowers underwater were also correlated with other omitted economic shocks that themselves could be responsible for higher default rates (Adelino, Schoar and Severino 2016).

Our results using a *nonparametric* identification strategy complement prior work by Scharlemann and Shore (2016) (henceforth SS), who use a *parametric* identification strategy to also examine the effect of principal reduction in HAMP. That paper's research design exploits the fact that principal reduction is a kinked function of LTV. Principal reduction in HAMP reduces borrower LTV to a cutoff of 115, and SS's preferred specification relies on borrowers far from the cutoff, with pre-modification LTV values as high as 240. This empirical strategy is parametric because it assumes that the relationship between default and LTV would be globally linear in the absence of principal reduction. Such a specification is biased if there are any non-linearities in the relationship between the outcome and the running variable. To address this type of potential bias from functional form assumptions, the identification results for RD and regression kink designs call for estimation strategies to flexibly estimate the regression function by relying only on data close to the cutoff (Hahn, Todd and Van der Klaauw 2001; Nielsen, Sorensen and Taber 2010). We use local linear regression and an optimal bandwidth procedure to achieve nonparametric identification in our study.[16]

In spite of the differences in methodology, our research design and SS's research design both imply that principal reduction has at most small impacts on foreclosures. Although our point estimates are not directly comparable due to differences in the size of the principal reduction treatments we study, one common metric to compare our estimates is the cost to the government per foreclosure avoided. SS estimate a cost of $320,000, which is slightly smaller than the most optimistic point in our confidence interval but still six times larger than prevailing estimates of the social cost of foreclosure. Overall, our findings reinforce their policy conclusion that principal reduction is not a cost-effective strategy for reducing defaults. Furthermore, our paper also examines the effects of principal reduction on consumption and

[16]SS explain that one of the challenges to achieving nonparametric identification by implementing a regression kink design at their cutoff is that there is little identifying variation at this cutoff. SS write: "It should not be surprising that we lose power in the region very near the kink. Borrowers who are near but on opposite sides of the kink receive nearly identical treatments. One must look relatively far from the kink to find borrowers with substantial differences in principal forgiveness, and consequently different default rates." This challenge forces SS to rely on data far from the kink in their central estimates, rather than using data close to the kink as required by the identification results for RD and regression kink designs. In contrast, in the RD design that we study, there is substantial variation in treatment at the cutoff.

payment reduction on default, to which we turn next.

4 Effect of Principal Reduction on Consumption

In this section we explore the effect of principal reductions on consumption. Using a difference-in-differences empirical strategy we find that principal reductions affecting wealth but not liquidity have no significant impact on consumer spending.

4.1 Identification: Panel Difference-in-Difference Empirical Strategy

Our analysis of consumption motivates a change in research design to a panel difference-in-differences strategy for two reasons. First, our RD strategy is under-powered for studying changes in consumption. Economically meaningful consumption changes cannot be ruled out using an RD design. As we discuss in more detail in Section 4.3, even a small change in consumption on the order of five cents for each dollar of principal forgiven would be meaningful relative to average marginal propensities to consume out of housing wealth changes studied in other contexts, whereas the predicted impacts on default from the prior literature were much larger. The second reason is that the panel nature of the spending measures from our credit bureau and banking data allow us to exploit an alternative strategy that offers better precision. Lagged spending measures allow us to adjust for underlying differences between borrowers receiving different modification types within a wider bandwidth than with the RD. These factors favor a panel difference-in-differences design, though we also report results from the RD strategy.[17]

Our panel difference-in-differences design uses as a control group the set of underwater borrowers who were eligible for principal reductions, but who instead received only payment reduction modifications. This design relies on the fact that borrowers who receive payment reduction modifications experience the same short-term payment reductions as borrowers who receive principal reduction, but they receive substantially less generous long-term payment relief. Summary statistics for both groups are shown in online Appendix Table 2.[18] The size of short-term payment reductions are nearly identical across groups, but borrowers who receive payment and principal reduction modifications receive on average $67,000 more principal reduction, reducing the NPV of the payments owed under their mortgage contract by an additional $34,000. In accordance with the HAMP rules described in the previous section, borrowers who received principal forgiveness remained underwater (usually at 115 percent LTV). Thus, the treatment captures the effect of long-term debt forgiveness holding short-term payments and access to liquidity fixed.

[17] We also have lagged measures of default from the credit bureau data. However, a difference-in-differences design is not valid for studying default because pre-treatment differences in the levels of default are mechanically removed at modification date, at which point all loans become current. This means that the change in default for the control group is not a valid counterfactual for the change in the treatment group.

[18] Our main sample for this analysis includes underwater borrowers in the matched HAMP credit bureau dataset who are observed one year before and after modification and report positive credit card expenditure in at least one month during this window.

Our identification comes from cross-servicer and cross-lender variation in the propensity to provide principal reductions given observed borrower characteristics. Borrowers are not assigned to principal reduction modifications according to the NPV calculation alone because different lenders have different views about principal reduction and servicers are not always confident they have the contractual right to forgive principal or the capacity to manage the process.[19] Conditional on lender and servicer, all borrowers are treated alike. Servicers must submit a written policy to the Treasury department detailing when they will offer principal reduction modifications and attesting that they will treat all observably similar borrowers alike (U.S. Department of the Treasury 2014*b*). Intuitively, this strategy compares borrowers with loans from servicer-lenders that were more likely to offer principal reduction to borrowers whose servicer-lenders were less likely to offer principal reduction.

The key identifying assumption for the panel difference-in-differences design is that consumption trends would be the same in both groups in the absence of treatment. This assumption is plausible when the two groups exhibit parallel trends before treatment. We show this visually in Figure 3a, which plots mean credit card expenditure around modification date.[20] The same figure shows that principal reduction appears to have little effect—a result we explore in a regression framework.

Formally, our main specification is:

$$y_{i,g,t,s} = \gamma_g + \gamma_t + \gamma_{m(i),s} + \beta\left(Principal\ Reduction_g \times Post_t\right) + x'_{it}\delta + \varepsilon_{i,g,t,s}, \qquad (3)$$

where i denotes borrowers, $g \in$ {payment reduction, payment & principal reduction} the modification group, t the number of months since modification, s the calendar month, and m the household's Metropolitan Statistical Area (MSA). Our main outcome variables $y_{i,g,t,s}$ are monthly credit card and auto expenditure, which proxy for non-durable and durable spending, respectively. γ_g captures the modification group fixed effect, and γ_t captures a fixed effect for each month relative to modification. $Principal\ Reduction_g$ is a dummy variable equal to 1 for the group receiving modifications with principal reduction while $Post_t$ is a dummy variable equal to 1 for $t \geq 0$. The main coefficient of interest is β, which captures the difference-in-differences effect of principal reduction.

One potential concern is that different geographies were experiencing different trends in their house price recoveries, which affected borrower outcomes. To address this concern $\gamma_{m(i),s}$ captures MSA-by-calendar-month fixed effects. x_i is a vector of individual characteristics de-

[19]The contractual frictions are particularly acute with securitized loans. For example, Kruger (2018) shows that 22 percent of servicing agreements governing securitized pools explicitly forbid servicers from reducing principal balances as part of modifications. As a result, principal reduction in HAMP was less common among borrowers in securitized pools (Scharlemann and Shore 2016). Conversely, principal reduction is more common for loans held on banks' own balance sheets, where servicer-lender frictions are mitigated (Agarwal et al. 2011).

[20]Online Appendix Figure 8a normalizes expenditure to zero at modification date in order to more clearly show the parallel pre-trends. Online Appendix Figure 8b plots mean auto expenditure around modification date and similarly demonstrates parallel pre-trends.

signed to capture any residual heterogeneity between treatment and control groups.[21] These characteristics x_i are interacted with the $Post_t$ variable to allow for borrower characteristics to explain changes in underlying trends after modification $\left(x'_{it} = (x_i \quad x_i \times Post_t)'\right)$.

4.2 Results: Effect of Principal Reduction on Consumption

We find that neither credit card nor auto expenditures are affected by principal reduction in the year after modification. Our main results are reported in Panels A and B of Table 2. In both panels, column (1) reports the most sparse specification, while columns (2)-(6) add in additional fixed effects and controls. Across all specifications, the treatment effect of principal reduction on both monthly credit card and auto expenditure is small and statistically insignificant. Our preferred estimate using equation (3) is in column (6), which includes MSA-by-calendar-month fixed effects and interacts control variables with a post-modification dummy. In this specification, our point estimate is that principal reduction of \$67,000 increases borrower monthly credit card expenditure by \$2 and auto spending by \$11.

Robustness – We address two potential weaknesses of the credit bureau data by confirming that the result also holds in the JPMCI bank dataset. The first potential weakness is that credit card expenditure is inferred from other variables reported by servicers, as discussed in Section 2.1. The second is any measurement error introduced by our matching procedure. The JPMCI dataset covers only one servicer but does not suffer from either of these two potential limitations. It includes credit card data but not auto loan data. Online Appendix Figure 9 shows that the same pattern of credit card expenditure around modification date holds in the JPMCI data. Our estimated treatment effects are displayed in online Appendix Table 3. Here again we find the treatment effect of debt forgiveness on credit card expenditure is small and statistically insignificant.

We also explore the effect of principal reduction on consumption using our RD strategy. Our outcome variables are the change in mean credit card and auto spending from the 12 months before modification to the 12 months after modification. The reduced form plots are shown in online Appendix Figure 10. These plots confirm the weakness of this strategy for studying consumption impacts since the strategy suffers from lack of precision.[22]

We are unable to analyze the long-run effects of principal reduction on consumption within our sample window. We discuss potential long-run effects in Section 6.2.

Effect of Payment Reduction on Consumption – A natural concern with our

[21] This includes the predicted gain to lenders from providing principal reduction, the predicted gain interacted with a dummy variable equal to one when the gain is positive, borrower characteristics (credit score, monthly income, non-housing monthly debt payment), pre-modification loan characteristics (LTV, principal balance, PTI, monthly payment), property value, origination LTV, and monthly payment reduction. By controlling for the predicted gain to lenders of providing principal reduction, the main difference between our RD and difference-in-differences strategies is that the RD strategy instruments for treatment with the jump in the probability of receiving principal reduction at the cutoff while the difference-in-differences strategy uses all the variation conditional on the running variable.

[22] Translating these estimates to a marginal propensity to consume, as in Section 4.3, our confidence interval ranges all the way from -15 cents to 41 cents.

zero result is that our consumption series might not detect responses to important financial changes. However, the paths of credit card and auto spending around modification suggest that borrowers do seem to respond to short-term payment reductions. Both credit card and auto spending are declining before modification and recover after modification. The decline pre-modification is likely a result of financial stress experienced by the borrowers. The slope of expenditure changes sharply around modification, suggesting that lower payments help expenditure to recover.

The apparent positive effect of short-term payment reductions on auto spending is consistent with findings in Agarwal et al. (2017*a*). That paper exploits regional variation in the implementation of HAMP to estimate the effects of HAMP modifications which combine both short-term and long-term payment reductions. They find that the combined modifications are associated with increased auto spending. If the effect of long-term payment reductions in HAMP is zero, as suggested by our estimates, it makes sense to infer that short-term payment reductions are responsible for the consumption impact they estimate. In online Appendix B.2.2 we attempt to directly estimate the impact of short-term payment reductions on consumption using the payment-reduction RD identification strategy in Section 5, but we conclude that this strategy is under-powered for studying consumption impacts.

4.3 Economic Significance: The MPC From Principal Reduction

To help interpret the economic significance of our results, we convert our estimate for the impact on credit card and auto consumption into an MPC out of principal reduction. First, we scale up credit card spending to a measure of non-auto retail spending to be comparable to Mian, Rao and Sufi (2013). We do this by adjusting for credit card spending on cards where spending is not reported in the credit bureau data and then multiplying by the ratio of non-auto consumer retail spending to consumer credit card spending in 2012.[23] Second, we combine with our auto spending measure, annualize, and divide by the mean incremental amount of principal reduction in the treatment group.

Using this method, our point estimate is that households increased annual consumption by an insignificant 0.3 cents per dollar of principal reduction, with the upper bound of the 95 percent confidence interval corresponding to 0.9 cents. If we normalize by the reduction in the NPV of mortgage payments owed under the new mortgage contract rather than the dollar value of principal reduction, we get a point estimate of 0.7 cents and an upper bound of 1.8 cents. Our estimate of the MPC out of principal reductions for underwater borrowers (which affect wealth but not liquidity) is thus an order of magnitude smaller than typical estimates

[23]Specifically, our adjustment factor is the product of two ratios: (1) the ratio of the number of credit cards in TransUnion to the number of credit cards with spending reported in TransUnion, and (2) the ratio of non-auto consumer retail spending in 2012 to total consumer credit card spending in 2012. The first term uses our data, retail spend is from Census, and consumer credit card spending is from Federal Reserve Payment Study (U.S. Federal Reserve System 2014). This gives an adjustment factor of $(1.2)\,(2.5) = 3.1$. An alternative adjustment multiplying by the ratio of average household monthly non-auto retail spend to the average credit card spending we observe in our sample gives the same MPC point estimate.

of the MPC out of housing wealth increases (which affect both wealth and liquidity). We interpret the lessons from this result in Section 6.2.

5 Effect of Payment Reduction on Default

In this section we analyze the effect of liquidity provision on borrower default. In contrast to our results on the ineffectiveness of principal reduction, we find that short-term payment reduction with no change in long-term obligations significantly reduces default.

5.1 Variation in Payment Reduction Between Government-Subsidized and Private Modifications

We analyze the effect of short-term payment reduction by comparing borrowers with government-subsidized HAMP modifications to those with alternative private modifications. Although servicers were required to offer HAMP modifications to all eligible borrowers, as described in Section 3.1 not all borrowers were eligible.

To mitigate losses on loans ineligible for HAMP, lenders developed their own modification programs. During the Great Recession, mortgages could be partitioned into two approximately equally-sized groups. Loans which met certain underwriting criteria, including a maximum loan size and a minimum borrower FICO score, were usually owned or securitized by the GSEs. Loans which did not meet these criteria were usually underwritten and often securitized by other market actors, such as banks. We analyze borrowers receiving modifications designed by both types of mortgage owners using the JPMCI bank dataset described in Section 2.2. This sample includes both GSE and non-GSE borrowers whose mortgages are serviced by Chase and hence were eligible either for a modification designed by the GSEs or a modification designed by Chase. As we describe below, the GSE modification and the Chase modification are quite similar. For simplicity, we refer to both these types of non-HAMP modifications as "private" modifications.[24]

The design of HAMP and the private modifications we study reflect different views about the most effective way to reduce defaults. HAMP was designed with an explicit *31 percent Payment-to-Income (PTI) ratio target,* as we mention in Section 3.1. This target evolved from the National Housing Act of 1937, which established a PTI limit in the federal government's public housing program. Adopting this income ratio target in a modification program assumes that borrowers with high PTI ratios must need much larger payment reductions in order to avoid subsequent default than borrowers with lower PTI ratios.[25]

The first important feature of the alternative private modification programs is that, in contrast to the rigid PTI ratio target in HAMP, they use a *payment-reduction target.*[26] Under

[24] In practice, the distinction between "public" and "private" modifications is not black and white, especially for GSE-designed modifications which were developed in part while the GSEs were in federal conservatorship.

[25] Another interesting dimension of using income ratio targets to modify debt contracts is that these targets may have implications for labor supply (Ji (2018), Mulligan (2009)).

[26] Chase private modifications target a 30 percent payment reduction. GSE private modifications did not have a single payment reduction target, but Farrell et al. (2017) show that the GSE program effectively

this alternative view, the amount of payment reduction relative to the status quo is at least as important for reducing defaults as ensuring that the new payment is an "affordable" share of income according to the government's metric. For borrowers with PTI near the 31 percent target, private modifications result in immediate payment reductions that are *larger* than HAMP. (Conversely, for borrowers with pre-modification PTI greater than approximately 42 percent, modifications with only a payment reduction target of 30 percent would result in *smaller* payment reductions than HAMP).

The second important feature of these private modification programs is that they use maturity extension as a low-cost tool for achieving deeper immediate payment reductions without reducing long-term obligations. A range of contract terms can be modified in order to achieve a given amount of immediate payment reduction. As described in Section 3.1, HAMP reduces payments by first providing interest rate reductions, then maturity extensions, and finally principal forbearance as necessary to achieve the 31 percent PTI ratio target. In contrast, the private modifications we study target larger payment reductions by first extending maturity and amortization terms (which we call maturity extension for simplicity). In this way, deeper immediate payment reductions are offset by continued payments in the long-term.

Both of these distinctive features are visible in Figure 4 for Chase private modifications, which is our preferred estimation sample. (The modification design and point estimates for GSE private modifications are very similar, but the RD design has a technical issue which we discuss later in this section.) This figure summarizes payments under each program using the same plot structure as Figure 1. Figure 4a shows the annual path of payments relative to the status quo for government and private modifications offered to borrowers with PTI close to 31 percent. Government (HAMP) modifications for this sample have a modest reduction in payments until year 27, with small increases thereafter. Private modifications lead to deeper payment reductions for the first 22 years and then to higher payments for the remaining 18 years.

Figure 4b shows that compared to government modifications, private modifications offer a relative reduction in immediate mortgage payments but leave the NPV of total payments owed approximately unchanged. This analysis uses the same methodology as in Section 3.1. The main force driving the deeper payment reductions is maturity extension, and when the discount rate equals the interest rate (consistent with the empirical evidence of borrower behavior in Busse, Knittel and Zettelmeyer 2013), there is no change in the NPV of total payments owed. Because the private modification has deeper payment reductions for 22 years, it may be NPV-positive from the perspective of a more impatient borrower. Even for such borrowers, we attribute the causal impact of the private modification to its liquidity component because in Section 3 we find that future mortgage payments do not affect default

targeted a payment reduction of about 25 percent. While the HAMP program ended in December 2016, both the Chase and GSE modification programs described in this section are ongoing, and their rules are subject to change. Our description of the programs reflects guidelines in place during our analysis period.

rates.

The mix of other modification characteristics also changed at the cutoff, though these changes roughly offset in terms of NPV. As shown in the figure, borrowers receiving private modifications had less principal reduction. This is because private modifications are ineligible for the subsidized principal reduction in HAMP.[27] On the other hand, private modifications had more principal forbearance and slightly deeper interest rate reductions, offsetting the effect of reduced principal forgiveness from an NPV perspective.

5.2 Identification: Discontinuity in Payment Reduction at the 31 percent Payment-to-Income Ratio

We use variation generated by HAMP's 31 percent PTI eligibility cutoff and the different payment reduction targets used in HAMP and the private modification program. Since HAMP was designed to reduce monthly payments to 31 percent of a borrower's income, borrowers whose PTI ratio was already below 31 percent were ineligible for HAMP. These borrowers were only evaluated for private modifications.

The difference in modification program rules generates substantial variation in the amount of immediate payment reduction received by borrowers on either side of HAMP's 31 percent PTI eligibility cutoff. Below the cutoff, all borrowers receive private modifications with large payment reductions. Above the cutoff, about half of borrowers receive HAMP modifications with small payment reductions (since these borrowers are already close to HAMP's 31 percent PTI target), and about half receive private modifications with large payment reductions (since Chase had a minimum 30 percent payment reduction target). Although borrowers above the cutoff who receive HAMP modifications with a 31 percent PTI ratio target would have received larger payment reductions in the private program, HAMP rules prohibited servicers from offering private sector alternatives to any HAMP-eligible borrower. Thus the only borrowers above the cutoff receiving private modifications are those failing to meet one of the other eligibility criteria described in Section 3.1. The PTI cutoff therefore serves as an instrument for allocating borrowers between HAMP modifications with small payment reductions and private modifications with large payment reductions.

We make three sample restrictions for ease of exposition, though our central result is unchanged when we lift all three restrictions. First, we drop 241 observations between 31 percent and 31.1 percent PTI who receive an amount of payment reduction about halfway between that received by borrowers clearly above and clearly below the cutoff. Second, to focus on a subsample eligible for the full maturity extension treatment, we subset to mortgages whose current terms are less than or equal to 30 years. This maintains 89 percent of the sample. Finally, to simplify the calculation of expected payments before and after

[27] As we discuss above, because we found no independent impact of principal reductions on default in Section 3.3, we attribute the causal impact of this treatment on default to the immediate payment reduction portion of the treatment. In Section 5.3 we provide a bound on our estimate of the effect of payment reduction under the alternative assumption that the relative principal increase actually had an offsetting effect on default.

modification, we restrict to fixed rate mortgages. This maintains 48 percent of observations.

We provide the same tests as for our prior principal reduction RD strategy. Table 1b shows that borrowers in our sample are broadly representative of underwater delinquent borrowers during the recent crisis and similar on observables to the principal reduction sample. Online Appendix Figure 11 shows that predicted default rates based on pre-determined covariates trend smoothly through the cutoff. We provide additional detail on covariate balance in online Appendix B.2.1. Finally, online Appendix Figure 12 shows that borrower density is also smooth around the cutoff. The lack of bunching indicates that there is no manipulation of the running variable.

5.3 Results: Effect of Payment Reduction on Default

Figure 5a shows that borrowers below the cutoff receive payment reductions that are substantially more generous than those received by borrowers above the cutoff. Average payments fall by 32 percent below the cutoff and by only 13 percent above the cutoff. The figure also shows that payment reductions are approximately constant below the cutoff, consistent with the *payment reduction target* discussion above, and that payment reductions are increasing in PTI above the cutoff, consistent with the *PTI ratio target* discussion above. The difference in financial impacts at the cutoff are similar to those between government and private modifications discussed above: at the cutoff there is a sharp drop in immediate monthly payments with no significant change in the NPV of total payments owed.[28]

We use a fuzzy regression discontinuity strategy. The running variable V is the PTI ratio. Similar to equation (2), the estimand for the effect of a 1 percent payment reduction is given by

$$\tau = \frac{lim_{v\downarrow 31\%} E[Y|V=v] - lim_{v\uparrow 31\%} E[Y|V=v]}{lim_{v\downarrow 31\%} E[\Delta Pay|V=v] - lim_{v\uparrow 31\%} E[\Delta Pay|V=v]}. \tag{4}$$

We follow the same procedures as in Section 3.2 to estimate $\hat{\tau}$ using borrowers in the JPMCI bank dataset. The sample includes borrowers with PTI ratios between 25 and 80 whom we observe for at least two years after modification, and we define default as being 90 days delinquent at any point within these two years.[29] For symmetry around the 31 percent PTI cutoff, Figure 5 plots data for borrowers with PTI ratios between 25 percent and 37 percent.

We find that immediate payment reduction significantly reduces default rates. Figure 5b shows the reduced form, plotting the default rate on the y-axis. The figure provides visual evidence that payment reduction reduces default in two ways. First, the default rate falls sharply by 7.3 percentage points relative to a control mean of 32.1 percent at the eligibility cutoff. Second, the slope of the estimated default rates in the bottom panel mirrors the slope of the payment reductions in the top panel; default rates are approximately constant on the left-hand side of the cutoff, consistent with a constant amount of payment reduction, and

[28] Online Appendix Figure 13 reports the change in the NPV of payments owed, the amount of principal forgiveness, and the change in the interest rate.

[29] Borrower density thins out above PTI ratios of 80, and borrowers with PTI ratios below 25 are evaluated according to different program rules.

are falling on the right-hand side of the cutoff, consistent with a rising amount of payment reduction. This pattern is even more striking in online Appendix Figure 14, which plots the first stage and reduced form for borrowers with PTI ratios as high as 80 percent. Our point estimate of $\hat{\tau}$ from equation (4) shows that an extra one percent payment reduction reduces default rates in the two years after modification by 0.38 percentage point, or by 1.2 percent of the mean above the cutoff.

Our result on the effect of *immediate* liquidity provision is of similar magnitude to the effects found in two types of prior work analyzing the effect of *sustained* payment reductions. First, Agarwal et al. (2011) and Haughwout, Okah and Tracy (2016) analyze modifications provided to delinquent borrowers before the implementation of the HAMP program. They find that a 1.0 percent immediate payment reduction that also reduces long-term debt obligations is associated with a 0.9 percent to 1.3 percent reduction in default rates. Second, other authors have analyzed the effect of sustained payment reductions for non-distressed borrowers. Although differences in borrower characteristics and baseline default rates make it difficult to directly compare magnitudes between distressed and non-distressed borrowers, this literature has found that a 1.0 percent payment reduction is associated with a 1.1 percent to 2.0 percent reduction in default rates.[30] Because we find similar impacts from a reduction in *only* immediate payments, immediate liquidity provision may be a main driver of the default reductions documented in previous work.

Robustness – Online Appendix Figure 14 shows that when we remove all three of the sample restrictions described above we find that an extra one percent payment reduction reduces default rates by 0.25 percentage point. This reduction is 1.0 percent of the default rate above the cutoff, which is similar to our estimate of 1.2 percent in the baseline sample. Online Appendix Figure 15 shows that our estimates are stable for a variety of alternative bandwidths. Our point estimate is similar (rising to 1.4 percent) if we use the most optimistic point in our confidence interval from Section 3.3 to adjust for any potential effect of the relative increase in mortgage principal at the cutoff. Online Appendix B.2.1 provides additional detail on this calculation.

We also analyze the effects of immediate payment reduction for GSE-backed loans and find a similar reduction in delinquency. For borrowers ineligible for HAMP, the GSEs offered

[30]See Tracy and Wright (2016), Fuster and Willen (2017), and DiMaggio et al. (2017). These papers analyze the effect of payment reductions caused by downward adjustments of interest rates for borrowers with adjustable rate mortgages. As Fuster and Willen (2017) discuss, to the extent that borrowers are aware of their mortgage terms and follow the movement of underlying index rates, then these estimates may already be capturing only the liquidity effects of lower payments since borrowers would already have been anticipating and responding to the long-term payment reductions. Similarly, Scharlemann and Shore (2019) analyze the effect of sustained payment increases for non-distressed borrowers in a setting where these payment increases are pre-determined five years ahead of time (resulting from a step up in interest rates as part of a previous mortgage modification). The four aforementioned papers capture the effect of payments on default for the average borrower. Other work looking at the effect for borrowers that choose to refinance by Ehrlich and Perry (2015) and Abel and Fuster (2018) has found larger elasticities. Finally, this research examining the impact of mortgage payment reductions is complemented by Hsu, Matsa and Melzer (2018), who show that liquidity provided by unemployment insurance benefits also helps reduce mortgage default.

a private modification that was very similar to the Chase private modification in that it featured a payment reduction target and for most borrowers used maturity extensions prior to any interest rate reduction. Borrowers on the left-hand side receive payment reductions that are 22.4 percentage points more generous than borrowers on the right-hand side and have default rates that are 5.9 percentage points lower.[31] We estimate that a one percent reduction in payments reduces default by 0.26 percentage point. This reduction is 1.05 percent of the default rate above the cutoff, which is similar to our estimate of a reduction of 1.2 percent for non-GSE loans. It is reassuring that we recover essentially the same point estimate when using a completely different sample of borrowers.

However, the research design for the GSE-backed loans does not pass one of the standard RD identification checks and hence it is not our preferred specification. There are more borrowers on the *right-hand* side of the cutoff than on the left-hand side, as shown in online Appendix Figure 19. This issue does not arise because of borrower manipulation of PTI, which would induce more mass on the *left-hand* side where additional payment reduction is available. Rather, it arises because eligibility for GSE private modifications required a FICO lower than 620 for borrowers less than 60 days delinquent at the time of modification, while HAMP had no such screen. As a result, there is an imbalance: the left-hand side has no borrowers who are less than 60 days delinquent with FICO above 620, while the right-hand side has borrowers of all FICO scores and pre-modification delinquency statuses. We are unable to correct this problem because we do not observe FICO in the JPMCI data. Nevertheless, we believe the bias from these extra low-delinquency higher FICO borrowers is small and therefore find this research design to still be informative.[32] Thus, we find that payment reduction as implemented through both the Chase and GSE private alternatives to HAMP has a substantial effect on default for a wide range of types of borrowers.

6 Discussion and Interpretation

This section discusses the positive and normative lessons of our empirical findings for default and consumption.

6.1 Default

In our sample immediate liquidity substantially affects default, but reduction in mortgage

[31]The sample of GSE modifications covers the same 2011 through 2016 time period as the sample of Chase modifications. Online Appendix Figures 16, 17, 18 and 19 replicate our 31 percent PTI cutoff analysis for the GSE-backed loans sample. Recall that in online Appendix Figure 13 we showed that there was no change in NPV of payments owed at the cutoff for the non-GSE-backed sample because more interest rate reduction was offset by less principal reduction. Fannie Mae and Freddie Mac did not allow principal forgiveness (even for HAMP recipients) and so in the GSE-backed sample the reduction in the NPV of payments owed is smaller on the right-hand side.

[32]Two pieces of evidence suggest the bias in the point estimate is small. First, we plot predicted delinquency as a function of average observable borrower characteristics in online Appendix Figure 18. There is little change in predicted delinquency at the cutoff. Second, the marginal borrowers on the right-hand side are more creditworthy. This change in sample composition will lower average default rates on the right-hand side, leading us to understate the effectiveness of payment reduction.

principal does not. From a positive perspective, this is consistent with liquidity-focused explanations for default and contrasts with explanations for default as a response to negative equity.

However, there are at least three reasons to be cautious about extrapolating the "liquidity drives default" conclusion to other contexts: the treatment was first implemented in 2010, it did not bring borrowers above water, and few borrowers started extremely far underwater. First, it is possible that defaults prior to 2010 were more responsive to negative equity. For example, there is compelling evidence that speculators were instrumental in the run-up to the crisis in driving both house price dynamics and the initial rise in default rates.[33] The default decision by such speculators early in the crisis may have been driven by different forces than the decisions of households in our sample. Overall, Haughwout et al. (2011) estimate that investors accounted for at most one-third of defaults during this early period. Second, more generous principal reductions that completely eliminate negative equity—unlike the one we study which left borrowers underwater—may provide liquidity and may mechanically reduce defaults by allowing borrowers to sell their homes (Gupta and Hansman 2019). However, as we note in Section 3.3, forgiving 100 percent of negative equity would be an expensive way to reduce foreclosures unless it is targeted to borrowers nearly certain to default in its absence. Third, our results are only valid for borrowers in our analysis sample, and 90 percent of our sample has pre-modification LTV ratios below 168. This is because there were actually relatively few borrowers who were this deeply underwater during the Great Recession (see Table 1). However, Bhutta, Dokko and Shan (2017) find that it is this small share of extremely underwater borrowers who are most likely to respond to negative equity. For all of these reasons, we think that more investigation about the general prevalence of negative-equity driven default is an important topic for future research.

From a normative perspective, our key conclusion is that a policy that focuses on reducing payments can be superior to one that focuses on reducing principal. In our sample, principal reduction is ineffective for borrowers and costly to both lenders and taxpayers. Even at the most optimistic point in our confidence interval, taxpayers spent at least $365,000 per avoided foreclosure, far larger than common estimates of the social costs of foreclosures (U.S. Department of Housing and Urban Development 2010). In addition, we estimate in Section 3.3 that lenders lost at least $402,000 per foreclosure avoided.

In contrast, payment-focused modifications are able to successfully reduce defaults for borrowers—at zero cost to taxpayers and at *negative* cost to lenders. Prioritizing maturity extension—before changing other mortgage terms—enables these modifications to offer larger short-term payment reductions that are offset by continued payments in the long term. In Section 5, we calculate that moving a borrower from the right side of the HAMP eligibility discontinuity to the left side reduces borrower default probability by seven percentage points.

[33]See e.g., Albanesi, De Giorgi and Nosal (2017), DeFusco, Nathanson and Zwick (2017), Gao, Sockin and Xiong (2018), Mian and Sufi (2018), Nathanson and Zwick (2018).

Lower default rates indicate by revealed preference that borrowers find this contract that maximizes immediate liquidity provision more attractive than the alternative contract with less generous payment reductions.[34] For lenders, moving a loan from the right side to the left *increases* the NPV of expected payments received by about $6,000. This is because the reduction in defaults more than offsets the higher discount rate used for cash flows pushed further in the future. In addition, loans that receive private modifications require no taxpayer subsidy, so altogether moving a borrower across the cutoff is likely a Pareto improvement for borrowers, lenders, and taxpayers. We provide more detail on our calculations of borrower and lender benefits in online Appendix C.3.

The result that borrowers, lenders, and taxpayers are all better off from payment-focused modifications raises two questions. First, why didn't the private sector provide these alternative modifications to HAMP borrowers? After all, the government spent substantial resources subsidizing HAMP modifications above the 31 percent PTI cutoff with small payment reductions and high default rates. If alternative modifications existed that were better for borrowers and lenders and free for taxpayers, it seems natural that the private sector would have offered them instead. However, the key friction was that the government explicitly prohibited participating servicers from offering alternative private modifications to any HAMP-eligible borrower. One potential motivation for the requirement to offer "HAMP first" may have been an assumption that the private modifications would be less effective.[35] However, the PTI target in HAMP means that borrowers with pre-modification PTI ratios below 42 percent actually received payment reductions *smaller* than the payment reductions offered in private modifications. About 40 percent of all HAMP borrowers (625,000 borrowers) were in this region, and the government spent approximately $7 billion subsidizing potentially less effective modifications for them. Thus, the HAMP-first requirement may have crowded out private modifications that could have been more effective for a large fraction of HAMP borrowers.[36]

Second, how much could have been gained if modification programs had been redesigned to focus on immediate liquidity provision? For borrowers *at our RD cutoffs*, default is responsive to immediate liquidity but not to changes in total long-term debt obligations. Assessing the potential gains from re-designing modification programs requires extrapolating

[34] Our empirical results show that defaults are lower for at least two years after modification. Concluding that borrowers are better off implicitly assumes that treatment does not raise defaults outside of the time horizon we observe in the data. This assumption could fail because maturity extension slows the repayment of principal. A borrower who is underwater and defaults might end up in foreclosure. A typical borrower on the left-hand side of the 31 percent RD will be underwater for two additional years as a result of treatment. We provide calculations in online Appendix C.3.2 suggesting that the lifetime foreclosure rate is about 7 percent lower in the treatment group. Thus, the default-reducing benefit of additional payment reduction appears to outweigh the default-increasing risk from additional years spent underwater.

[35] Another potential justification was a concern that the private sector would cherry-pick which loans to send to HAMP in order to maximize their private returns.

[36] While our emphasis is on crowd-out on the intensive margin of the quality of modifications, Agarwal et al. (2017*a*) examine the extensive margin in terms of the quantity of modifications and find little evidence of crowd out.

these treatment effect estimates beyond the specific cutoffs we study. Table 1 shows that borrowers at these cutoffs are similar to typical delinquent borrowers in this time period, so on the basis of observable characteristics such an extrapolation may be warranted.

One way to assess whether the *causal* effects are likely to extrapolate is to examine the *cross-sectional* relationships between payment reduction, principal reduction, and default away from the cutoffs. For payment reduction, our finding that treatment reduces default would need to extrapolate to borrowers at other PTI levels. Additional payment reduction is associated with a further reduction in default for borrowers with PTI as high as 60 percent, as shown in online Appendix Figure 14. For principal reduction, our finding that treatment does not reduce default would need to extrapolate to borrowers with very low or high values of the running variable. Borrowers with high values of the running variable received almost two times as much principal reduction as borrowers barely on the right-hand side of the cutoff, as shown in online Appendix Figure 6b. Yet Figure 2b shows no decline at all in default rates at high values of the running variable. Thus, the cross-sectional variation away from the cutoff is consistent with the view that payment reduction is effective and principal reduction is ineffective throughout the sample. However, these borrowers may differ on unobservables, and so we caution that our estimates of the gains from redesigning modifications rely on an extrapolation assumption.

If such an extrapolation is valid, our results suggest a simple rule: the efficient modification structure should maximize short-term liquidity provision, in the spirit of Eberly and Krishnamurthy (2014). The costs of payment reductions must be borne by either lenders or taxpayers. Minimizing costs per dollar of immediate liquidity provision will maximize the amount of payment reduction (and hence default reduction) achieved for a given quantity of public or private funds. We use our empirical estimates to evaluate the cost of each modification step used by the programs during the crisis. This calculation suggests a hierarchy for achieving a given amount of payment reduction: the efficient modification would first use maturity extension, followed by temporary interest rate reduction, followed by principal forbearance, and never use principal forgiveness, as shown in online Appendix Figure 20.

We quantify the potential gains if these more efficient modifications had been offered to all HAMP borrowers. We examine two distributional extremes: allocating all the gains to lenders/taxpayers and allocating all the gains to borrowers. First, we calculate that the median amount of payment reduction in HAMP could have been provided at a $67,000 lower cost per modification to lenders and taxpayers.[37] Alternatively, if we allocate all the gains from redesigning modifications to reducing borrower payments, the same amount of lender and taxpayer cost can be used to achieve substantially more default reduction. We calculate that it would have been possible to cut default rates by one-third, avoiding 267,000 defaults in HAMP at *no additional cost* to lenders or taxpayers.[38] Of course, the potential gains

[37]This translates into a total potential unnecessary cost of $121 billion aggregating over all HAMP modifications.

[38]We provide more details on this calculation in online Appendix C.4. In this appendix we also compare

would be smaller if our results do not extrapolate to all 1.8 million HAMP borrowers, while they would be larger if they also held for all of the 10 million private and public modifications completed during the Great Recession.

6.2 Consumption

Our consumption results help shed light on the mechanisms underlying the robust relationship between housing wealth and consumption. A large literature examines the consumption response to house price changes and typically estimates an MPC of around five cents per dollar.[39] Two main explanations have been advanced for this relationship, as summarized in Cloyne et al. (2019). First, consumption responses could reflect an increase in wealth.[40] Second, they could reflect a relaxation of collateral constraints. Because house price changes typically affect both wealth *and* collateral, it has been difficult to separate these effects.

Our setting allows us to distinguish between the wealth and liquidity-based explanations for housing MPCs. Only positive home equity can be used as collateral. Thus, a reduction in mortgage principal that leaves a borrower underwater increases that borrower's NPV of wealth (by reducing the NPV of their debt obligations) but does not relax their collateral constraint. Hence, our setting isolates the wealth channel holding the collateral channel fixed.

Because we find that the MPC from principal reduction is effectively zero, our results suggest that the wealth channel is weak and that relaxing collateral constraints is necessary for housing wealth to stimulate consumption. To our knowledge, ours is the first estimate of the consumption response to a wealth change that is not accompanied by a change in current liquidity. This estimate complements prior work that has investigated the opposite type of natural experiment: increases in access to housing collateral with no change in wealth.[41] This literature finds that the collateral channel can drive substantial responses to changes in home equity. Together, these results suggest that relaxing collateral constraints is not just a *sufficient* but also a *necessary* condition for housing wealth changes to affect consumption.

This zero MPC finding has lessons both for models and for policy. From a modeling perspective, our results provide evidence that the timing of liquidity matters. This is a key implication of incomplete market models with borrowing constraints.[42] A substantial literature has implemented tests for incomplete markets by showing that current consumption

the cost of private modifications to the cost of our proposed efficient modification.

[39]See e.g. Aladangady (2017), Campbell and Cocco (2007), Carroll, Otsuka and Slacalek (2011), Guren et al. (2018), and Mian, Rao and Sufi (2013).

[40]In the context of house price changes, it is unclear whether increases in nominal wealth reflect increases in real wealth. Because higher house prices compensate households for higher implicit rental costs, house price increases are more likely to reflect increases in real wealth for older homeowners more likely to downsize (Sinai and Souleles (2005), Campbell and Cocco (2007)). In our context, principal forgiveness translates into real wealth for any household who pays off the principal but not households who immediately re-default.

[41]See Agarwal and Qian (2017), Cloyne et al. (2019), Defusco (2017), and Leth-Petersen (2010). See also Berger, Turner and Zwick (2016), Argyle, Nadauld and Palmer (2019), and Fadlon, Ramnath and Tong (2019) who show respectively that home purchase, auto purchase, and labor supply decisions are very sensitive to liquidity.

[42]See Berger et al. (2018) for a recent example of such a model investigating house prices and consumption.

responds to *current* liquidity (e.g. Johnson, Parker and Souleles 2006; Zeldes 1989). We provide complementary evidence by showing that current consumption is *unresponsive* to changes in *future* liquidity.[43]

To provide more formal support for this discussion, we build a partial equilibrium life-cycle model of consumption and default. We leave the details of the model to online Appendix D and discuss the main findings here. The model contains one simple addition to the standard life-cycle consumption model in Carroll et al. (2018): households own a home with a long-term mortgage and can only borrow against their home equity subject to a collateral constraint. When households are far underwater, they are far from the point where home equity can be monetized.[44]

The inability to access liquidity can indeed explain why principal reduction fails to increase consumption in this type of simple incomplete markets model. One way to investigate this effect is to compare the consumption response to one dollar of cash versus one dollar of housing wealth gained by principal reduction, as shown in Figure 3b. As in prior empirical results (Mian, Rao and Sufi 2013), borrowers near their collateral constraint have a high MPC out of housing wealth gains. (In fact, this was one strong motivation for the policy interest in principal reduction.) However, borrowers far underwater are unresponsive to housing wealth changes even though they are highly responsive to cash transfers. The inability to monetize housing wealth drives a wedge between an underwater borrower's MPC out of cash and their MPC out of housing wealth. Housing wealth is special because it can only be monetized above a collateral constraint.

Our result contrasts with debt overhang models in which forced deleveraging leads to depressed consumption. For example, in Eggertsson and Krugman (2012) and Guerrieri and Lorenzoni (2017), debt is modeled as a one-period bond. In this setting, borrowers who find themselves beyond the borrowing constraint are forced to immediately cut consumption in order to delever. If mortgages were short-term loans, underwater borrowers would need to immediately repay their outstanding debt until they were above water. In this scenario, principal reduction would increase consumption by reducing the amount of forced repayment. But with long-term mortgages, nothing forces borrowers to immediately delever when they are far underwater. Modeling housing debt as a long-term contract removes a mechanical link between debt levels and consumption and reduces the expected effectiveness of mortgage debt reduction policies. Other recent papers to consider the effect of debt and housing

[43] Although this body of evidence is consistent with models with incomplete markets where households optimize subject to liquidity constraints, it is also consistent with various behavioral models. For example, the "spender" households in Campbell and Mankiw (1989) or the present-biased liquidity constrained households in Laibson (1997) would also fail to increase spending in response to a principal reduction that had no effect on their immediate disposable income. We cannot rule out other such models where current liquidity plays a central role.

[44] Indeed, because lenders typically require an equity buffer for new loans, borrowers need to go even beyond the 100 percent LTV threshold before being able to monetize housing wealth. In the model, we set the collateral constraint such that borrowers can only borrower against their homes up to an LTV ratio of 80 percent. We provide evidence to support this assumption during our sample period in online Appendix D.2.

wealth in settings with long-term contracts include Berger et al. (2018), Chen, Michaux and Roussanov (2013), Kaplan, Mitman and Violante (2017), and Justiniano, Primiceri and Tambalotti (2015).

From a policy perspective, our results highlight that when borrowing constraints matter for real outcomes, programs can be ineffective if they fail to target these constraints. We find that the relationship between housing wealth and consumption breaks down when borrowers are underwater because collateral constraints continue to bind. Hence principal reduction will fail to stimulate consumption for underwater borrowers because households cannot increase borrowing to monetize these gains. However, Figure 3b suggests that providing direct liquidity to low-wealth, underwater borrowers would successfully stimulate consumption.[45] At this MPC, our model suggests that if principal reduction subsidies had instead been spent on direct transfers to borrowers, the partial equilibrium spending increase would have been ten times more than even the upper bound of our estimates for the consumption response to principal forgiveness.

Although principal reduction as implemented during the crisis is an ineffective way to provide immediate stimulus, there are two reasons to avoid concluding that principal reduction can never affect consumption. First, in the long-term there will likely be some impact on consumption. Eventually, principal reduction translates into lower payments and increased borrowing capacity. Although we cannot analyze this effect empirically within our sample window, our model shows that consumption would likely begin responding about five years after modification for the average borrower. To the extent that principal reductions were pursued for short-run macroeconomic stabilization, long-run consumption responses will have limited benefit.

Second, it is possible that more generous principal reductions that did push borrowers above their collateral constraint might have led to some immediate consumption response. However, this would have been an inefficient way to raise consumption. Figure 3b shows that all of the consumption increase would be coming from the region near the collateral constraint. This means that a policy of targeting deeply underwater borrowers with more generous writedowns will expend substantial resources in a region with no stimulative effect.[46]

7 Conclusion

In this paper, we explore how borrower liquidity and wealth affect default and consumption decisions through the lens of mortgage modifications in the Great Recession. Using quasi-experimental research designs, we show principal reduction that leaves short-term mortgage payments unchanged—wealth without liquidity—has no detectable effect on default or

[45] This high MPC is qualitatively consistent with the sensitivity of spending to payment reduction from our event study graphs in Figure 3a. It is also quantitatively consistent with DiMaggio et al. (2017) who find that underwater borrowers have MPCs out of cash more than twice as high as non-underwater households.

[46] One implication of this result is that a principal reduction policy very early in the crisis—before collateral constraints had tightened and before price declines had pushed borrowers into negative equity—would have had a positive effect on consumption.

consumption. In sharp contrast, we find that short-term payment reduction with no change in the net present value of payments owed—liquidity without wealth—significantly reduces default. Taken together, these results suggest that liquidity is the key driver of consumption and default decisions for borrowers in our sample.

What then have we learned since the financial crisis, when principal reduction was viewed as a promising policy tool? For stimulating consumption, even though highly-leveraged *above water* borrowers have a high MPC, principal reduction is ineffective for *underwater* borrowers because they are unable to access this wealth. For reducing defaults, we show that by focusing on borrower liquidity, distressed debt modifications can be redesigned with substantial potential gains. Specifically, instead of principal reduction, our results suggest policymakers should prioritize reducing current monthly payments. Unfortunately, according to this metric, many private and public mortgage modifications fell short. For example, fewer than half of borrowers with private sector or GSE modifications in 2008 received any payment reduction (Barr 2018). Had they focused on providing short-term liquidity, modifications could have offered substantially more payment reduction to borrowers at no additional cost to investors or taxpayers. Altogether, applied to the main government program for distressed borrowers during the Great Recession, our results imply that 267,000 defaults could have been avoided.

References

Abel, Joshua, and Andreas Fuster. 2018. "How Do Mortgage Refinances Affect Debt, Default, and Spending? Evidence from HARP." Federal Reserve Bank of New York Staff Report 841.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino. 2016. "Loan Origination and Defaults in the Mortgage Crisis: The Role of the Middle Class." *Review of Financial Studies*, 29(7): 1635–1670.

Agarwal, Sumit, and Wenlan Qian. 2017. "Access to Home Equity and Consumption: Evidence from a Policy Experiment." *The Review of Economics and Statistics*, 99(1): 40–52.

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, and Douglas D. Evanoff. 2011. "The Role of Securitization in Mortgage Renegotiation." *Journal of Financial Economics*, 102(3): 559–578.

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

Agarwal, Sumit, Gene Amromin, Souphala Chomsisengphet, Tim Landvoigt, Tomasz Piskorski, Amit Seru, and Vincent Yao. 2017*b*. "Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinancing Program." National Bureau of Economic Research Working Paper 21512.

Aladangady, Aditya. 2017. "Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata." *American Economic Review*, 107(11): 3415–3446.

Albanesi, Stefania, Giacomo De Giorgi, and Jaromir Nosal. 2017. "Credit Growth and the Financial Crisis: A New Narrative." National Bureau of Economic Research Working Paper 23740.

Argyle, Bronson, Taylor D Nadauld, and Christopher Palmer. 2019. "Monthly Payment Targeting and the Demand for Maturity." National Bureau of Economic Research Working Paper 25668.

Barr, Michael. 2018. "Responding to the Global Financial Crisis: What We Did and Why We Did It." *Brookings Institution Conference Transcript*.

Berger, David, Nicholas Turner, and Eric Zwick. 2016. "Stimulating Housing Markets." National Bureau of Economic Research Working Paper 22903.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. 2018. "House Prices and Consumer Spending." *The Review of Economic Studies*.

Bhutta, Neil, Jane Dokko, and Hui Shan. 2017. "Consumer Ruthlessness and Mortgage Default during the 2007 to 2009 Housing Bust." *The Journal of Finance*, 72(6): 2433–2466.

Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer. 2013. "Are Consumers Myopic? Evidence from New and Used Car Purchases." *American Economic Review*, 103(1): 220–256.

Campbell, John Y., and Joao F. Cocco. 2007. "How Do House Prices Affect Consumption? Evidence from Micro Data." *Journal of Monetary Economics*, 54(3): 591–621.

Campbell, John Y., and N. Gregory Mankiw. 1989. "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence." National Bureau of Economic Research, Inc NBER Chapters.

Campbell, John Y., Nuno Clara, and Joao F. Cocco. 2018. "Structuring Mortgages for Macroeconomic Stability." Working Paper.

Carroll, Chris, Nathan Palmer, Matthew N. White, Jacqueline Kazil, and David Low. 2018. "econ-ark/HARK: 0.8.0 (Version pre)."

Carroll, Christopher D., Misuzu Otsuka, and Jiri Slacalek. 2011. "How Large Are Housing and Financial Wealth Effects? A New Approach." *Journal of Money, Credit and Banking*, 43(1): 55–79.

Chen, Hui, Michael Michaux, and Nikolai L. Roussanov. 2013. "Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty." Social Science Research Network SSRN Scholarly Paper ID 2024392, Rochester, NY.

Cloyne, James, Kilian Huber, Ethan Ilzetzki, and Henrik Kleven. 2019. "The Effect of House Prices on Household Borrowing: A New Approach." *American Economic Review*, Forthcoming.

Defusco, Anthony A. 2017. "Homeowner Borrowing and Housing Collateral: New Evidence from Expiring Price Controls." *The Journal of Finance*, 73(2): 523–573.

DeFusco, Anthony A, Charles G Nathanson, and Eric Zwick. 2017. "Speculative Dynamics of Prices and Volume." National Bureau of Economic Research Working Paper 23449.

DiMaggio, Marco, Amir Kermani, Benjamin Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao. 2017. "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging." *American Economic Review*, 107(11): 3550–3588.

Dobbie, Will, and Jae Song. 2019. "Targeted Debt Relief and the Origins of Financial Distress: Experimental Evidence from Distressed Credit Card Borrowers." Working Paper.

Eberly, Janice, and Arvind Krishnamurthy. 2014. "Efficient Credit Policies in a Housing Debt Crisis." *Brookings Papers on Economic Activity*, 49(2): 73–136.

Eggertsson, Gauti B., and Paul Krugman. 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach." *The Quarterly Journal of Economics*, 127(3): 1469–1518.

Ehrlich, Gabriel, and Jeffrey Perry. 2015. "Do Large-Scale Refinancing Programs Reduce Mortgage Defaults? Evidence From a Regression Discontinuity Design." Working Paper.

Fadlon, Itzik, Shanthi P Ramnath, and Patricia K Tong. 2019. "Household Responses to Transfers and Liquidity: Evidence from Social Security's Survivors Benefits." National Bureau of Economic Research Working Paper 25586.

Farrell, Diana, Kanav Bhagat, Peter Ganong, and Pascal Noel. 2017. "Mortgage Modifications after the Great Recession: New Evidence and Implications for Policy." Working Paper.

Fuster, Andreas, and Paul Willen. 2017. "Payment Size, Negative Equity, and Mortgage Default." *American Economic Journal: Economic Policy*, 9(4): 167–191.

Gao, Zhenyu, Michael Sockin, and Wei Xiong. 2018. "Economic Consequences of Housing Speculation." *Working Paper.*

Geanakoplos, John, and Susan Koniak. 2009. "Matters of Principal." The New York Times, March 4, 2009.

Geithner, Timothy. 2014. *Stress Test.* Broadway Books.

Gerardi, Kristopher, Kyle F. Herkenhoff, Lee E. Ohanian, and Paul S. Willen. 2015. "Can't Pay or Won't Pay? Unemployment, Negative Equity, and Strategic Default." National Bureau of Economic Research, Inc NBER Working Paper 21630.

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

Goldfarb, Zachary. 2012. "Economists, Obama Administration at Odds Over Role of Mortgage Debt in Recovery." The Washington Post, November 22, 2012.

Gorea, Denis, and Virgiliu Midrigan. 2017. "Liquidity Constraints in the U.S. Housing Market." National Bureau of Economic Research Working Papers 23345.

Greenwald, Daniel, Tim Landvoigt, and Stijn Van Nieuwerburgh. 2018. "Financial Fragility with SAM?" Working Paper.

Guerrieri, Veronica, and Guido Lorenzoni. 2017. "Credit Crises, Precautionary Savings, and the Liquidity Trap." *The Quarterly Journal of Economics*, 132(3): 1427–1467.

Gupta, Arpit, and Christopher Hansman. 2019. "Selection, Leverage and Default in the Mortgage Market." *Working Paper*.

Guren, Adam M, Alisdair McKay, Emi Nakamura, and Jon Steinsson. 2018. "Housing Wealth Effects: The Long View." *Working Paper*, 87.

Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw. 2001. "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design." *Econometrica*, 69(1): 201–209.

Haughwout, Andrew, Donghoon Lee, Joseph S. Tracy, and Wilbert van der Klaauw. 2011. "Real Estate Investors, the Leverage Cycle, and the Housing Market Crisis." FRB of New York Staff Report No. 514.

Haughwout, Andrew, Ebiere Okah, and Joseph Tracy. 2016. "Second Chances: Subprime Mortgage Modification and Redefault." *Journal of Money, Credit and Banking*, 48(4): 771–793.

Hedlund, Aaron. 2015. "Failure to Launch: Housing, Debt Overhang, and the Inflation Option During the Great Recession." Working Paper.

Holden, Steve, Austin Kelly, Douglas McManus, Therese Scharlemann, Ryan Singer, and John D. Worth. 2012. "The HAMP NPV Model: Development and Early Performance." *Real Estate Economics*, 40(s1): S32–S64.

Hsu, Joanne W., David A. Matsa, and Brian T. Melzer. 2018. "Unemployment Insurance as a Housing Market Stabilizer." *American Economic Review*, 108(1): 49–81.

Imbens, Guido, and Karthik Kalyanaraman. 2012. "Optimal Bandwidth Choice for the Regression Discontinuity Estimator." *Review of Economic Studies*, 79(3): 933–959.

Ji, Yan. 2018. "Job Search under Debt: Aggregate Implications of Student Loans." Social Science Research Network SSRN Scholarly Paper ID 2976040, Rochester, NY.

Johnson, David S, Jonathan A Parker, and Nicholas S Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review*, 96(5): 1589–1610.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2015. "Household Leveraging and Deleveraging." *Review of Economic Dynamics*, 18(1): 3–20.

Kaplan, Greg, Kurt Mitman, and Giovanni L Violante. 2017. "The Housing Boom and Bust: Model Meets Evidence." National Bureau of Economic Research Working Paper 23694.

Kruger, Samuel. 2018. "The Effect of Mortgage Securitization on Foreclosure and Modification." *Journal of Financial Economics*, 129(3): 586–607.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *The Quarterly Journal of Economics*, 112(2): 443–77.

Lee, David S., and Thomas Lemieux. 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature*, 48(2): 281–355.

Leth-Petersen, Soren. 2010. "Intertemporal Consumption and Credit Constraints: Does Total Expenditure Respond to an Exogenous Shock to Credit?" *American Economic Review*, 100(3): 1080–1103.

Mian, Atif, and Amir Sufi. 2011. "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis." *American Economic Review*, 101(5): 2132–2156.

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *The Quarterly Journal of Economics*, 128(4): 1687–1726.

Mian, Atif R., and Amir Sufi. 2018. "Credit Supply and Housing Speculation." Social Science Research Network SSRN Scholarly Paper ID 3209564, Rochester, NY.

Mulligan, Casey B. 2009. "Means-Tested Mortgage Modification: Homes Saved or Income Destroyed?" National Bureau of Economic Research Working Paper 15281.

Nathanson, Charles G., and Eric Zwick. 2018. "Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market." *The Journal of Finance*, 73(6): 2587–2633.

Nielsen, Helena Skyt, Torben Sorensen, and Christopher Taber. 2010. "Estimating the Effect of Student Aid on College Enrollment: Evidence from a Government Grant Policy Reform." *American Economic Journal: Economic Policy*, 2(2): 185–215.

Palmer, Christopher. 2015. "Why Did So Many Subprime Borrowers Default During the Crisis: Loose Credit or Plummeting Prices?" Working Paper.

Piskorski, Tomasz, and Alexei Tchistyi. 2010. "Optimal Mortgage Design." *The Review of Financial Studies*, 23(8): 3098–3140.

Scharlemann, Therese, and Stephen H. Shore. 2019. "The Effect of Mortgage Payment Size on Default and Prepayment: Evidence from HAMP Resets." *Working Paper*.

Scharlemann, Therese C., and Stephen H. Shore. 2016. "The Effect of Negative Equity on Mortgage Default: Evidence From HAMP's Principal Reduction Alternative." *The Review of Financial Studies*, 29(10): 2850–2883.

Sinai, Todd, and Nicholas S. Souleles. 2005. "Owner-Occupied Housing as a Hedge Against Rent Risk." *The Quarterly Journal of Economics*, 120(2): 763–789.

Tracy, Joseph, and Joshua Wright. 2016. "Payment Changes and Default Risk: The Impact of Refinancing on Expected Credit Losses." *Journal of Urban Economics*, 93(C): 60–70.

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

United States Department of the Treasury. 2014*a*. "Making Home Affordable Data File User Guide." United States Department of the Treasury.

United States Department of the Treasury. 2014*b*. "Making Home Affordable Program Performance Report Through The Second Quarter of 2014." United States Department of the Treasury.

United States Department of the Treasury. 2015. "Home Affordable Modification Program Base Net Present Value (NPV) Model v7.0 Model Documentation." United States Department of the Treasury.

United States Department of the Treasury. 2017. "Making Home Affordable: Program Performance Report Through the First Quarter of 2017." United States Department of the Treasury.

United States Federal Reserve System. 2014. "The 2013 Federal Reserve Payments Study." United States Federal Reserve System.

Zeldes, Stephen. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy*, 97(2): 305–46.

Figure 1: Financial Impact of Modifications with and without Principal Reduction

(a) Annual Impacts on Payments

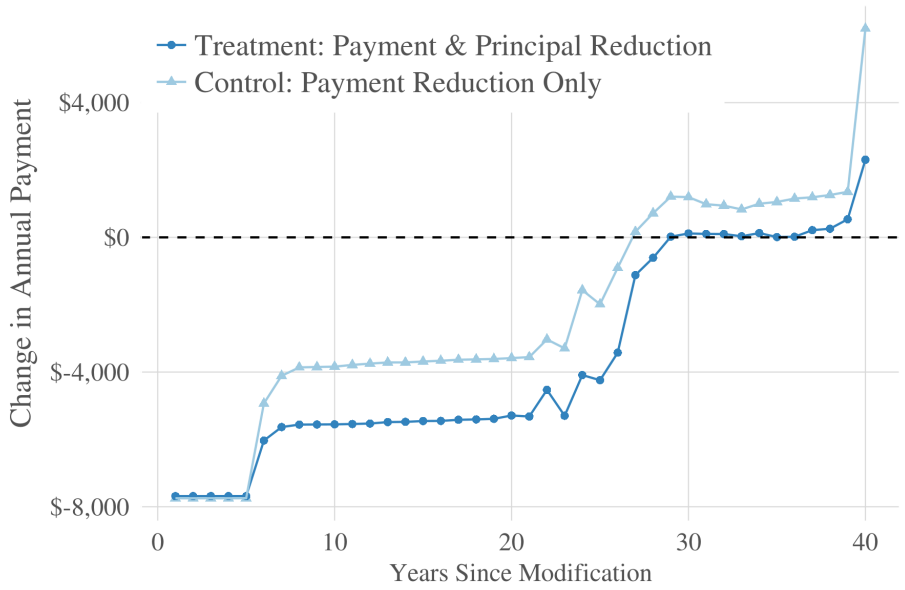

(b) Summary Impact

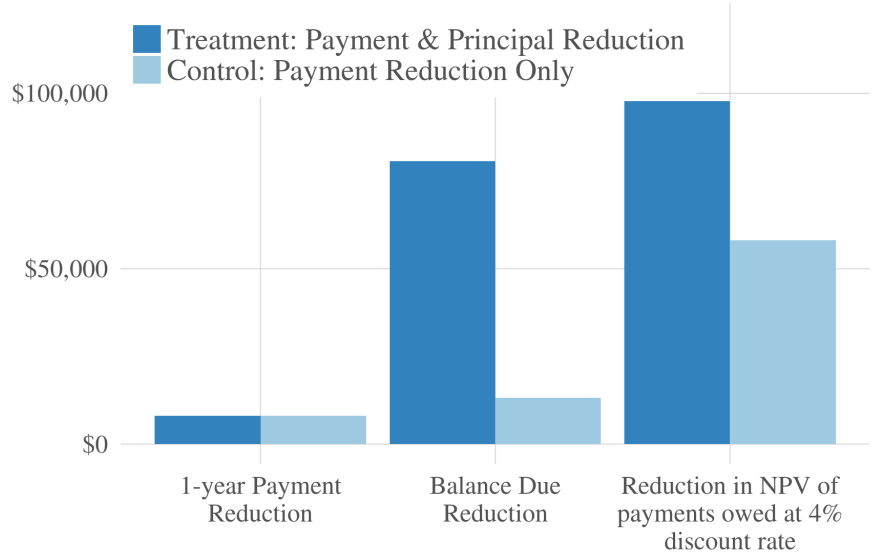

Notes: This figure compares modifications with principal reduction to modifications without principal reduction. Panel (a) plots the difference in average annual payments for borrowers receiving each type of modification relative to the payments borrowers owed under their unmodified mortgage contracts in the matched HAMP credit bureau dataset. The change in payments is winsorized at the 95th percentile; see online Appendix Figure 21 for an unwinsorized version of the same plot. Panel (b) summarizes the financial impacts of modifications along various dimensions: the change in the one-year payment, the change in the unpaid balance, and the change in the net present value of mortgage payments owed, discounted at a four percent interest rate. See Section 3.1 for details.

Figure 2: Effect of Principal Reduction on Default

(a) First Stage -- Receive Principal Reduction

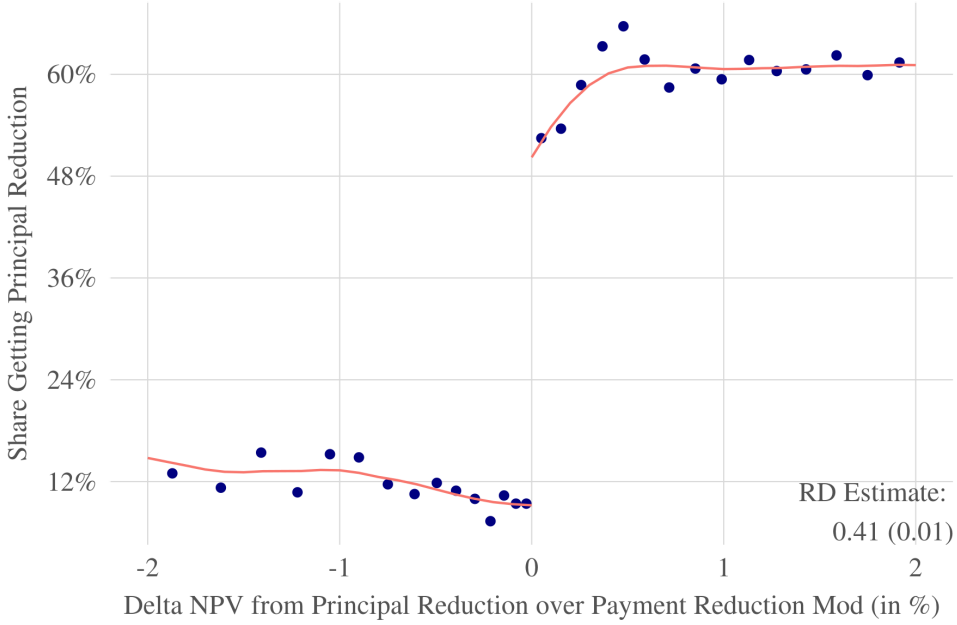

(b) Reduced Form -- Mortgage Default

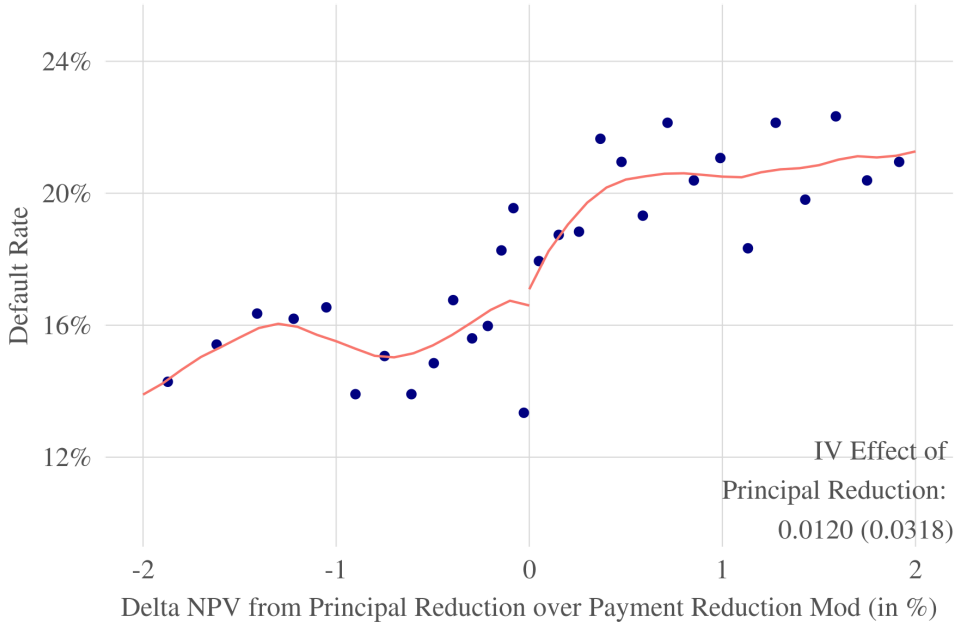

Notes: This figure evaluates the impact of principal forgiveness using a regression discontinuity at the Net Present Value cutoff in the matched HAMP credit bureau dataset. The horizontal axis shows the normalized predicted gain to lenders of providing principal reduction to borrowers from equation (1). The dots are conditional means for 15 bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. Panel (a) plots the share of borrowers receiving principal reduction and panel (b) plots the share defaulting, which is defined as 90 days delinquent between the modification date and March 2015, when our dataset ends. Construction of the IV estimate $\hat{\tau}$ in panel (b) is described in Section 3.2.

Figure 3: Effect of Principal Reduction on Consumption

(a) Credit Card Spend Around Modification

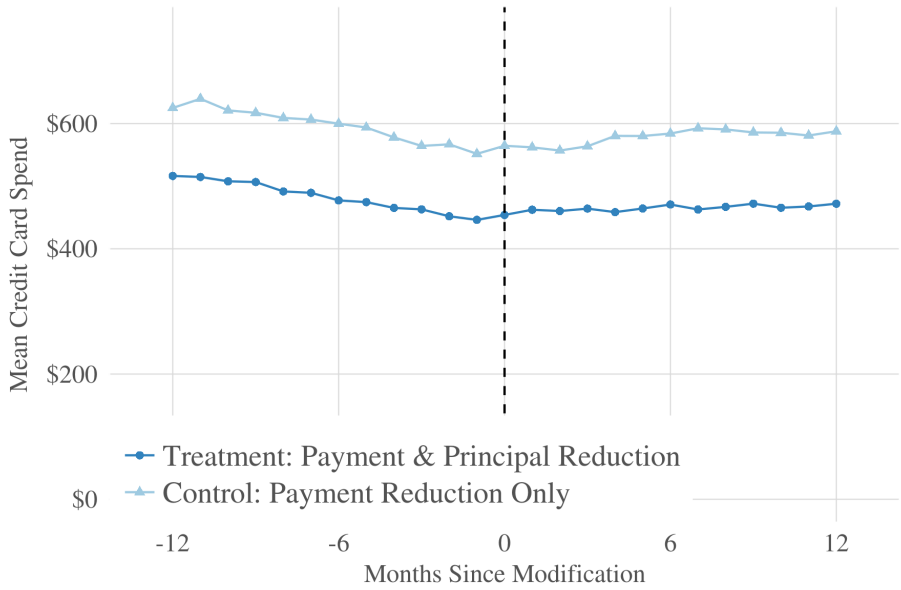

(b) Marginal Propensity to Consume (MPC) by Home Equity Position in Model

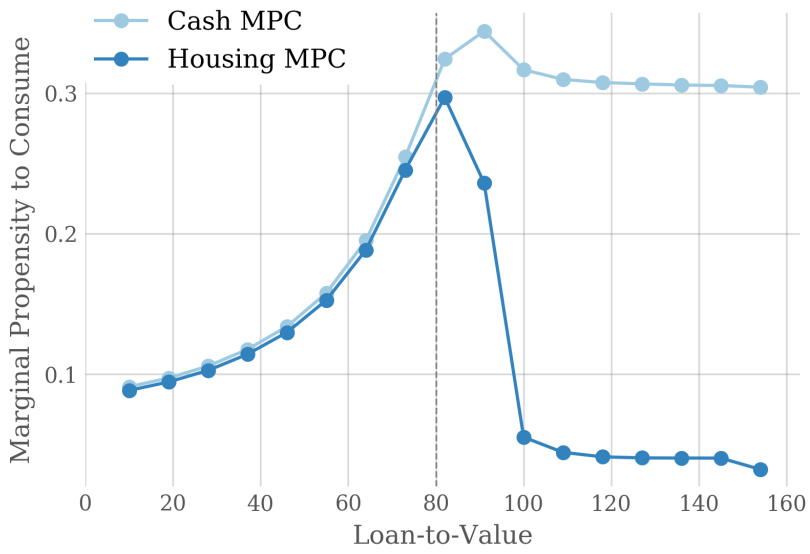

Notes: This figure empirically and theoretically evaluates the impact of principal forgiveness on consumption. Panel (a) shows the event study of monthly credit card expenditure around modification for borrowers receiving each type of modification in the matched HAMP credit bureau dataset. See Section 4 for details. Panel (b) plots our model estimates for the marginal propensity to consume (MPC) out of an additional dollar of cash or an additional dollar of housing wealth generated by mortgage debt forgiveness. This panel assumes a collateral constraint of 80 percent of LTV and cash-on-hand (assets + annual income) of 86 percent of permanent income. See Section 6.2 for further discussion.

Figure 4: Financial Impacts of Private and Government Modifications

(a) Annual Impacts on Payments

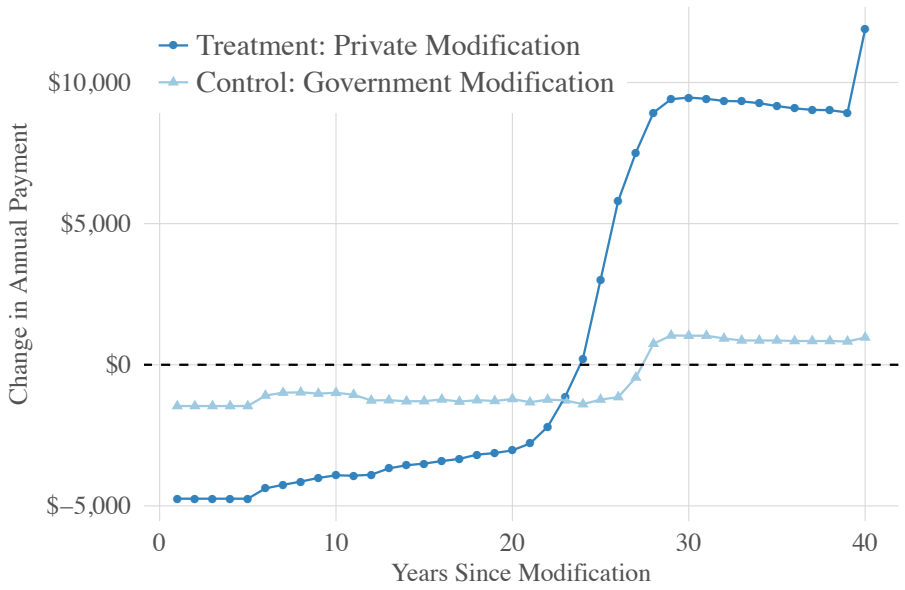

(b) Summary Impacts

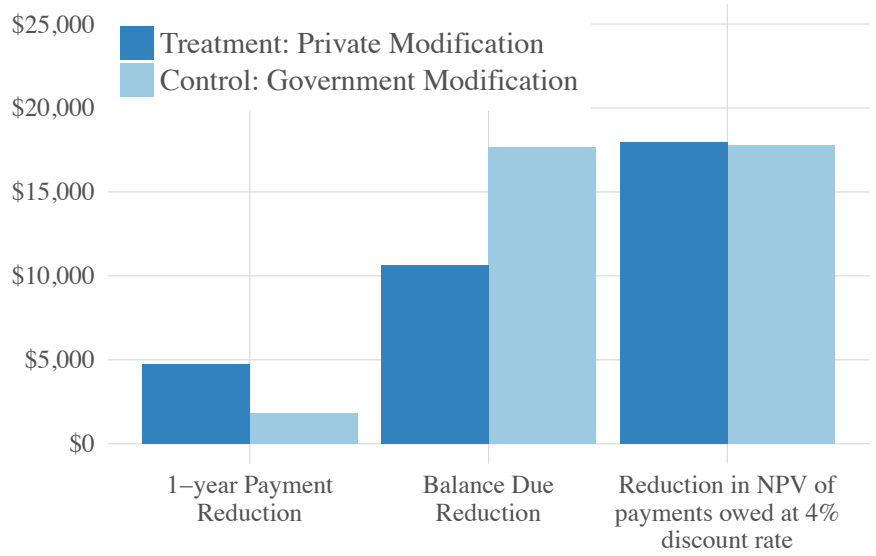

Notes: This figure compares private modifications to government-subsidized HAMP modifications near the HAMP eligibility cutoff. Panel (a) plots the difference in average annual payments for borrowers receiving each type of modification relative to the payments borrowers owed under their unmodified mortgage contracts in the JPMCI bank dataset. Panel (b) summarizes the financial impacts of modifications along various dimensions: the change in the one-year payment, the change in the unpaid balance, and the change in the net present value of mortgage payments owed, discounted at a four percent interest rate. See Section 5 for details.

Figure 5: Effect of Payment Reduction on Default

(a) First Stage -- Change in Mortgage Payment from Modification

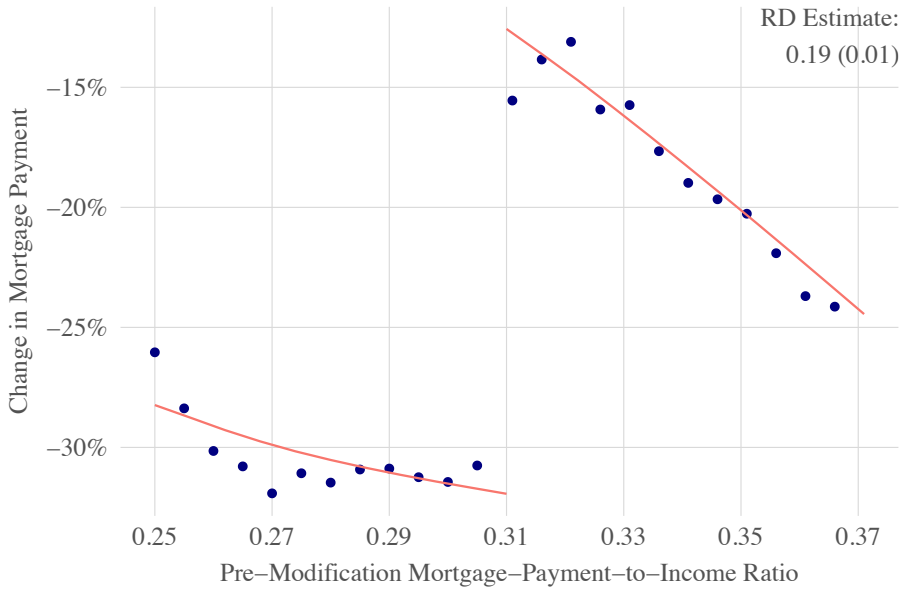

(b) Reduced Form -- Mortgage Default

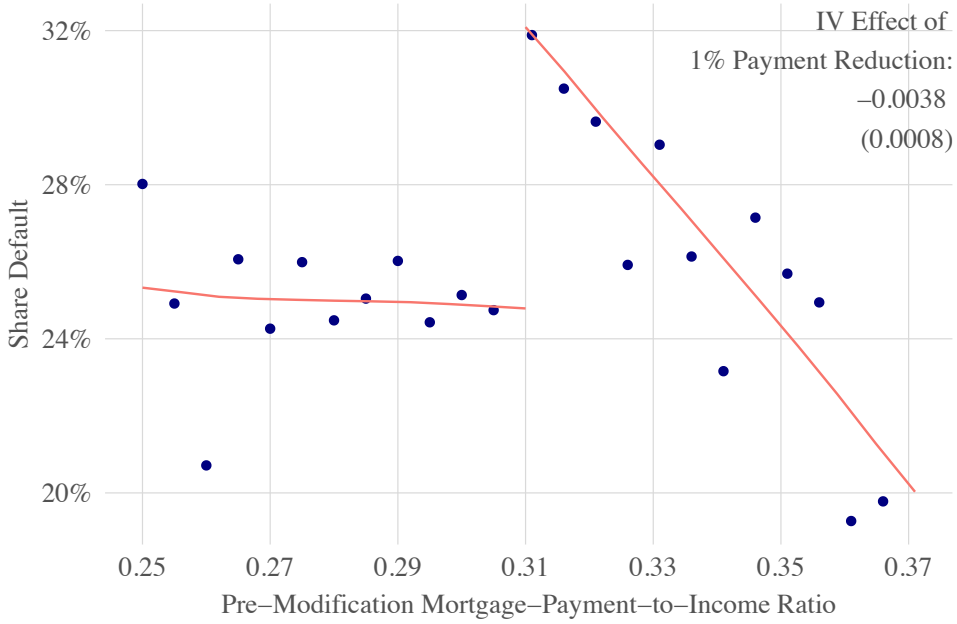

Notes: This figure evaluates the impact of payment reduction on default using a regression discontinuity design around the HAMP eligibility cutoff at the 31 percent Payment-to-Income (PTI) ratio in the JPMCI bank dataset. The horizontal axis shows borrower PTI. The dots are conditional means for 12 equally spaced bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. On the vertical axis, panel (a) plots mean payment reduction and panel (b) plots the two-year default rate, which is defined as being 90 days delinquent at any point within two years of the modification date. Construction of the IV estimate $\hat{\tau}$ is described in Section 5.2.

Table 1: Representativeness

(a) Principal Reduction Regression Discontinuity Sample

	RD Analysis Sample				PSID Delinquent Households			
	Mean	p10	p50	p90	Mean	p10	p50	p90
Income	58,938	28,930	56,416	97,069	64,000	21,000	55,000	120,000
Home Value	257,983	100,000	240,000	440,000	190,000	50,000	140,000	350,000
Loan to Value Ratio	129	105	121	168	101	52	94	166
Monthly Mortgage Payment	1,843	900	1,700	3,000	1,349	459	1,100	2,528
Mortgage Interest Rate	0.058	0.030	0.060	0.080	0.058	0.000	0.060	0.090
Mortgage Term Remaining (Years)	26.0	23.0	25.0	34.5	23.1	10.0	25.0	30.0
Months Past Due	8.6	0.0	6.0	21.0	5.0	2.0	3.0	11.5
Male (d)	0.60	0.00	1.00	1.00	0.68	0.00	1.00	1.00
Age	48.6	36.0	46.0	66.0	43.2	31.0	42.5	57.0
Value of Liquid Assets					3,238	0	250	5,000
N	9,725				190			

(b) Payment Reduction Regression Discontinuity Sample

	RD Analysis Sample				PSID Delinquent Households			
	Mean	p10	p50	p90	Mean	p10	p50	p90
Income	67,811	27,095	54,726	125,196	64,000	21000	55,000	120,000
Home Value	190,341	49,000	139,000	398,000	190,000	50000	140,000	350,000
Loan to Value Ratio	129	63	107	205	101	52	94	166
Monthly Mortgage Payment	1,327	495	1,053	2,553	1,349	459	1,100	2,528
Mortgage Interest Rate	0.068	0.050	0.066	0.092	0.058	0.000	0.060	0.090
Mortgage Term Remaining (Years)	22.5	15.0	24.0	26.5	23.1	10.0	25.0	30.0
Months Past Due	9.1	1.0	7.0	24.0	5.0	2.0	3.0	11.5
N	12,939				190			

Notes: This table compares borrowers in our regression discontinuity samples to delinquent borrowers in the 2009 and 2011 Panel Study of Income Dynamics (PSID) Supplements on Housing, Mortgage Distress, and Wealth Data as reported in Gerardi et al. (2015). The principal reduction sample includes borrowers with v within 0.61 percent of the cutoff (from equation 1) and the payment reduction sample includes borrowers with PTI within 6 percent of the cutoff. All values are before modification. Panel (b) does not include gender, age, or liquid assets since these are not observed for this sample. The PSID sample includes heads of households who are mortgagors, ages 24-65, are labor force participants, and are 60 or more days late on their mortgage as of the survey date. The summary statistics are repeated in panel (a) and panel (b). Liquid assets include checking and savings account balances, money market funds, certificates of deposit, Treasury securities, and other government saving bonds. (d) indicates a dummy variable.

Table 2: Impact of Principal Reduction on Expenditure

(a) Credit Card Expenditure ($/month)

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (Principal Reduction x Post)	0.686	0.721	0.811	2.068	0.557	2.240
	(3.621)	(3.619)	(3.685)	(3.855)	(3.887)	(3.912)
MSA Fixed Effects		Yes				
Calendar Month Fixed Effects			Yes			
MSA by Calendar Month Fixed Effects				Yes	Yes	Yes
Controls					Yes	Yes
Controls x Post Interactions						Yes
Dependent Variable Mean	483.82	483.82	483.82	483.82	485.55	485.55
Observations	1,678,612	1,678,612	1,678,612	1,678,612	1,642,328	1,642,328
Adjusted R^2	0.003	0.018	0.005	0.015	0.081	0.081

(b) Auto Expenditure ($/month)

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (Principal Reduction x Post)	13.541	13.545	13.062	14.648	15.333	11.054
	(8.824)	(8.824)	(8.896)	(9.011)	(9.134)	(9.308)
MSA Fixed Effects		Yes				
Calendar Month Fixed Effects			Yes			
MSA by Calendar Month Fixed Effects				Yes	Yes	Yes
Controls					Yes	Yes
Controls x Post Interactions						Yes
Dependent Variable Mean	185.84	185.84	185.84	185.84	186.62	186.62
Observations	1,678,612	1,678,612	1,678,612	1,678,612	1,642,328	1,642,328
Adjusted R^2	0.001	0.001	0.001	0.003	0.004	0.004

Notes: This table reports difference-in-differences estimates of the effect of principal reduction on expenditure in the matched HAMP credit bureau dataset. The dependent variable in panel (a) is monthly credit card expenditure, while the dependent variable in panel (b) is monthly auto expenditure computed based on balances of new auto loans. The coefficient of interest, *Treatment*, is the estimated change in the difference between outcomes of mortgages receiving modifications with and without principal reduction during the year after modification. All specifications include fixed effects for modification type and months since modification. Controls include the predicted gain to lenders of providing principal reduction, the predicted gain interacted with a dummy for this value being positive, FICO score, monthly income, pre-modification loan characteristics (LTV, principal balance, DTI, monthly payment), property value, LTV at origination, non-housing monthly debt payment, and monthly payment reduction. The sample includes underwater borrowers who are observed one year before and after modification and report positive credit card expenditure in at least one month during this window. The dependent variable mean is reported for borrowers receiving principal reduction modifications in the year before modification. Standard errors, in parentheses, are clustered at the borrower level ($n_{borrower} = 69,496$). See Section 4 for additional detail on the specification, outcome measures, and sample.

Liquidity vs Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession – Online Appendix

Peter Ganong and Pascal Noel

Contents

A Appendix Figures and Tables

Figure 1: Match Rate around Principal Reduction Discontinuity

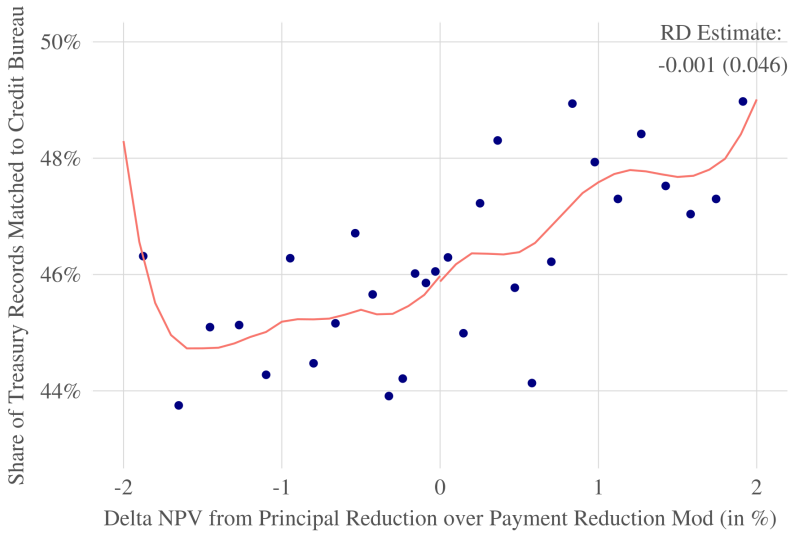

Notes: This figure plots the share of borrowers in the Treasury HAMP dataset successfully matched to their credit bureau records. The horizontal axis shows the normalized predicted gain to lenders of providing principal reduction to borrowers from equation (1). The dots are conditional means for 15 bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (2).

Figure 2: Modification Terms Summary

C. Interest at the rate of __2.0_____% will begin to accrue on the New Principal Balance as of ____1/1/2012___ and the first new monthly payment on the New Principal Balance will be due on __1/15/2012___. My payment schedule for the modified Loan is as follows:

Years	Interest Rate	Interest Rate Change Date	Monthly Principal and Interest Payment Amount	Estimated Monthly Escrow Payment Amount*	Total Monthly Payment*	Payment Begins On	Number of Monthly Payments
1-5	2.00%	01/01/2012	$1,000.06	$312.50, may adjust periodically	$1,312.50, may adjust periodically	01/15/2012	60
6	3.00%	01/01/2013	$1,143.71	$312.50, May adjust periodically	$1,455.21, May adjust periodically	01/15/2013	12
7	4.00%	01/01/2014	$1,291.06	$312.50, May adjust periodically	$1,603.56, May adjust periodically	01/15/2014	12
8-35	5.00%	01/01/2015	$1,444.00	$312.50, May adjust periodically	$1,756.50. May adjust periodically	01/15/2015	336

Notes: This figure shows the modified payment terms as explained to borrowers in the modification agreement, which they are required to sign. Example terms are shown for a mortgage with a post-modification principal balance of $300,000, temporary interest rate of 2 percent, mortgage term of 35 years, and escrow payments equal to 1.5 percent of the property value ($250,000).

Figure 3: Mortgage Delinquency over Time

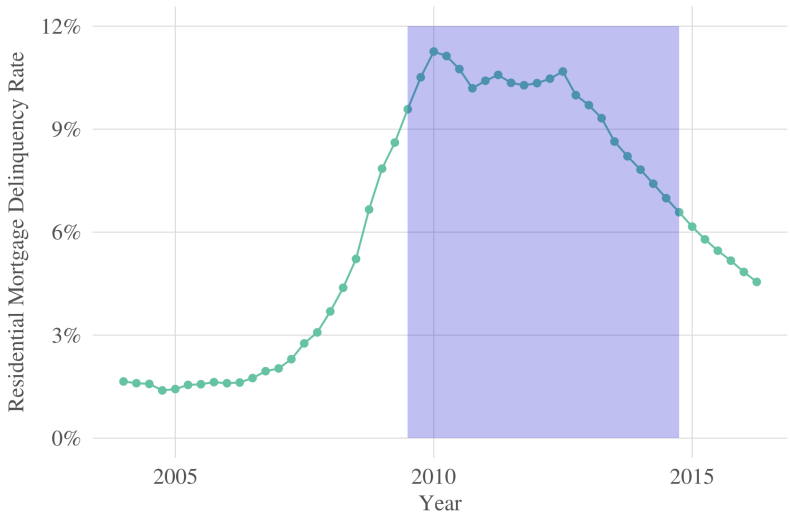

Notes: This figure plots the share of U.S. residential mortgages more than 30 days delinquent as reported by the Federal Reserve Board. The shaded region denotes the period where borrowers in our principal reduction sample had their first pre-modification delinquencies.

Figure 4: Pre-Modification Characteristics around Principal Reduction Discontinuity

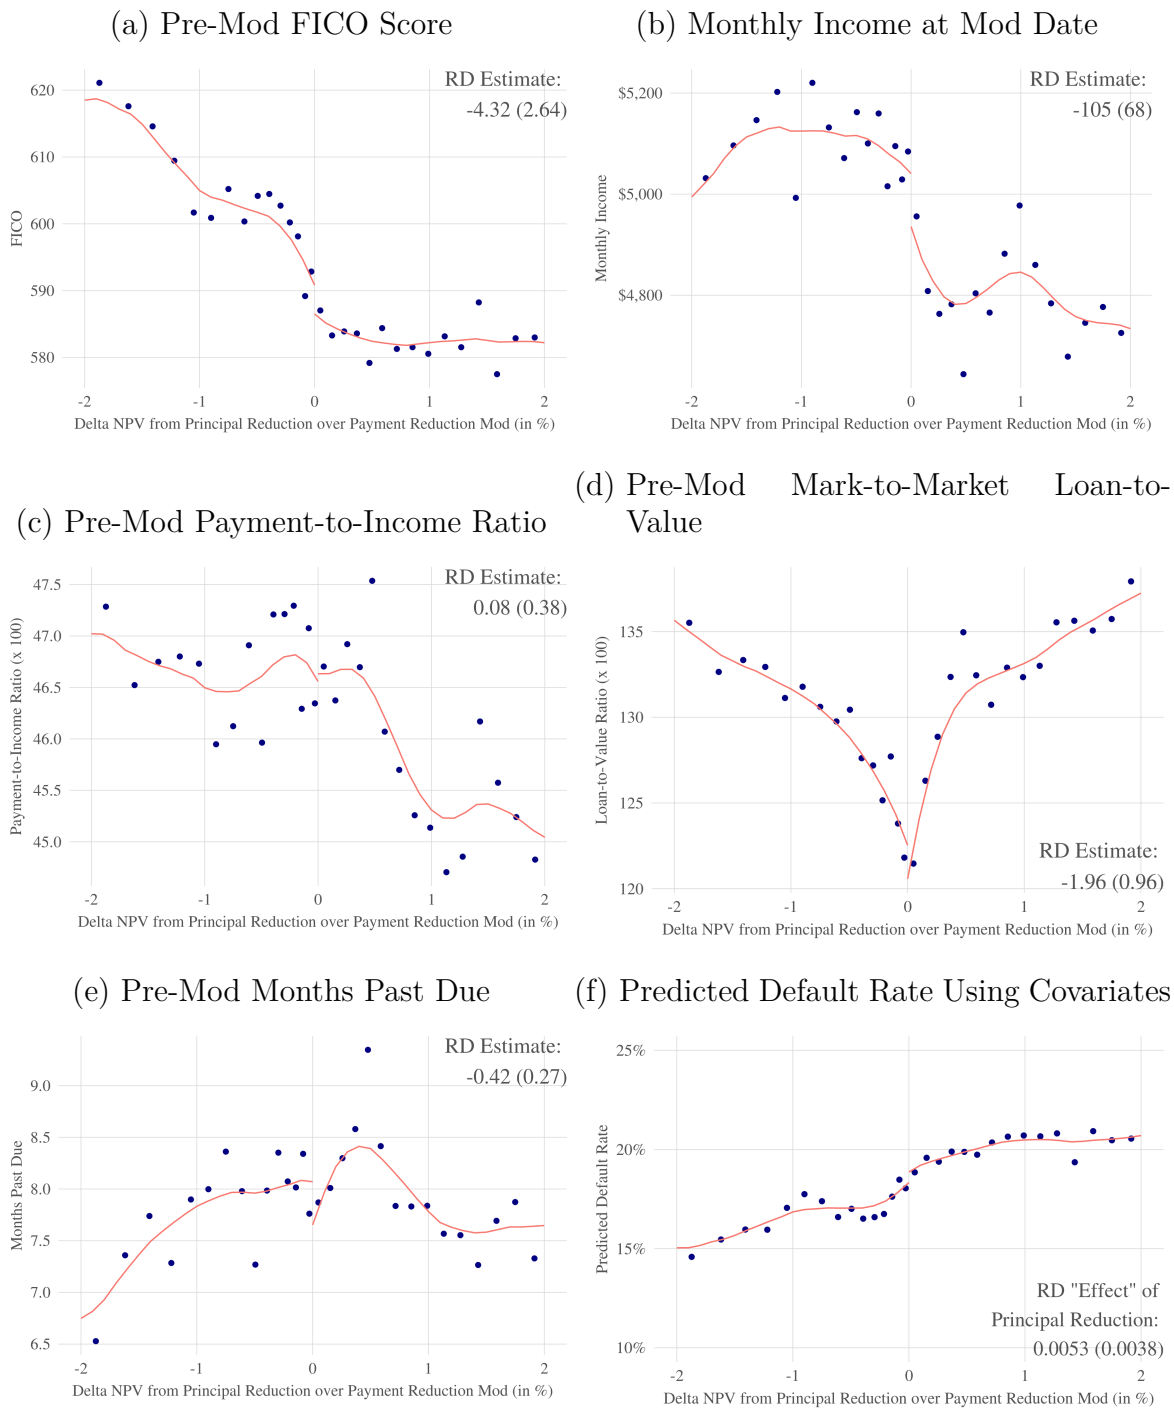

Notes: This figure shows average pre-treatment characteristics around the regression discontinuity cutoff in the matched HAMP credit bureau dataset. The horizontal axis shows the normalized predicted gain to lenders of providing principal reduction to borrowers from equation (1). The vertical axis in the first five panels shows borrower credit score, monthly income, the ratio of monthly mortgage payments to monthly income, the ratio of unpaid principal balance to the market value of the house (mark-to-market loan-to-value ratio), and the number of monthly mortgage payments the borrower is past due at application date. The final panel shows predicted default rates from a linear regression of default on the first five borrower characteristics. The dots are conditional means for 15 bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (2). See Appendix B.1.3 for details.

Figure 5: Borrower Density and Take-Up around Principal Reduction Discontinuity

(a) Histogram of Running Variable

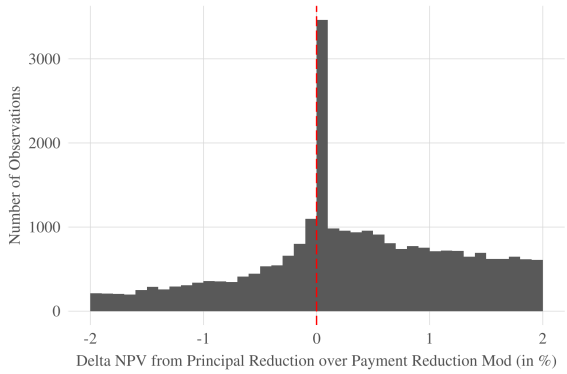

(b) Histogram of Running Variable Excluding Zeros

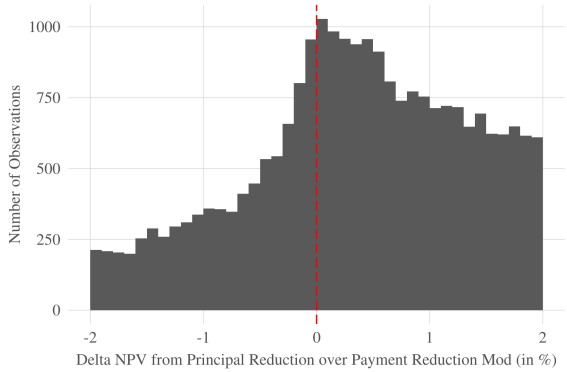

(c) Take-up Rate

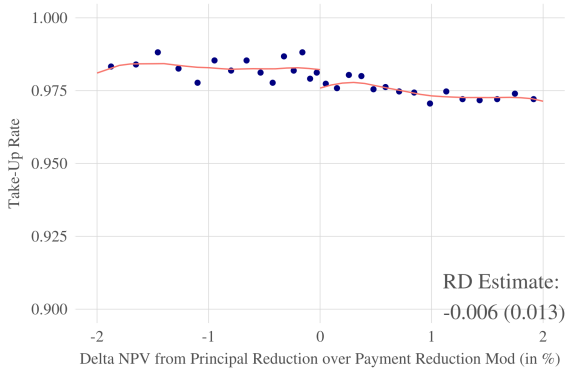

Notes: Panel (a) plots the histogram of the running variable from our regression discontinuity strategy in the matched HAMP credit bureau dataset. The horizontal axis shows the normalized predicted gain to lenders of providing principal reduction to borrowers from equation (1). HAMP program officers in the U.S. Treasury Department explain that the mass at exactly zero is due to data misreporting. Some servicers reported a single number as the calculation for both the payment reduction and principal reduction modifications, meaning that the estimated gains from principal reduction were calculated to be zero. Panel (b) plots the same histogram dropping observations exactly at zero, which is our analysis sample. Appendix B.1.3 discusses four additional arguments for why the mass at zero is unlikely to pose a challenge for the validity of the regression discontinuity research design. Panel (c) shows the take-up rate conditional on borrowers being offered a modification in the Treasury HAMP dataset.

Figure 6: Treatment Size around Principal Reduction Discontinuity

Notes: This figure shows the treatment size at the regression discontinuity cutoff in the matched HAMP credit bureau dataset. The horizontal axis shows the normalized predicted gain to lenders of providing principal reduction to borrowers from equation (1). Panel (a) shows the change in the net present value (NPV) of mortgage payments owed under the modified contract relative to the status quo mortgage contract, discounted at a 4 percent interest rate, panel (b) shows the change in the loan-to-value ratio, panel (c) shows the change in initial monthly housing payments, and panel (d) shows the average amount of principal reduction per borrower. The dots are conditional means for 15 bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (2). Construction of the IV estimate $\hat{\tau}$ is described in Section 3.2.

Figure 7: Regression Discontinuity Robustness to Alternative Bandwidths around Principal Reduction Discontinuity

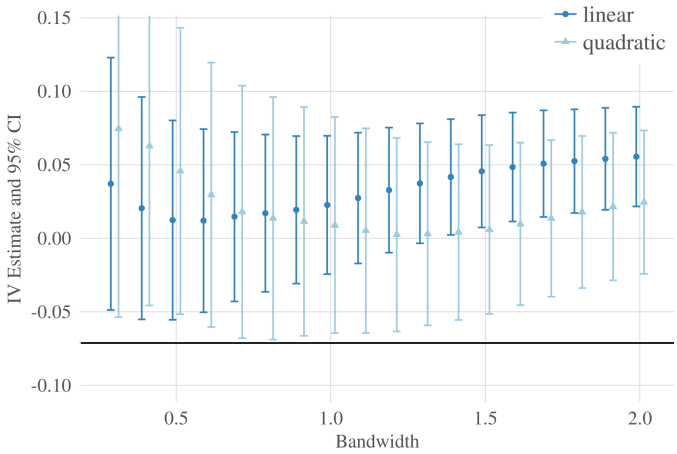

Notes: This figure shows the estimated impact of principal reduction on default under various specifications and bandwidths in the matched HAMP credit bureau dataset. Each line plots the IV estimate and associated 95 percent confidence interval from a local linear or quadratic regression on either side of the cutoff. The optimal bandwidths for the linear specification from Imbens and Kalyanaraman (2012) and Calonico, Cattaneo and Titiunik (2014) are 0.61 and 0.63, respectively. The optimal bandwidths for a quadratic specification from Imbens and Kalyanaraman (2012) and Calonico, Cattaneo and Titiunik (2014) are 0.80 and 0.67, respectively. The black horizontal line is the predicted impact of principal reduction on default from Treasury's redefault model.

Figure 8: Spending around Modifications with and without Principal Reduction

(a) Credit Card Spending around Modification – Normalized

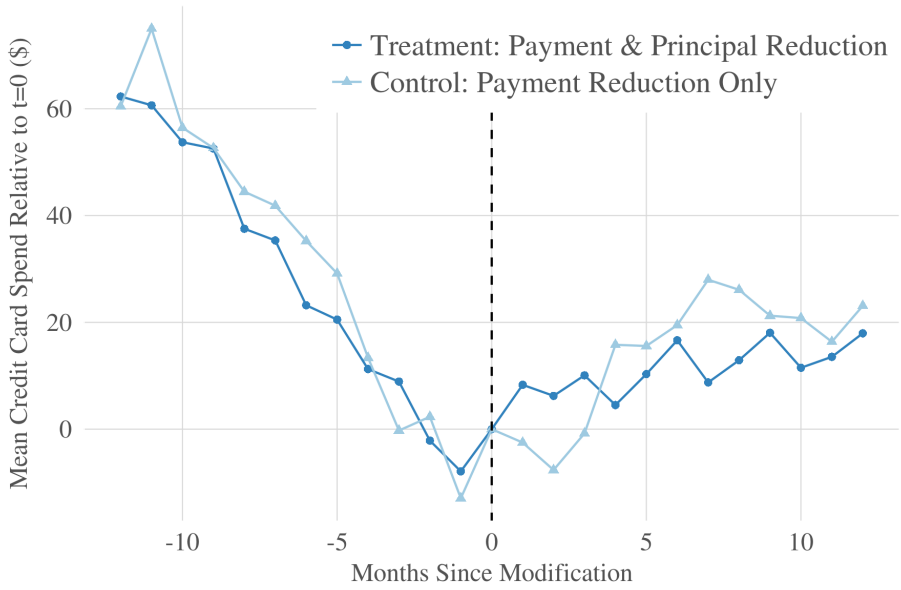

(b) Auto Spending around Modification

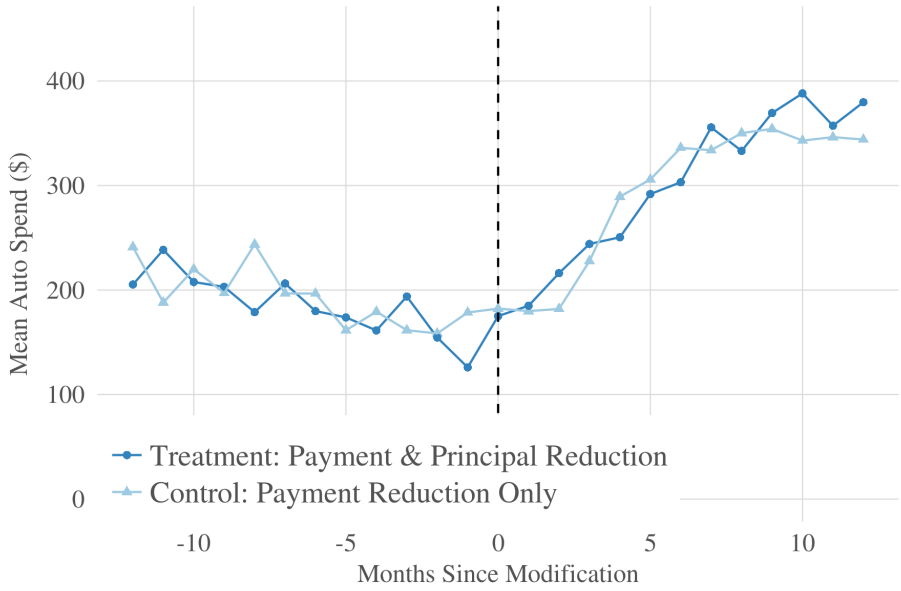

Notes: This figure shows the event study of monthly spending around modification for borrowers receiving each type of modification in the matched HAMP credit bureau dataset. The top panel plots credit card expenditure in dollars as measured from credit bureau records relative to the month of modification (discussed in Section 2.1). The bottom panel shows the event study of monthly auto spending around modification. Auto spending is measured from new auto loans, as described in Section 2.1. See Appendix Table 2 for sample summary statistics.

Figure 9: Spending around Modifications with and without Principal Reduction using Bank Data

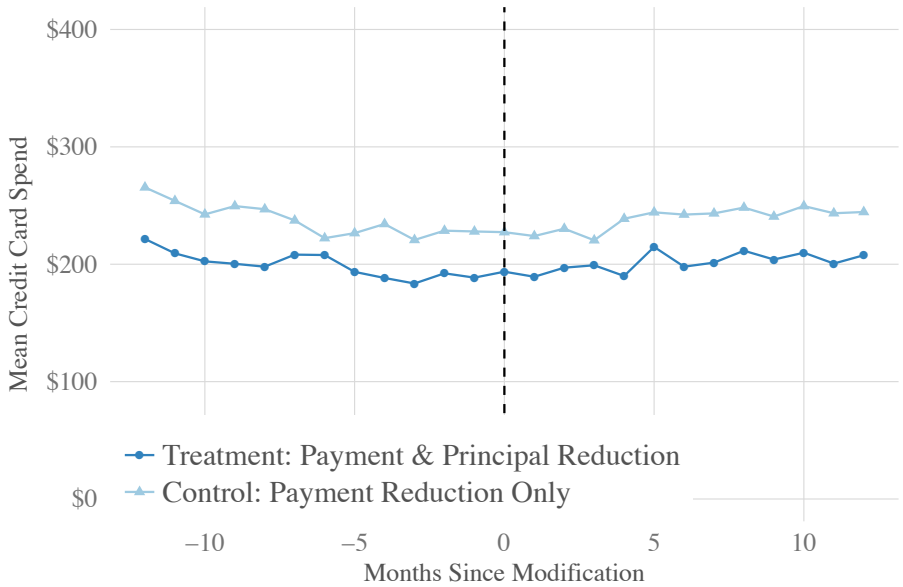

Notes: This figure shows the event study of monthly credit card expenditure around modification for borrowers receiving each type of modification in the JPMCI bank account dataset. For further details see Sections 2.2 and 4.2.

Figure 10: Credit Card and Auto Spend around Principal Reduction Discontinuity

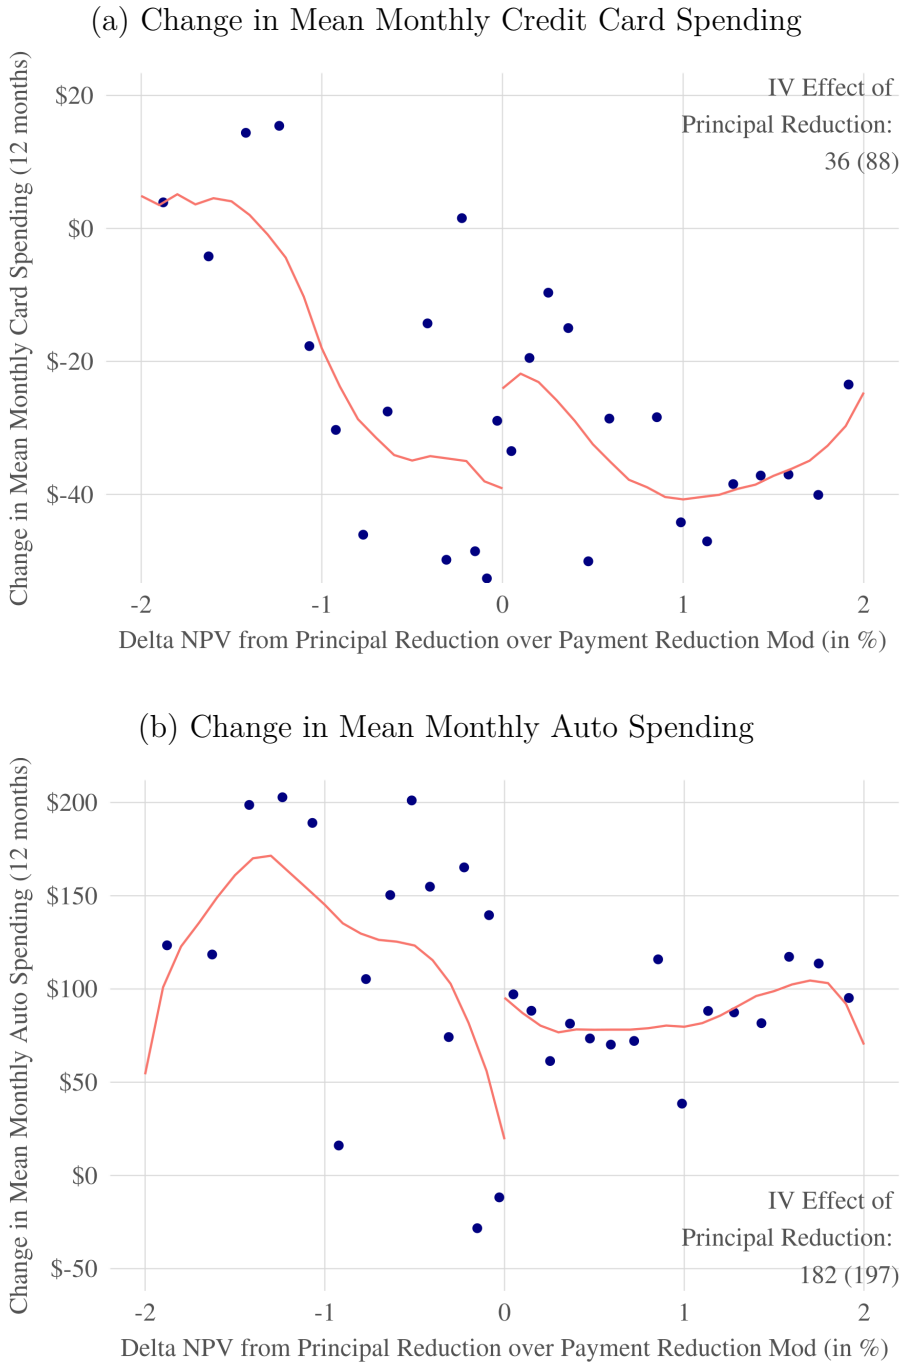

Notes: This figure shows the estimated impact of principal reduction on expenditure using the fuzzy regression discontinuity strategy in the matched HAMP credit bureau dataset. The horizontal axis shows the normalized predicted gain to lenders of providing principal reduction to borrowers from equation (1). The vertical axis on the top panel shows the average change in credit card expenditure between the 12 months before modification and the 12 months after modification. Credit card expenditure is measured from credit bureau records as discussed in Section 2.1. The vertical axis in the bottom panel shows the average change in auto spending between the 12 months before modification and the 12 months after modification. Auto spending is measured from new auto loans, as described in Section 2.1. The dots are conditional means for 15 bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. Construction of the IV estimate $\hat{\tau}$ is described in Section 3.2.

Figure 11: Pre-Modification Characteristics around Payment Reduction Discontinuity

(a) Monthly Income at Mod Date

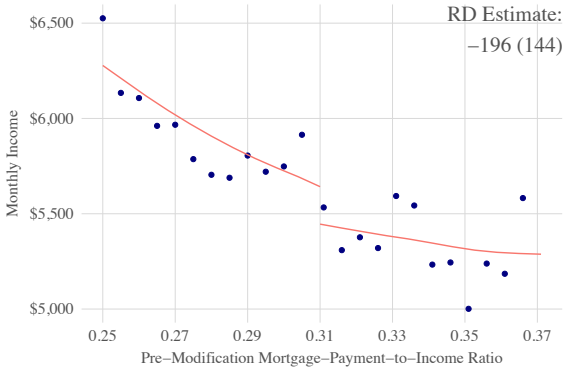

(b) Pre-Mod Monthly Payment

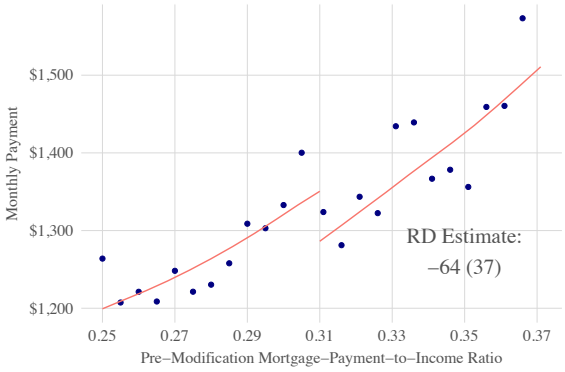

(c) Pre-Mod Mark-to-Market Loan-to-Value

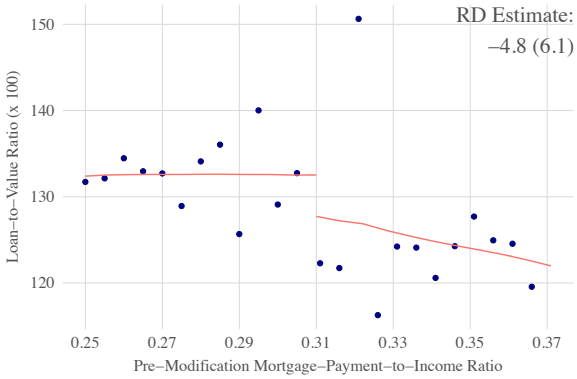

(d) Pre-Mod Months Past Due

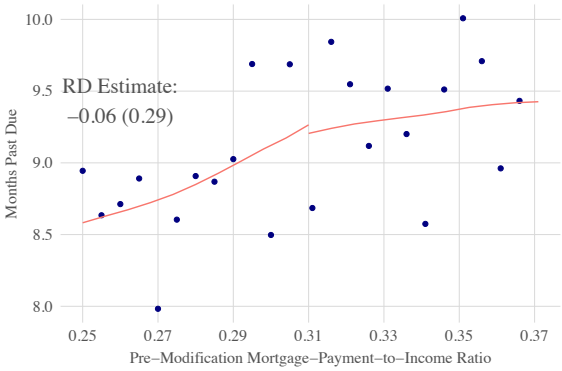

(e) Predicted Default Rate Using Covariates

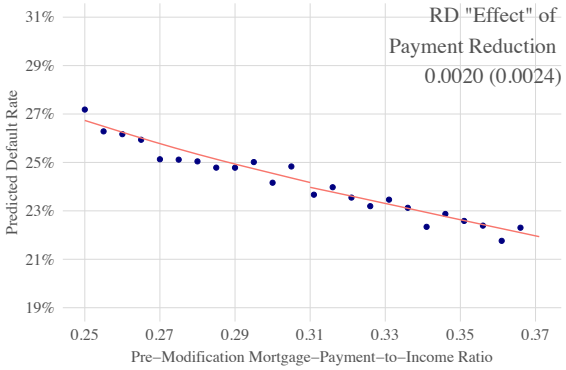

Notes: This figure shows average pre-treatment characteristics around the 31 percent PTI regression discontinuity cutoff in the JPMCI bank dataset for non-GSE-backed loans. The horizontal axis shows pre-modification borrower PTI. The vertical axis in the first four panels shows monthly income, monthly payment, the ratio of unpaid principal balance to the market value of the house (mark-to-market loan-to-value ratio), and the number of months past due at modification date. The final panel shows predicted default rates from a linear regression of default on the first four borrower characteristics. The dots are conditional means for 12 bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (4).

Figure 12: Borrower Density around the Payment Reduction Discontinuity

(a) Histogram of Running Variable – Near Discontinuity

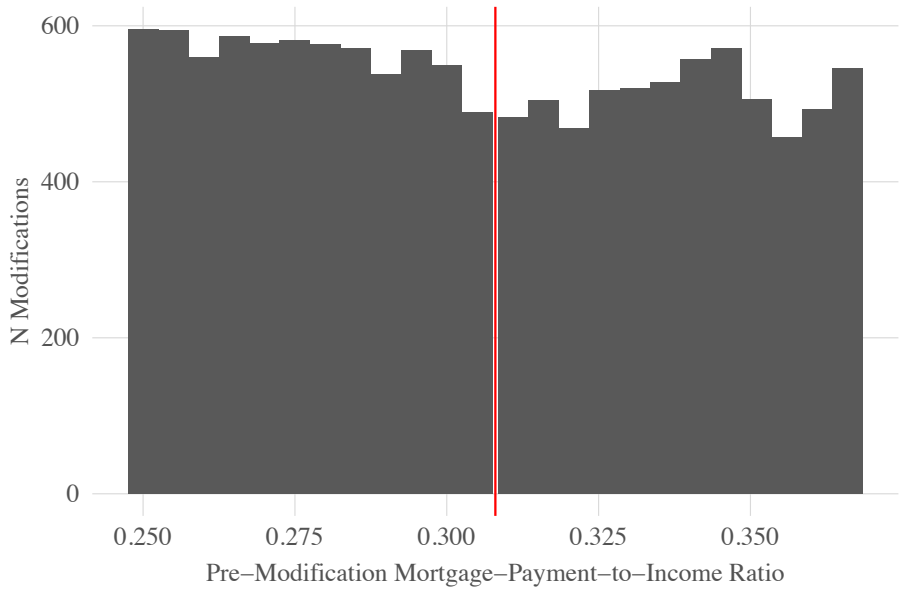

(b) Histogram of Running Variable – Full Support

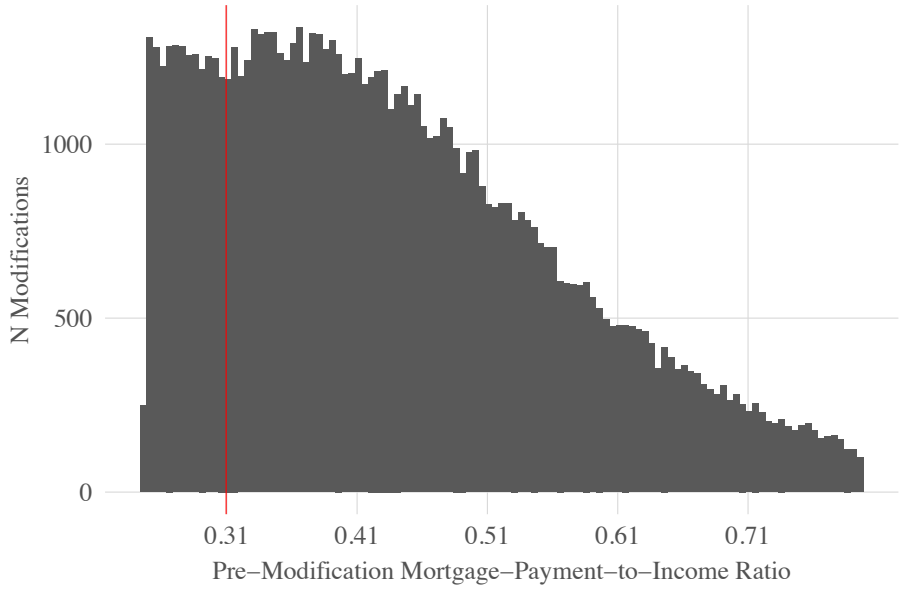

Notes: This figure plots the histogram of the running variable from our 31 percent PTI regression discontinuity strategy in the JPMCI bank dataset for non-GSE-backed loans. The horizontal axis shows pre-modification borrower PTI. The top panel shows borrowers in the main analysis sample. This sample is restricted to pre-modification PTI ratio between 25 percent and 37 percent (dropping the 241 observations between 31.0 percent and 31.1 percent), pre-modification terms 30 years or less, and fixed rate loans. This is our main analysis sample. The bottom panel shows the density for the full sample.

Figure 13: Treatment Size around Payment Reduction Discontinuity

(a) ΔNet Present Value of Payments Owed

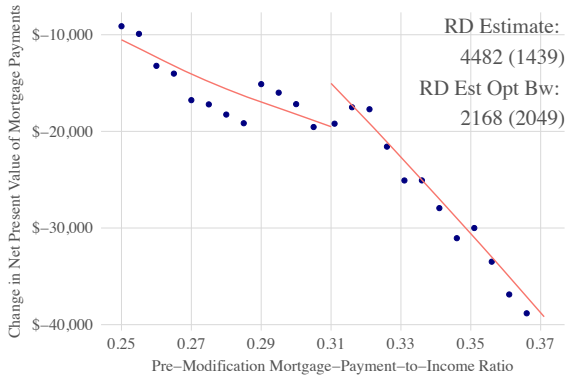

(b) Mortgage Principal Forgiveness

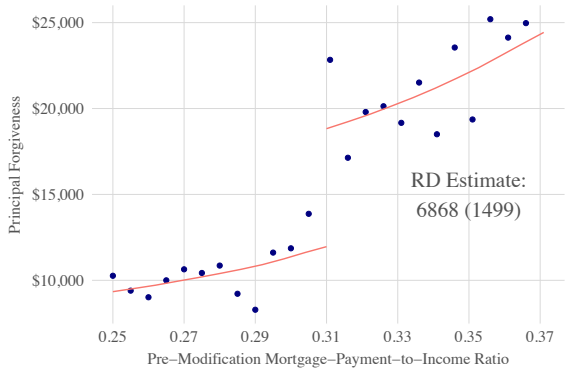

(c) ΔInterest Rate

Notes: This figure describes the treatment in terms of long-term obligations around the 31 percent PTI discontinuity in the JPMCI bank dataset for non-GSE-backed loans. The dots are conditional means for 12 bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (4) using the IK-optimal bandwidth for delinquency of 0.061. Panel (a) shows the change in the NPV of payments owed under the mortgage contract before and after modification. The IK-optimal bandwidth *for this outcome variable* is 0.039 and the label also includes a second RD estimate using this optimal bandwidth of 0.039. Panel (b) shows mortgage principal forgiveness. Panel (c) shows the change in the interest rate.

Figure 14: Effect of Payment Reduction on Default: Robustness to Broader Sample

(a) First Stage – Change in Mortgage Payment from Modification

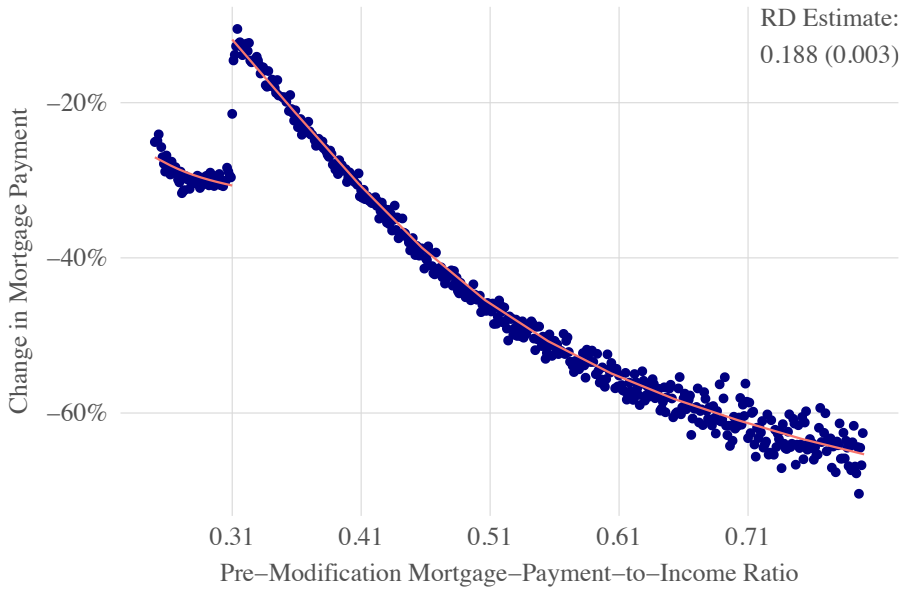

(b) Reduced Form -- Mortgage Default

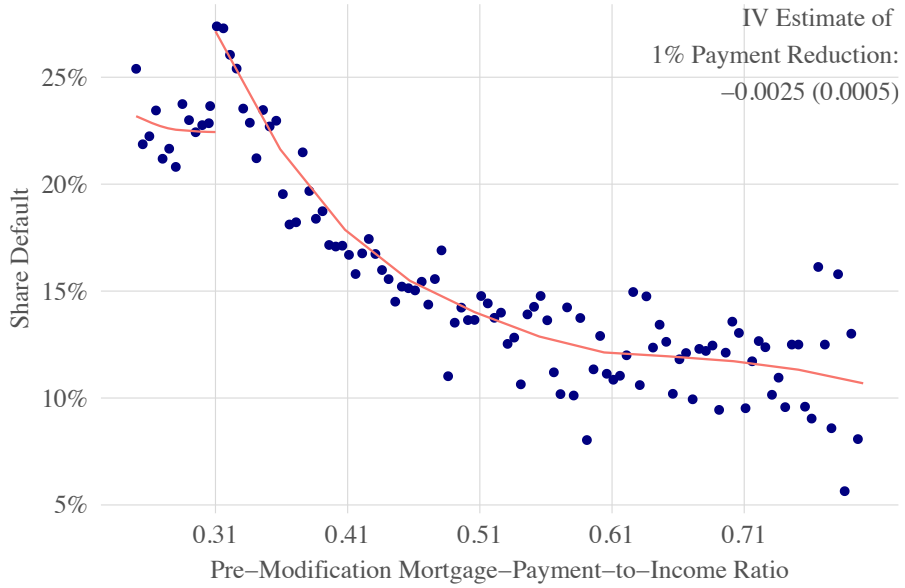

Notes: This figure shows the estimated effect of payment reduction on default using the 31 percent PTI regression discontinuity in the JPMCI bank dataset for a broader sample of non-GSE-backed loans. It includes loans with pre-modification terms greater than 30 years, loans with variable interest rates, and borrowers with PTI between 31 percent and 31.1 percent, all of which are dropped in the main analysis. The top panel plots the first stage, with payment reduction on the vertical axis and borrower PTI on the horizontal axis. The dots are conditional means for equally spaced bins on each side of the cutoff. Bins are four times narrower than in Figure 5a in order to visually capture the loans between 31 percent and 31.1 percent with a separate dot. All other plot details are the same as Figure 5.

Figure 15: Effect of Payment Reduction on Default Using Alternative Bandwidths

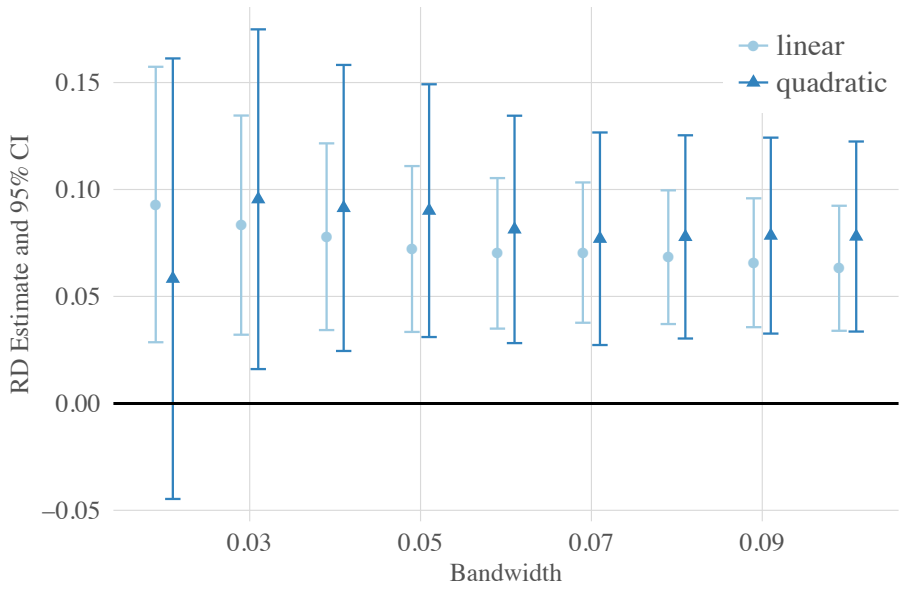

Notes: This figure plots the estimated reduced form jump in default and the associated 95 percent confidence interval at the 31 percent PTI regression discontinuity cutoff calculated using alternative bandwidths in the JPMCI bank dataset for non-GSE-backed loans. Our primary specification uses a bandwidth of 0.06. The plotting convention here differs slightly from Appendix Figure 7, which reports an IV estimate for a variety of bandwidths. Here we instead report the reduced form estimate because the estimated first stage in Figure 5a (and therefore the IV estimate reported in Figure 5b) are somewhat sensitive to bandwidth choice.

Figure 16: Effect of Payment Reduction on Default for GSE-Backed Loans

(a) First Stage -- Change in Mortgage Payment from Modification

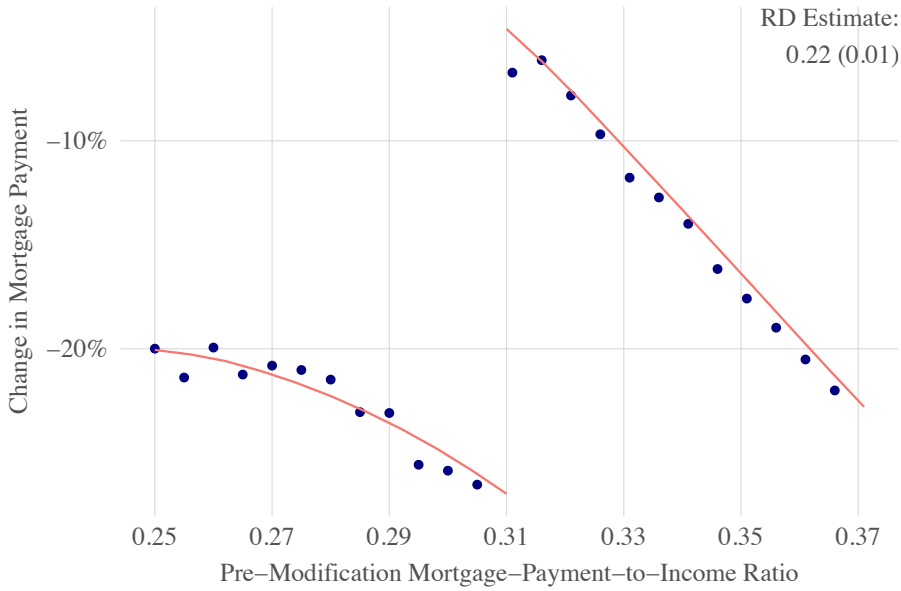

(b) Reduced Form -- Mortgage Default

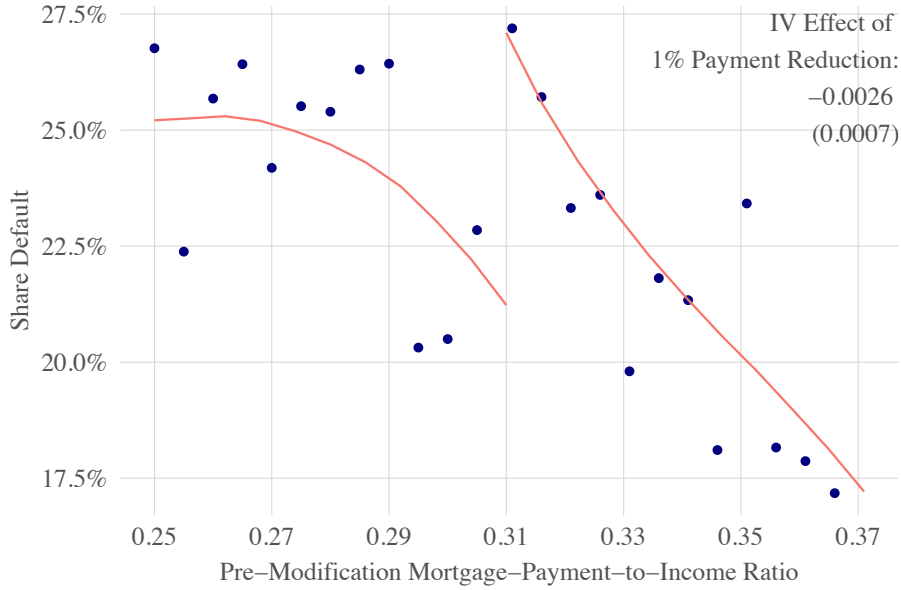

Notes: This figure evaluates the impact of payment reduction on default using a regression discontinuity at the 31 percent payment-to-income (PTI) in the JPMCI bank dataset for GSE-backed loans. The horizontal axis shows borrower PTI. The dots are conditional means for 12 equally spaced bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (4). The top panel plots mean payment reduction and the bottom panel plots the default rate on the vertical axis, which is defined as being 90 days delinquent at any point within two years of the modification date. Construction of the IV estimate $\hat{\tau}$ is described in Section 5.2.

Figure 17: Treatment Size around Payment Reduction Discontinuity for GSE-Backed Loans

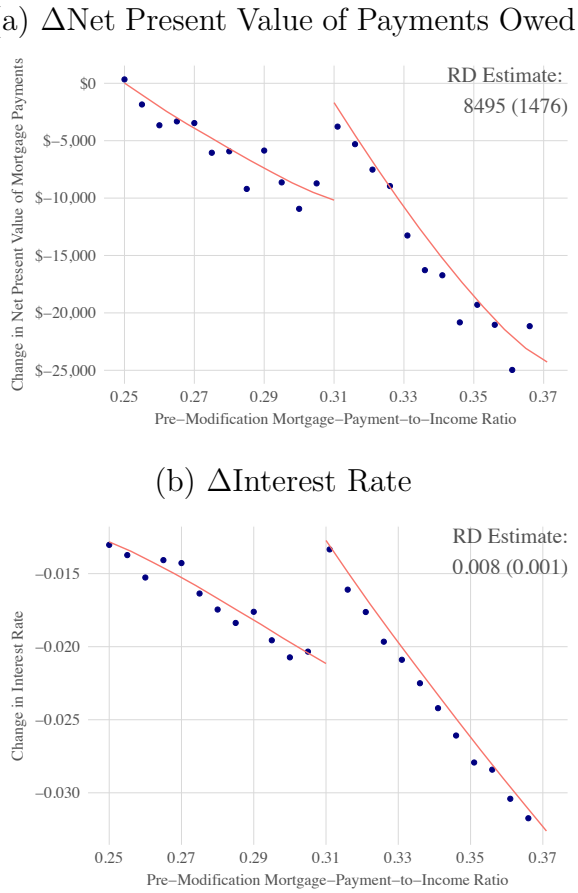

Notes: This figure describes the treatment in terms of long-term obligations around the 31 percent PTI discontinuity in the JPMCI bank dataset for GSE-backed loans. The dots are conditional means for 12 bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (4). Panel (a) shows the change in the NPV of payments owed under the mortgage contract for all loans. Panel (b) shows the change in the interest rate. We do not include principal forgiveness because the GSEs did not offer mortgage principal forgiveness for either private modifications or HAMP modifications.

Figure 18: Pre-Modification Characteristics around Payment Reduction Discontinuity for GSE-Backed Loans

(a) Pre-Mod Mark-to-Market Loan-to-Value

(b) Pre-Mod Months Past Due

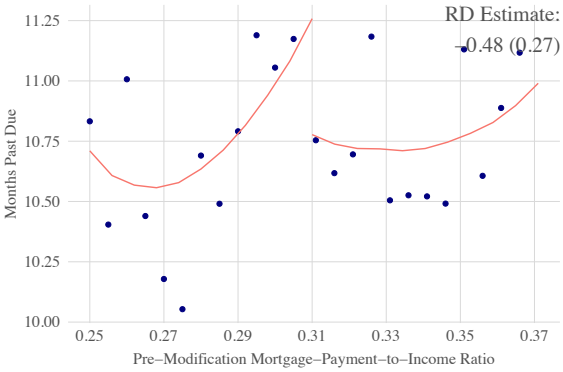

(c) Monthly Income at Mod Date

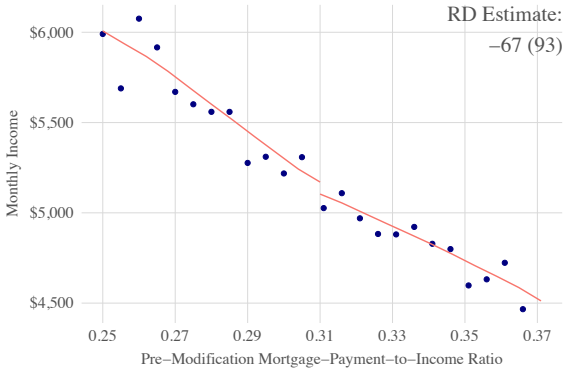

(d) Pre-Mod Monthly Payment

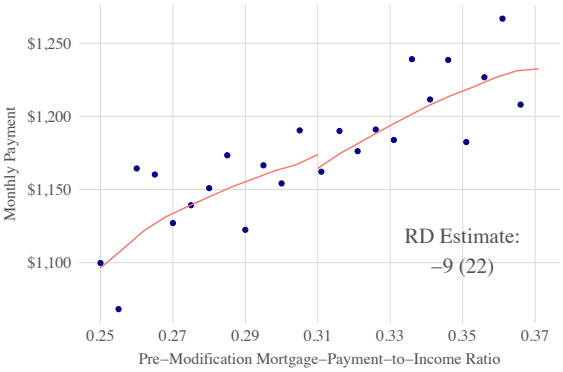

(e) Predicted Default Rate Using Covariates

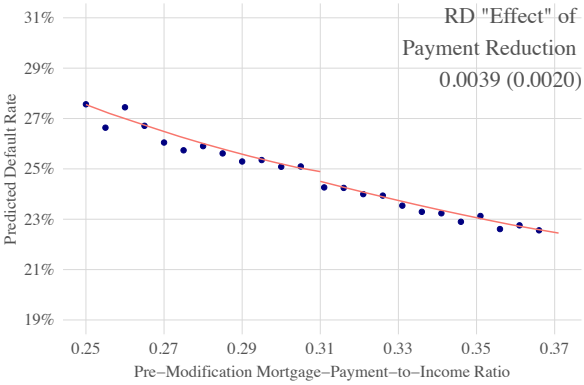

Notes: This figure shows average pre-treatment characteristics around the 31% PTI regression discontinuity cutoff in the JPMCI bank dataset for GSE-backed loans. The horizontal axis shows pre-modification borrower PTI. The vertical axis in the first four panels shows the ratio of unpaid principal balance to the market value of the house (mark-to-market loan-to-value ratio), the number of months past due at modification date, monthly income, and monthly payment. The final panel shows predicted default rates from a linear regression of default on the first four borrower characteristics. The dots are conditional means for 12 bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (4).

Figure 19: Borrower Density around the Payment Reduction Discontinuity for GSE-Backed Loans

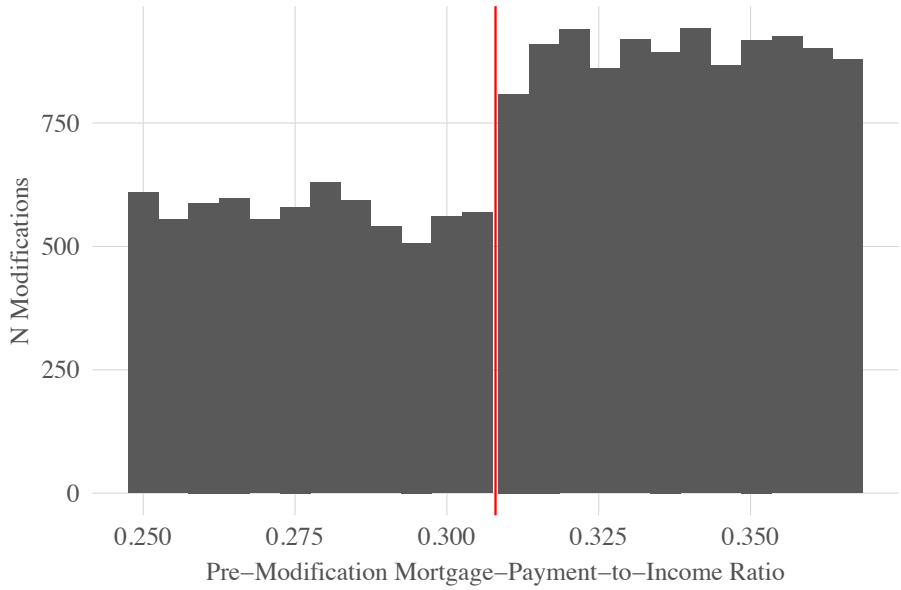

Notes: This figure plots the histogram of the running variable from our 31 percent PTI regression discontinuity strategy in the JPMCI bank dataset for GSE-backed loans. The horizontal axis shows pre-modification borrower PTI. See Section 5.3 for details on why there are more borrowers to the right of 0.31.

Figure 20: Efficient Default-Minimizing Modifications

(a) Cost to Lender of Reducing Payments by 10% for Each Modification Step

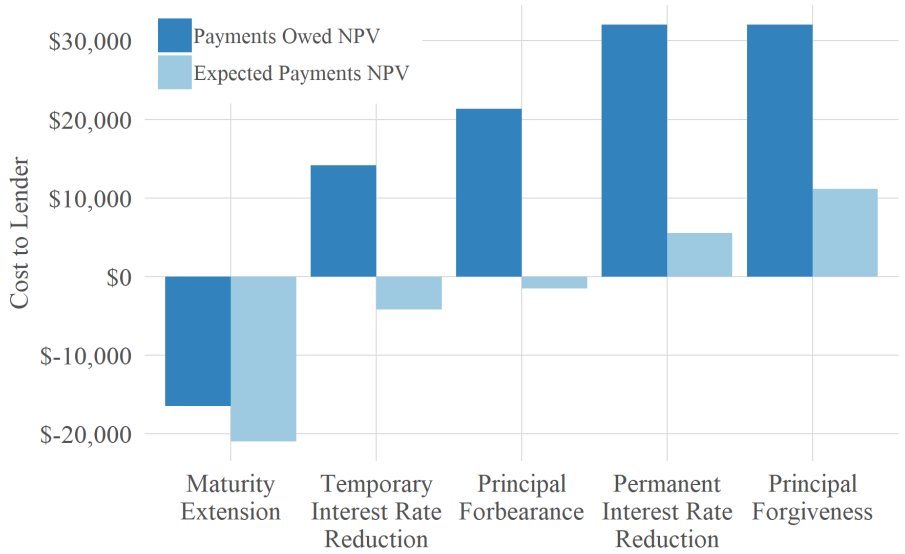

(b) Expected Payments NPV Cost of Payment Reduction for Various Sequences of Modification Steps

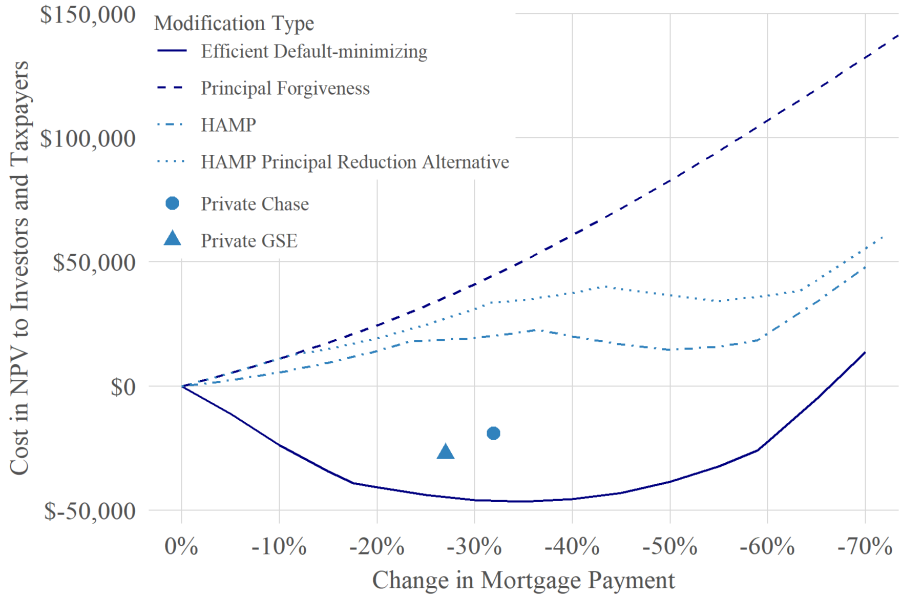

Notes: This figure shows the cost of different mortgage modifications for an illustrative mortgage with the average characteristics of loans at the HAMP eligibility cutoff (a 6.7 percent fixed interest rate, a 23-year term, and a mean unpaid balance of $248,000). Panel (a) shows the cost of reducing payments by 10 percent for five different possible modification steps. The light blue bars show the change in the NPV of *payments owed* under the mortgage contract and the dark blue bars show the change in the NPV of *expected payments* to the lender incorporating the yield curve and the impact of modification on default and prepayment risk. Panel (b) shows the expected payments NPV cost of various modification strategies. The last two programs in the legend are a dot—rather than a line—because they target a specific amount of payment reduction. See Appendix C.4 for details.

Figure 21: Financial Impact of Modifications with and without Principal Reduction

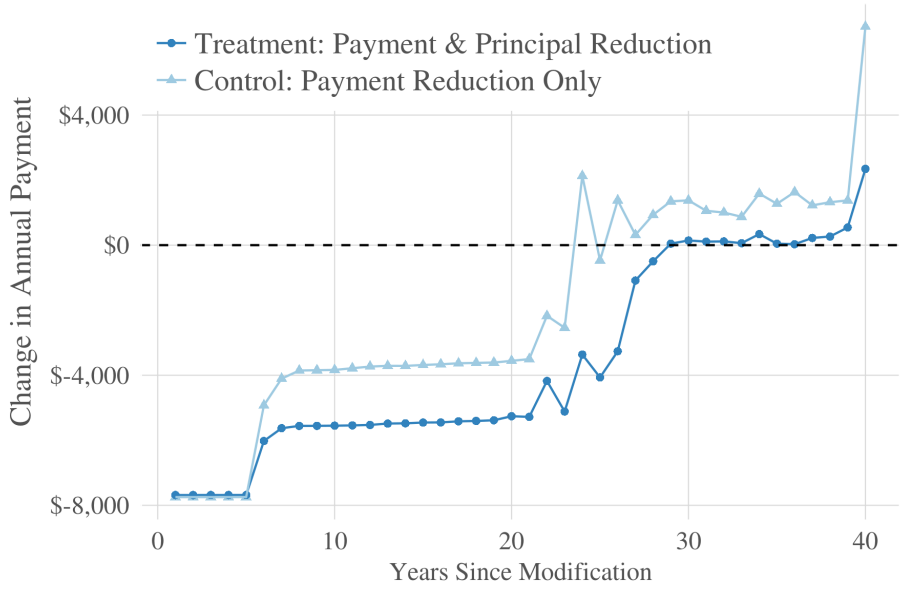

Notes: This figure is an unwinsorized version of Figure 1a. The plot shows the difference in average annual payments for borrowers receiving each type of modification relative to the payments borrowers owed under their unmodified mortgage contracts in the matched HAMP credit bureau dataset. The conventional HAMP waterfall includes interest rate reduction, followed by maturity extension to 40 years, followed by principal forbearance. However, some servicers offer principal forbearance prior to maturity extension, so some borrowers have large payments on the amount forborn due at the end of loan terms between 22 and 27 years. These large payments introduce variability into the average change in payment due.

Figure 22: Actual and Predicted Mortgage Default Around Principal Reduction Discontinuity, Unnormalized Running Variable

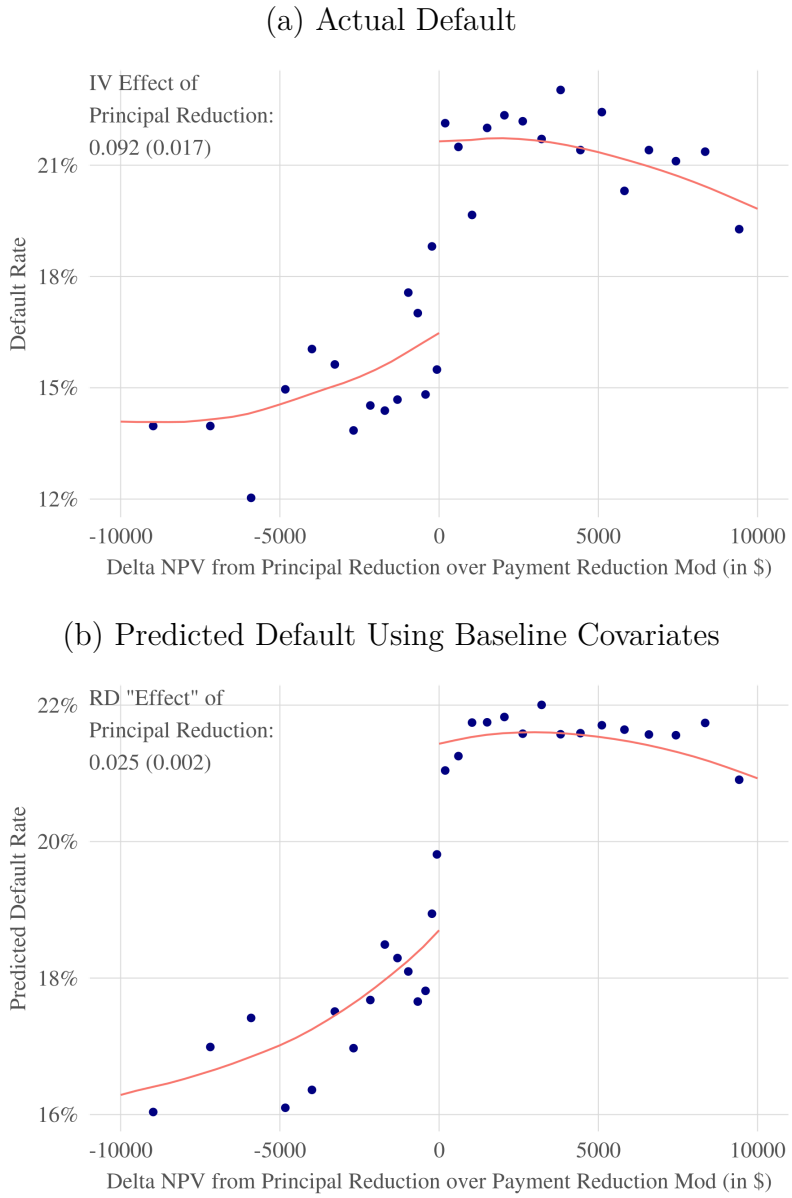

Notes: This figure shows the reduced form and predicted default rates based on pre-determined borrower covariates in the matched HAMP credit bureau dataset using an unnormalized version of the running variable. The horizontal axis shows the unnormalized predicted gain to lenders of providing principal reduction to borrowers, i.e. only the numerator from equation (1). The dots are conditional means for 15 bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. Panel (a) plots the default rate. Construction of the IV estimate $\hat{\tau}$ is calculated as described in Section 3.2, except the unnormalized running variable is used. Panel (b) plots predicted default rates from a linear regression of default on observed borrower characteristics (FICO score, monthly income, payment-to-income ratio, loan-to-value ratio, and months past due prior to modification). The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (2). This specification fails the balance test using baseline covariates; predicted default is *higher* among the group that gets more principal reduction. Thus, this specification is uninformative about the effect of principal reduction on default.

Figure 23: Impact of Principal Reduction on Default, Unmatched Sample

(a) First Stage – Receive Principal Reduction

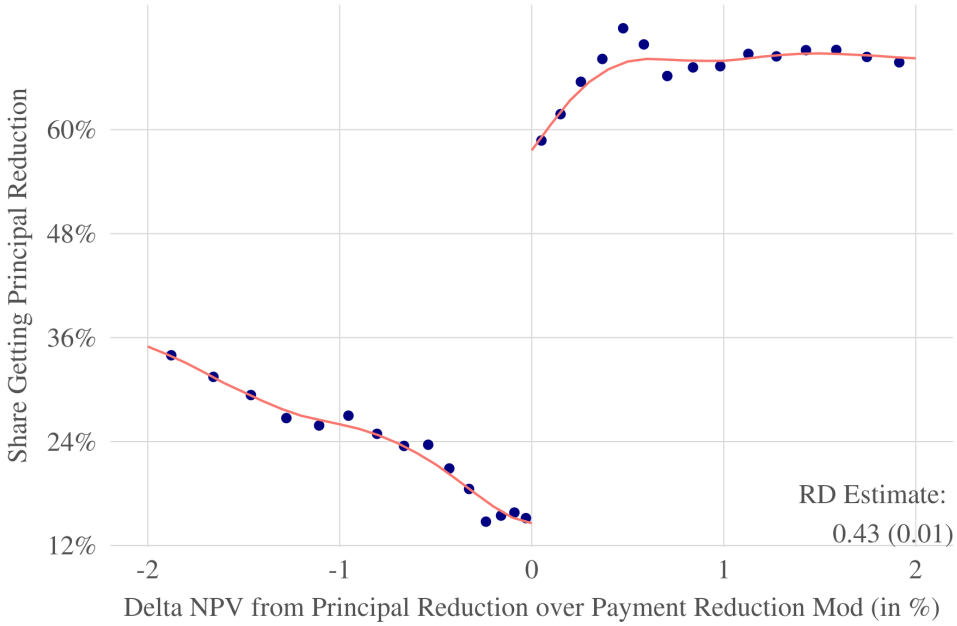

(b) Reduced Form -- Mortgage Default

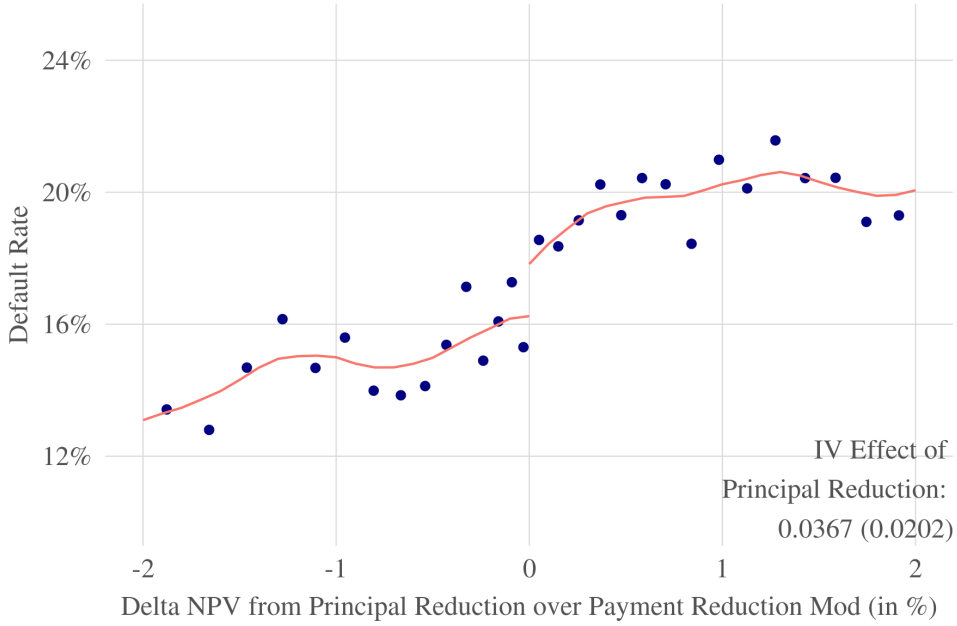

Notes: This figure replicates Figure 2a using the full HAMP public file. This file includes loans which do not have a match in the TransUnion data. See Appendix B.1.3 for details.

Figure 24: Pre-Modification Characteristics around Principal Reduction Discontinuity, Unmatched Sample

(a) Pre-Mod FICO Score

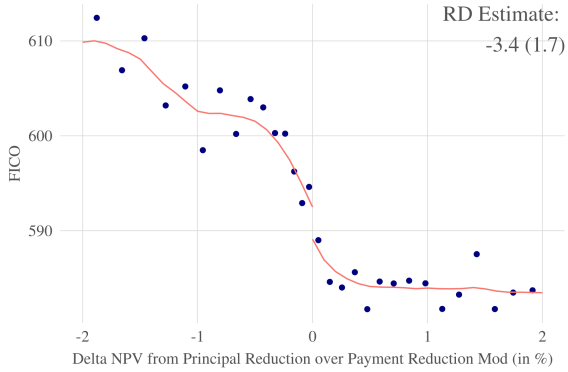

(b) Monthly Income at Mod Date

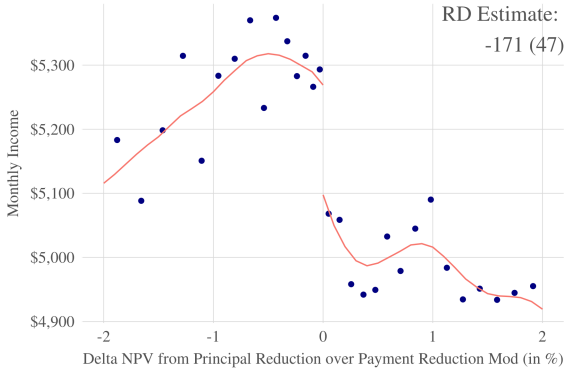

(c) Pre-Mod Payment-to-Income Ratio

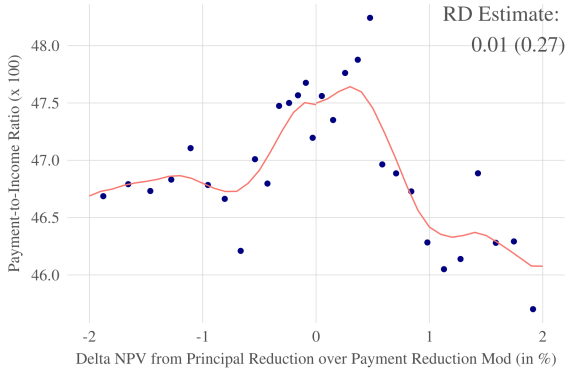

(d) Pre-Mod Mark-to-Market Loan-to-Value

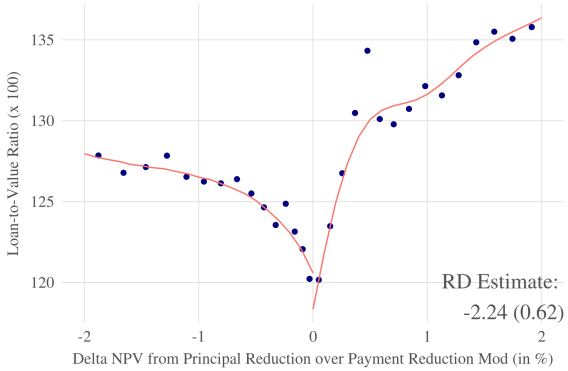

(e) Pre-Mod Months Past Due

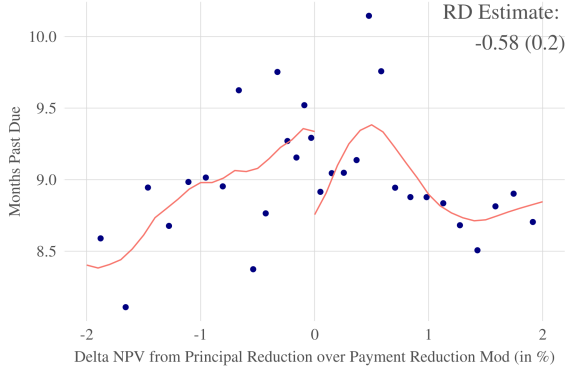

(f) Predicted Default Rate Using Covariates

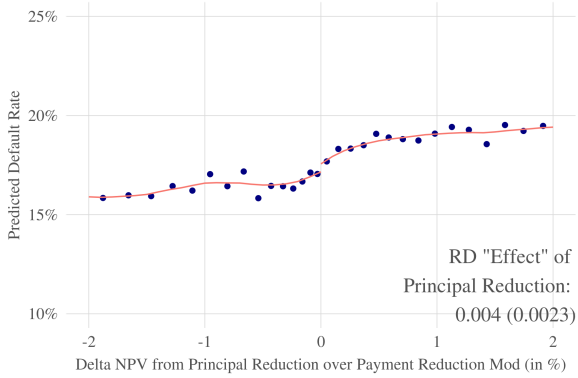

Notes: This figure replicates Appendix Figure 4 using the full HAMP public file. This file includes loans which do not have a match in the TransUnion data. See Appendix B.1.3 for details.

Figure 25: Effect of 10% Payment Reduction on NPV: Robustness

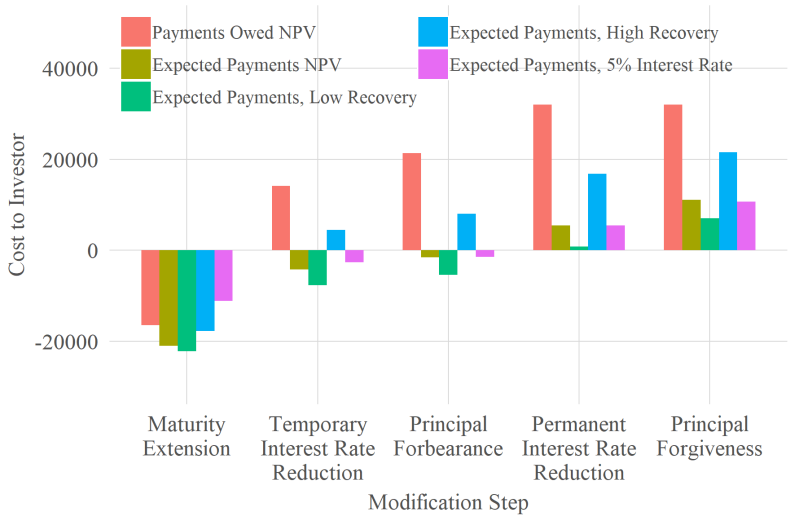

Notes: This figure shows the impact of a 10 percent payment reduction on the NPV of the loan to the investor under various assumptions. The red and yellow bars reproduce Appendix Figure 20a. The yellow bars assume a 39 percent self-cure rate on post-modification defaults and a 56 percent loss if the loan is liquidated. The green bars assume a self-cure rate of 18 percent and a liquidation loss of 61 percent. The blue bars assume a 61 percent self-cure rate and a 48 percent liquidation loss. See Appendix C.3.1 for the data sources for each of these assumption. The purple bars use the same assumptions as baseline, except a 5 percent initial interest rate.

Figure 26: Effect of Principal Reduction on Foreclosure Initiation

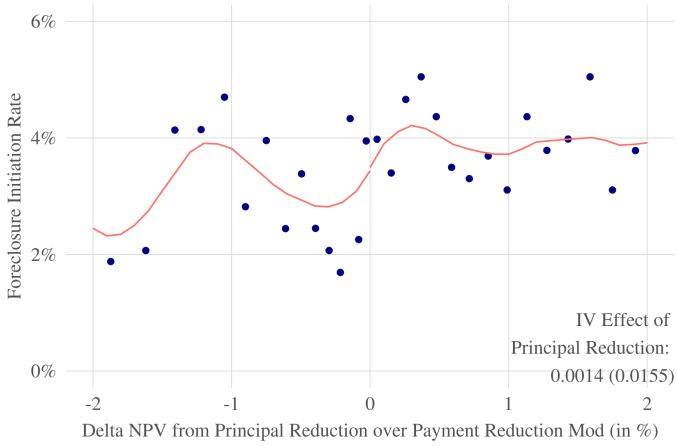

Notes: This figure shows the effect of principal reduction on foreclosure initiation in the matched HAMP credit bureau dataset. The foreclosure initiation rate is plotted on the vertical axis and the normalized predicted gain to lenders of providing principal reduction is on the horizontal axis. The dots are conditional means for 15 bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. Construction of the IV estimate $\hat{\tau}$ is described in Section 3.2.

Figure 27: Effect of Payment Reduction on Credit Card Expenditure Using the Payment Reduction Discontinuity

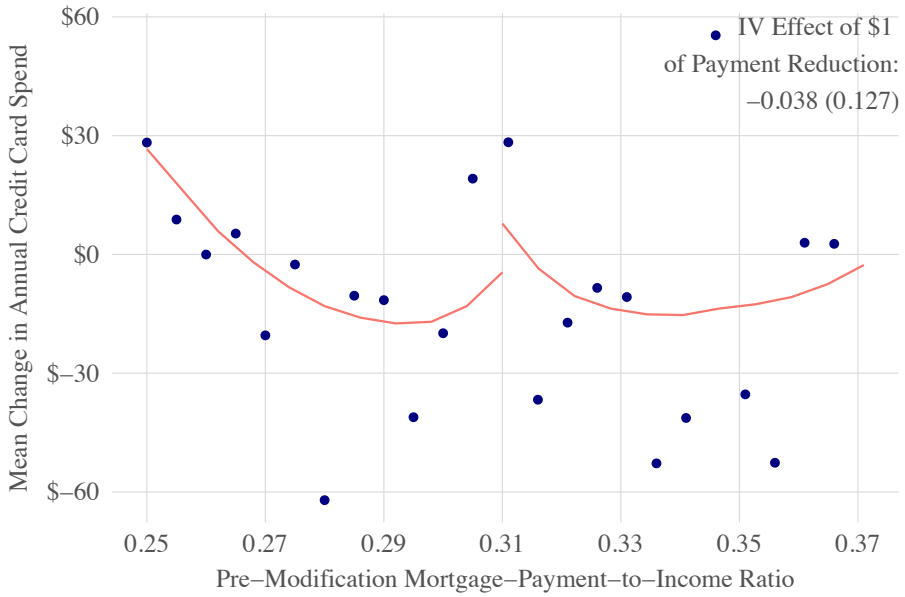

Notes: This figure shows the reduced form of the estimated impact of payment reduction on credit card expenditure using the 31 percent PTI regression discontinuity strategy in the JPMCI bank dataset. The dots are conditional means for 12 bins on each side of the cutoff. The line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. Construction of the IV estimate $\hat{\tau}$ is described in Section 5.2. This strategy is unable to detect economically meaningful changes in expenditure.

Figure 28: Projected 40-Year Mortgage Interest Rates

(a) Actual and Projected Loan Interest Rates

(b) Actual and Modeled Spread Between 30- and 40-year Rates

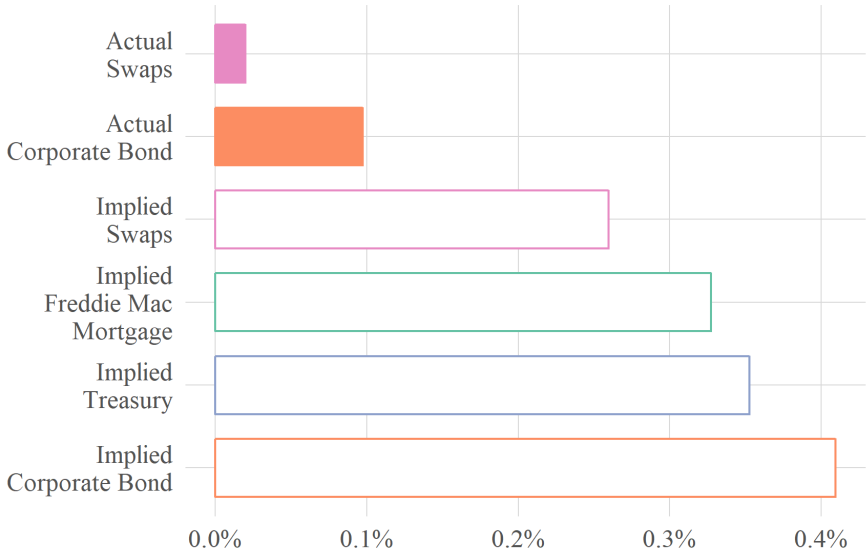

Notes: Panel (a) shows interest rates for various loan terms. Solid dots are data, lines are the best fit of $y = log(x)$ to the solid dots, and hollow dots are projections of 40-year interest rates. Green dots show mortgage rates from the Freddie Mac Conforming Loan Survey, red squares show corporate bond spot rates, blue triangles show Treasury note rates, and purple diamonds show fixed-for-floating interest rate swaps. Panel (b) shows estimates of the interest premium for a 40-year loan over a 30-year loan using four methodologies. It shows a premium of 10 basis points using actual corporate bond spot rates in a solid bar, a premium of 32 basis points extrapolated from shorter-term Freddie Mac mortgage rates in a hollow bar, a premium of 34 basis points extrapolated from shorter-term Treasury rates in a hollow bar, and a premium of 2 basis points using actual swap rates in a solid bar. (For reference, the panel also shows the extrapolated premium using corporate bond rates and swap rates.) See Appendix C.3.1 for calculation details and description of implied swaps and implied corporate bonds.

Figure 29: Default Risk Arising From Maturity Extension Through Additional Time Underwater

(a) Years Spent Underwater

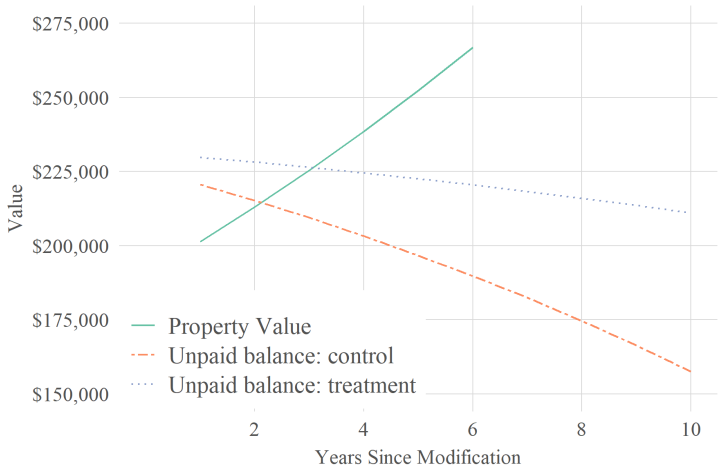

(b) Projected Long-Term Default

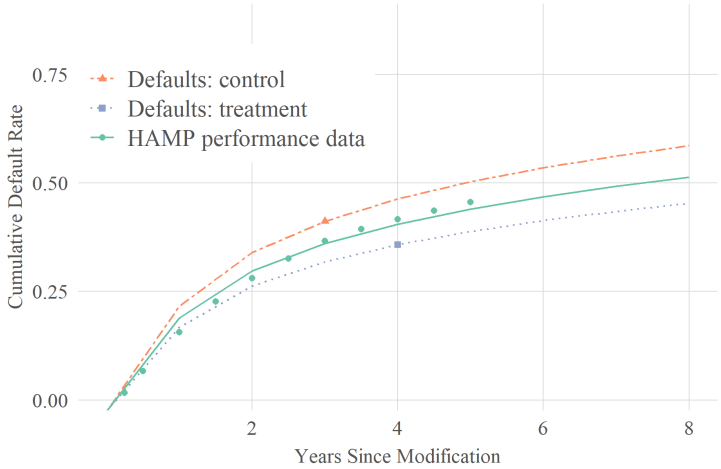

Notes: See Appendix C.3.2.

Figure 30: Amount of Payment Reduction and Default

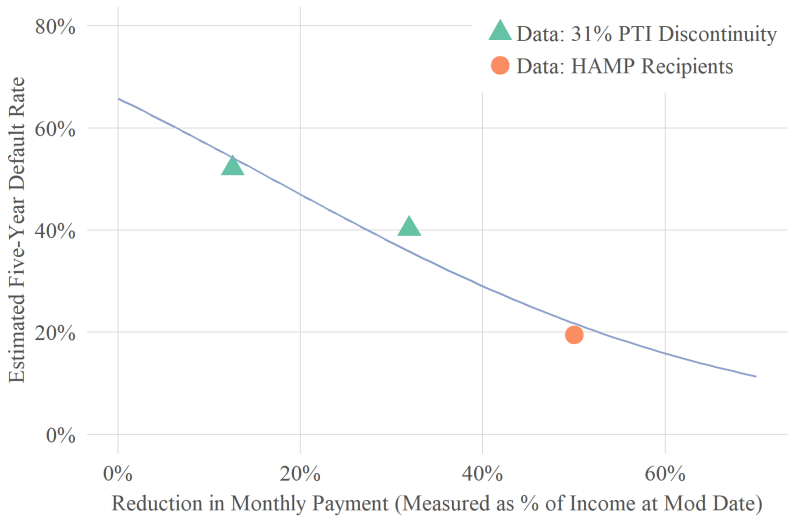

Notes: This figure shows estimated five-year default rates for various amounts of payment reduction. The green triangles are from the two sides of the discontinuity in Figure 5b and the orange circle is borrowers with PTI of 55 percent from Appendix Figure 14. We take the two-year default rates and multiply them by 1.62, which is the ratio of five-year default rates to two-year default rates among HAMP modifications performed in 2010. The line is a best fit of a logit model to the three data points.

Figure 31: NPV Cost of Payment Reduction for Various Sequences of Modification Steps

(a) Add Social Value of Payment Reduction

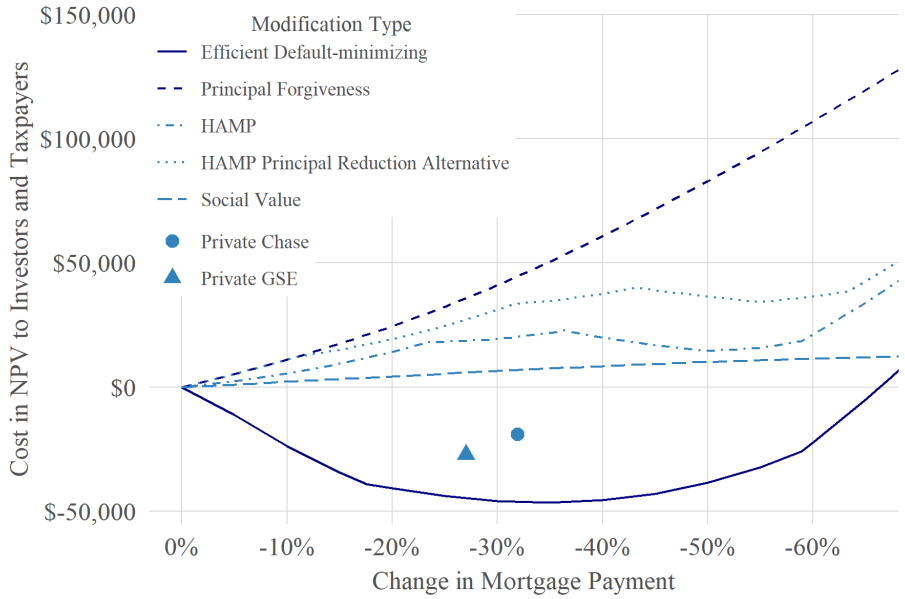

(b) NPV of Payments Owed

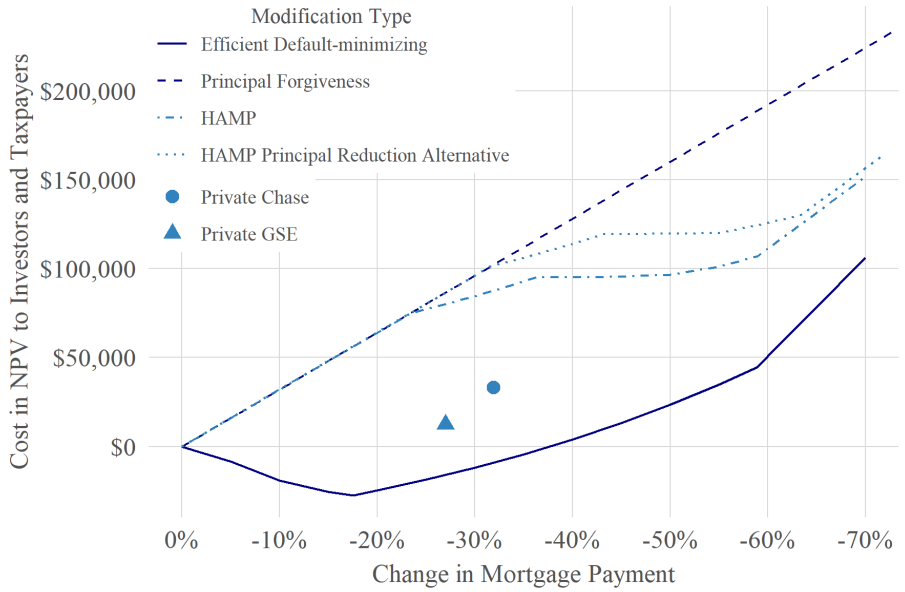

Notes: The top panel takes Figure 20b and adds a line reflecting the social value of payment reduction, assuming a $51,000 social cost per foreclosure as estimated in U.S. Department of Housing and Urban Development (2010). The bottom panel recomputes Figure 20b using the NPV of payments owed.

Figure 32: Consumption Functions with Cash-on-Hand and Collateral Grants at Various Dates

(a) Consumption Function out of Future Cash-on-Hand

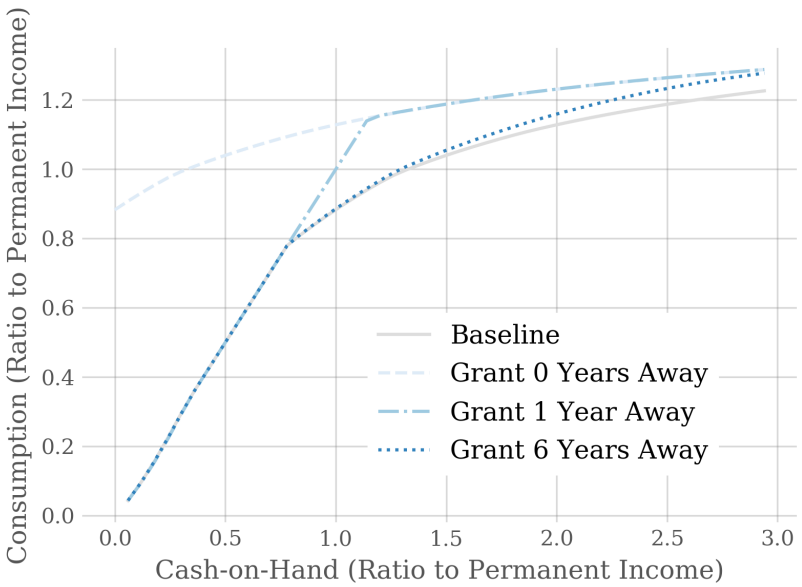

(b) Consumption Function out of Future Collateral

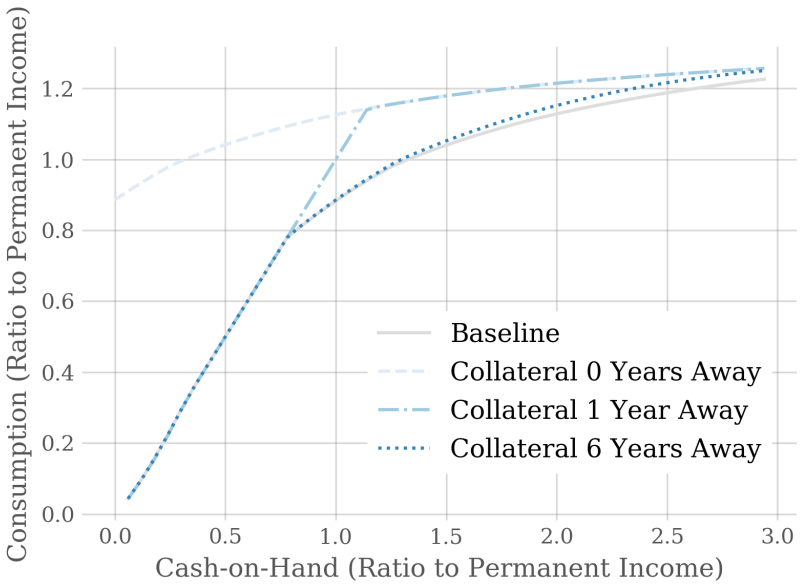

Notes: The top panel plots the consumption function out of cash-on-hand under various alternative scenarios from the model described in Appendix D. Both the horizontal and vertical axes are measured relative to permanent income. The baseline case considers a household with no home equity (and hence no current borrowing capacity). The lines show the consumption functions in the current period when the household is granted one year's worth of permanent income in the current period, in one year, and in six years. The bottom panel shows the equivalent consumption functions for the case when the household is granted collateral, rather than wealth, at various dates.

Figure 33: Consumption Function Out of Principal Forgiveness

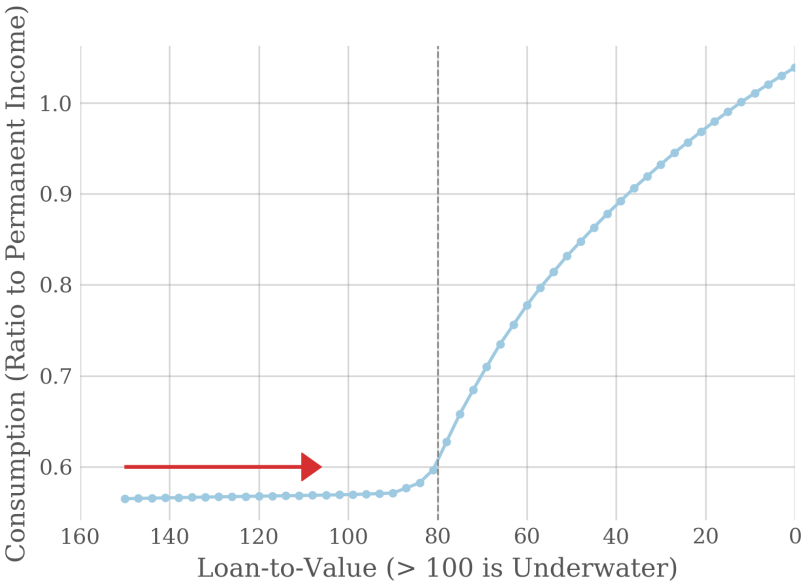

Notes: This figure plots the consumption function out of principal reduction. We begin borrowers at a 150 percent loan-to-value (LTV) ratio and give increasing amounts of principal reduction as necessary to hit a given LTV ratio. To mimic our empirical setting, mortgage payments for households who have not defaulted are fixed for five years, after which payments fall according to the new mortgage balance. The red arrow shows the treatment for the average borrower in the government program. This figure assumes a collateral constraint of 80 percent of LTV and cash-on-hand (assets + annual income) of 86 percent of permanent income. See Appendix D for details.

Figure 34: Mortgage Credit Availability

(a) Mortgage Originations by Credit Score

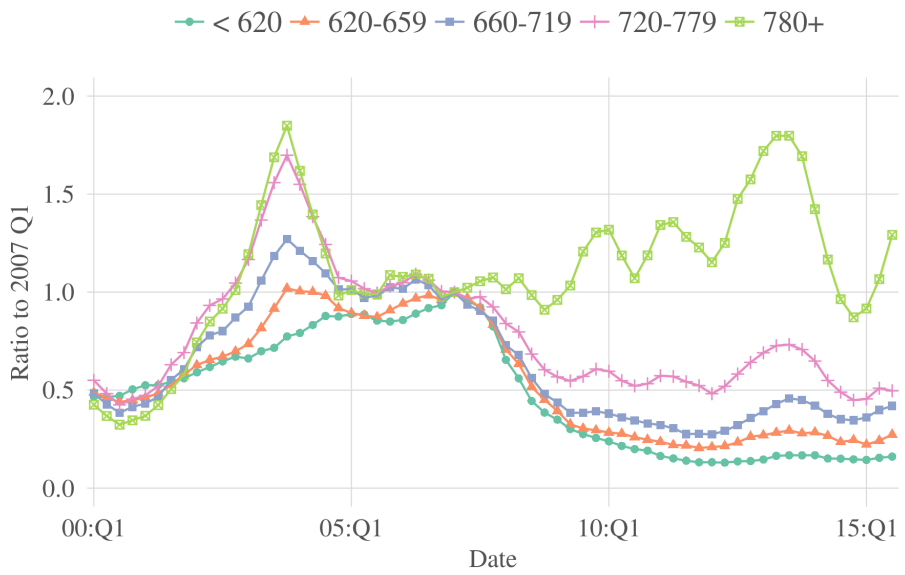

(b) Combined Loan-to-Value for New Home Equity Lines of Credit

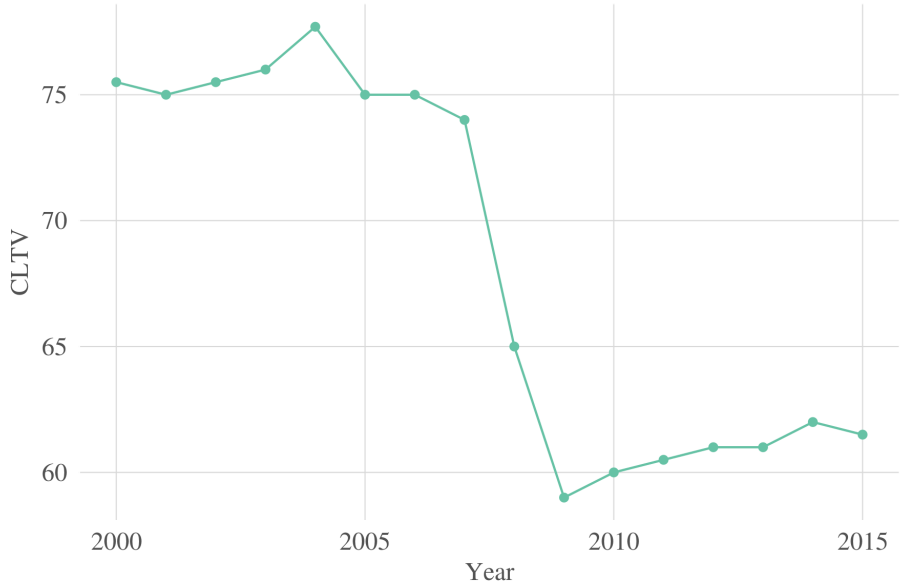

Notes: The top panel plots mortgage origination by borrower credit score from the New York Fed Consumer Credit Panel (Federal Reserve Bank of New York 2015). This includes first mortgages, second mortgages, and home equity installment loans. The bottom panel plots the average combined loan-to-value (CLTV) ratio for new home equity lines of credit (HELOCs) as reported by Corelogic (2016).

Figure 35: Effect of Modeled Principal Reduction on Borrowing Limits and Mortgage Payments

(a) Borrowing Limits

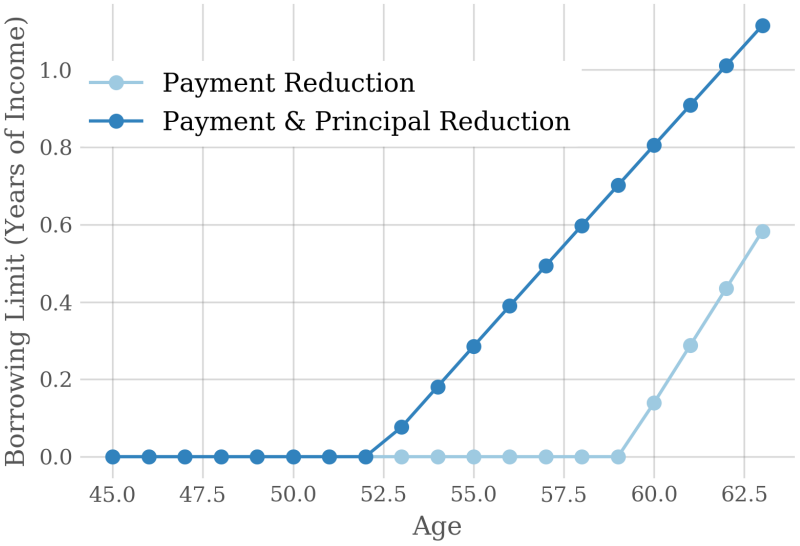

(b) Mortgage Payments

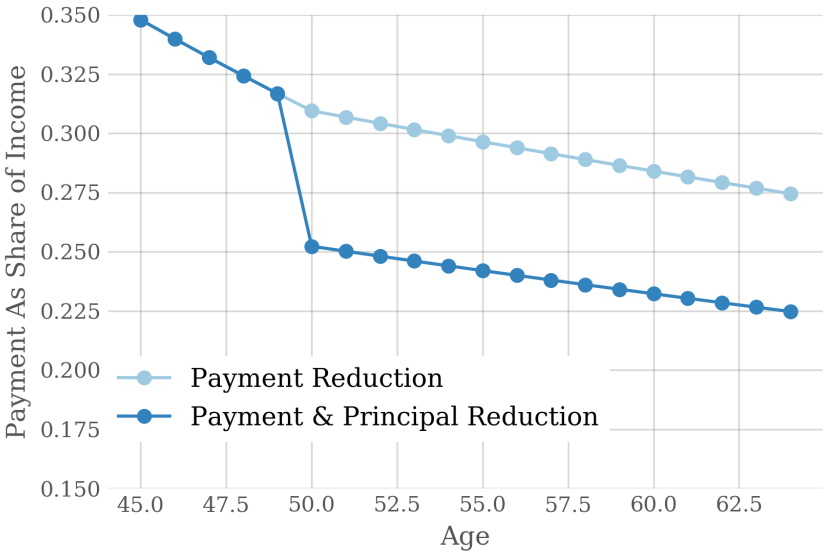

Notes: This figure shows the effect of the modeled principal reduction policy on borrowing limits and mortgage payments. We assume homeowners receive modifications at age 45. We set initial LTV equal to 150. For our treatment group, we then reduce their mortgage balance by $70,000, bringing them to an LTV of 106 in the first year. To mimic our empirical setting, mortgage payments for households who have not defaulted are fixed for five years, after which payments fall according to the new mortgage balance.

Figure 36: Default and Mortgage Debt Levels

(a) Default if Income + Assets < Threshold

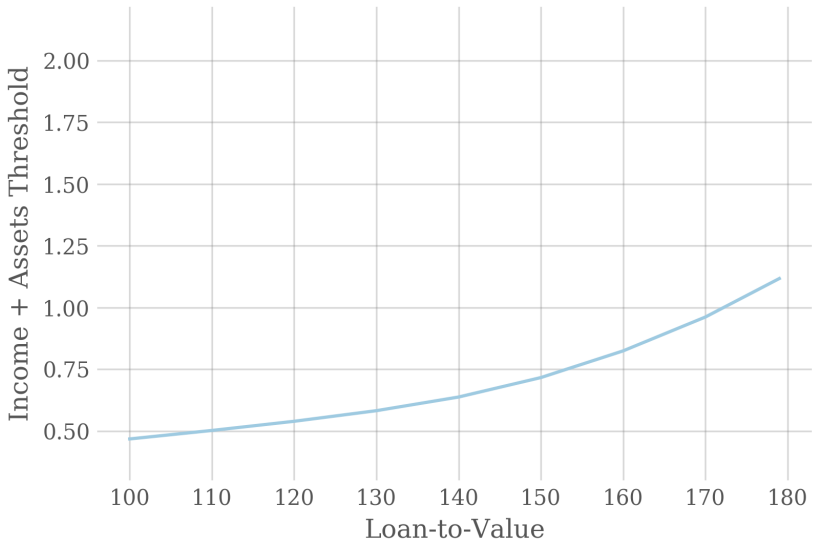

(b) Default Rate

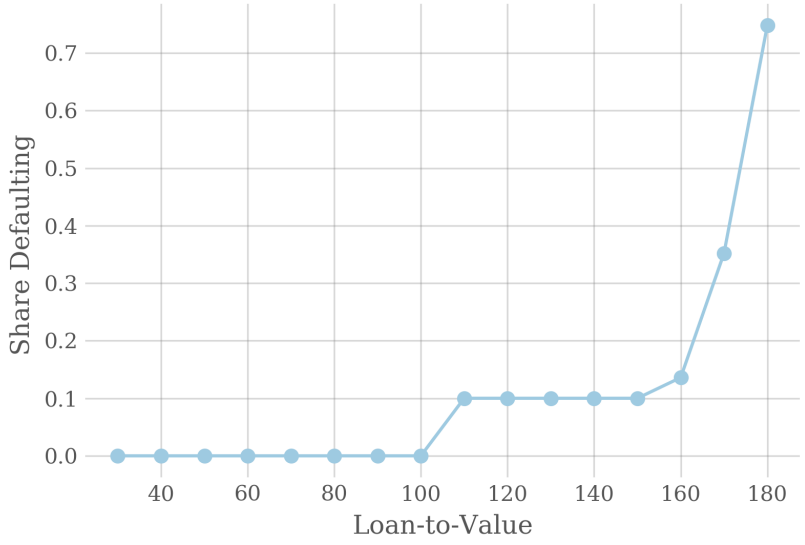

Notes: The top panel plots the cutoff thresholds for borrower default decisions. The vertical axis is relative to permanent income. The line shows the baseline assumptions as described in equation (8). For borrowers with a given LTV ratio, the line shows the cash-on-hand (income plus assets) threshold below which borrowers decide to exercise their default option. The bottom panel plots default rates by LTV ratio. Default rates are calculated by taking the default thresholds shown in the top panel and integrating over the distribution of income shocks described in equation (9).

Figure 37: Default with Heterogeneous Utility Cost of Default

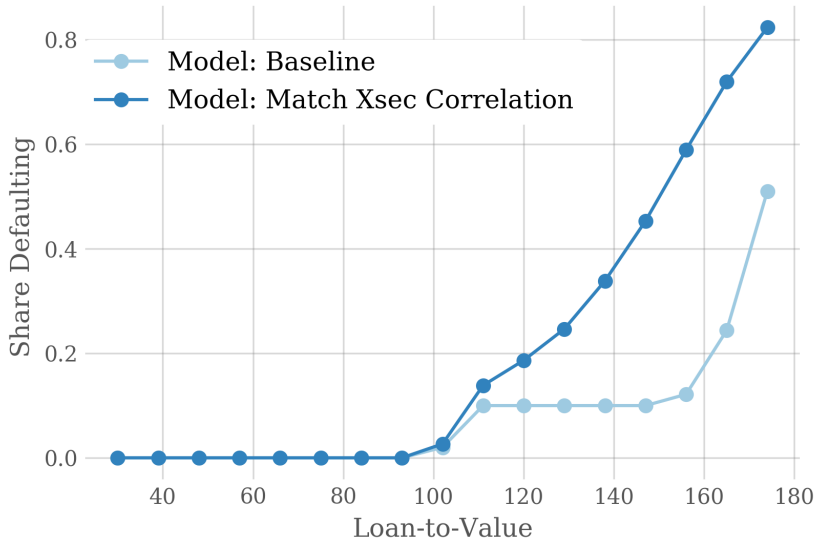

Notes: This figure plots default rates by LTV ratio in the model under alternative parameterizations. The LTV is moved according to the same policy exercise described in the notes to Appendix Figure 35a. The baseline parameterization corresponds to that in Appendix Table 5. The "Match Xsec Correlation" series assumes a distribution of default costs across the population instead of a constant default cost.

Table 1: HAMP Summary Statistics Pre- and Post-Credit Bureau Match

	Pre-Match		Post-Match		Normalized
	Mean	SD	Mean	SD	Difference
Income	54,549	24,638	51,186	23,612	-0.14
Home Value	197,080	123,940	178,250	114,753	-0.15
Loan to Value Ratio	150	35	151	35	0.02
Monthly Mortgage Payment	1,716	875	1,551	789	-0.19
Monthly Payment to Income Ratio	0.48	0.12	0.47	0.12	-0.11
Mortgage Interest Rate	0.063	0.020	0.063	0.020	-0.00
Mortgage Term Remaining (Years)	25.9	4.6	25.8	4.7	-0.01
ARM (d)	0.49	0.50	0.46	0.50	-0.05
Months Past Due	10.9	11.5	9.5	10.5	-0.12
Credit Score	584	74	582	75	-0.03
Male (d)	0.57	0.50	0.56	0.50	-0.02
Age	48.8	10.8	48.6	10.9	-0.01
Monthly Payment Reduction ($)	737	544	641	483	-0.18
Monthly Payment Reduction (%)	42	20	41	20	-0.07
Principal Forgiveness Amount	53,046	70,413	46,025	62,173	-0.10
Received Principal Forgiveness (d)	0.59	0.49	0.59	0.49	-0.01
Post Modification LTV	134	34	135	35	0.03
Post Modification DTI	0.30	0.04	0.30	0.04	-0.05
Post Modification Default (d)	0.201	0.400	0.201	0.400	0.00
N	221,487		105,365		

Notes: This table shows characteristics for all HAMP borrowers who were underwater and evaluated for both modification types during our sample window. Our regression discontinuity and panel difference-in-differences analyses each use different subsets of the matched sample. The normalized difference in the final column is the difference in means divided by the pre-match standard deviation. All values are before modification unless otherwise noted. (d) indicates a dummy variable.

Table 2: Summary Statistics for Difference-in-Differences Analysis

	Payment Reduction		Payment and Principal Reduction	
	Mean	SD	Mean	SD
Principal Forgiveness Amount	13,257	36,479	80,460	64,675
NPV Payment Reduction	63,000	60,050	97,397	74,112
Monthly Payment Reduction ($)	676	478	674	505
Monthly Payment Reduction (%)	38.5	18.3	41.9	21.4
Loan to Value Ratio	150	33	152	37
Post Modification LTV	148	34	122	29
Monthly Payment to Income Ratio	0.47	0.11	0.47	0.12
Post Modification DTI	0.31	0.03	0.30	0.05
Income	55,597	23,753	52,924	23,692
Credit Score	598	83	578	72
Home Value	205,275	118,748	174,994	111,218
Monthly Mortgage Payment	1,725	803	1,593	780
Mortgage Interest Rate	0.061	0.018	0.064	0.019
Mortgage Term Remaining (Years)	26.2	4.5	26.1	4.4
Male (d)	0.58	0.49	0.55	0.50
Age	48.3	11.2	48.8	10.8
N	35,485		33,457	

Notes: This table shows summary statistics for the matched HAMP credit bureau sample analyzed in the panel difference-in-differences research design discussed in Section 4. The sample includes underwater borrowers who are observed in the credit bureau records one year before and after modification and report positive credit card expenditure in at least one month during this window. All variables are before-modification unless otherwise noted. (d) indicates a dummy variable.

Table 3: Impact of Principal Reduction on Credit Card Expenditure Using Bank Data

This table reports difference-in-differences estimates of the effect of principal reduction on credit card expenditure in the JPMCI bank account dataset. The coefficient of interest, *Treatment*, is the estimated change in the difference between outcomes of mortgages receiving modifications with and without principal reduction during the year after modification. All specifications include fixed effects for modification type and months since modification. Controls include pre-modification loan characteristics (LTV, principal balance), property value, and LTV at origination. The sample includes all HAMP borrowers with a mortgage and a credit card with Chase who are observed one year before and after modification. The dependent variable mean in the year before modification is reported for borrowers receiving principal reduction modifications. Standard errors, in parentheses, are clustered at the borrower level ($n_{borrower} = 10,741$). See the text for additional detail on the specification, outcome measures, and sample.

Credit Card Expenditure ($/month)

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (Principal Reduction x Post)	0.932	0.807	5.658	−0.524	−5.828	−2.227
	(6.969)	(7.229)	(8.921)	(9.966)	(10.002)	(10.210)
MSA Fixed Effects		Yes				
Calendar Month Fixed Effects			Yes			
MSA by Calendar Month Fixed Effects				Yes	Yes	Yes
Controls					Yes	Yes
Controls x Post Interactions						Yes
Dependent Variable Mean	228.02	232.14	228.02	232.14	232.14	232.14
Observations	268,525	254,084	268,525	254,084	254,084	254,084
Adjusted R^2	0.001	0.020	0.006	0.011	0.025	0.026

Table 4: Pareto Improvement at Payment Reduction Discontinuity: Robustness

Scenario	dNPV($)	dNPV (%)	Breakeven Discount Rate (%)
Preferred Estimate	6229	3.83	6.13
Robustness to Default Assumptions			
Low Default Reduction	2597	1.60	5.37
High Default Reduction	9861	6.07	6.77
Optimistic Recovery	-1028	-0.63	4.53
Pessimistic Recovery	9106	5.60	6.93
Robustness to Discounting Assumptions			
Flatter Yield Curve (Actual Swaps Spread)	10137	6.24	6.13
Steeper Yield Curve (Implied Treasury Spread)	5917	3.64	6.13
Discount at Treasury Rates	9838	6.05	6.13
Discount at Swap Rates	15845	9.75	6.13
Robustness to Prepayment Assumptions			
Low Prepayment	6064	3.73	5.97
High Prepayment	7014	4.31	10.0
Crosswalk to Payments Owed NPV			
Payments Owed	-2168	-1.33	3.59
Payments Owed, with Default	9053	5.57	5.70
Payments Owed, with Default and Yield Curve	5279	3.78	5.70

Notes: This table assesses the change in the Net Present Value (NPV) of expected payments to the lender of assigning a mortgage to the left-hand side of the 31 percent Payment-to-Income discontinuity instead of the right-hand side for a variety of scenarios. It also reports the percent change in the NPV and the annual discount rate a lender would need in order to be indifferent between assigning a mortgage to treatment or control. The baseline specification incorporates default risk, prepayment risk, and the yield curve. The first three panels of the table vary the assumptions about the probability of default, the recovery rate given default, the rate used to discount cash flows, and the prepayment rate. The final panel crosswalks the baseline specification to the alternative "Payments Owed NPV" discussed elsewhere in the text. See Appendix C.3.1 for details.

Table 5: Baseline Model Parameter Values

Description	Parameter	Value	Source
Life-cycle income growth	Γ_s	1.025 to 0.7	Carroll (1997)
Std. dev. income shocks	σ_δ	0.14	Carroll (1992)
Large income shock probability	p	0.1	Guvenen et al. (2014)
Large income shock size	b	0.5	Guvenen et al. (2014)
Real interest rate	r	0.02	Freddie Mac
Collateral constraint	ϕ	0.2	FHFA, Corelogic
Real house price growth	g	0.009	FHFA 1990-2010
Property tax rate	τ_p	0.015	Himmelberg et al. (2005)
Maintenance costly	τ_m	0.025	Himmelberg et al. (2005)
Utility cost of default	ψ	5.4	Match 10% Default
Risk aversion	γ	4	
Discount factor	β	0.96	

Notes: see Appendix D.3.6 for details.

Table 6: Housing Wealth MPC in Model and External Benchmarks

	MPC (Cents)		
	Model	External Benchmark	Source
Average	8	9	Mian, Rao, Sufi (2013)
LTV = 95	15	18	Mian, Rao, Sufi (2013)

Notes: This table shows the marginal propensity to consume out of changes in housing wealth in the model relative to the estimates in the external benchmark from Mian, Rao and Sufi (2013) (adjusted for homeowners). The model estimates are for age 45 borrowers with different initial LTVs. We endow each agent with cash-on-hand equal to two years of permanent income, which is the median non-housing wealth for all homeowners in the 2007 SCF (2007 is chosen as the base year to mimic estimates in Mian, Rao and Sufi (2013), which cover the 2006-2009 period). We then calculate the MPC for these agents at different LTV values. The "Average" row weights MPCs by LTV according to the distribution of LTV in 2007 reported in Carter (2012).

Table 7: MPC out of Principal Reduction in the Model

Scenario	MPC (cents)
Data	0.3
Model Parameterizations with Small Response	
Baseline Model	0.3
Low Cash-on-Hand	0.0
Age At Mod = 35	0.9
High discount rate ($\beta = 0.9$)	0.8
Low risk aversion ($\gamma = 2$)	0.9
Unused HELOCs	0.9
Model Parameterizations with Larger Response	
High Cash-on-Hand (PIH)	3.4
Collateral Constraint ϕ= 0	4.8
Expected 5% House Price Growth	6.2
Expected 5% House Price Growth and $\phi = 0$	24.2
Alternative Policy Simulations	
Write Down to 90% LTV	1.0
Write Down to 90% LTV and $\phi = 0$	14.1

Notes: This table compares the MPC out of principal reduction in the model under alternative parameterizations to the MPC calculated in our data (discussed in Section 4.2). The "Baseline Model" corresponds to the parameterization shown in Appendix Table 5 and the modeling of principal reduction policy discussed in Section D.3.4. "Low Cash-on-Hand" corresponds to initial cash-on-hand $m_t = 0.5$ units of permanent income. The "Unused HELOCs" row corresponds to an experiment where the household is given a credit line worth \$20,000 (or 0.38 units of permanent income), and then given principal reduction. The "High Cash-on-Hand (PIH)" row corresponds to initial cash-on-hand $m_t = 3.0$ units of permanent income. The "Expected 5% House Price Growth" row corresponds to an expected permanent annual real house price growth of 5 percent.

B Empirical Appendix

B.1 Effect of Principal Reduction

B.1.1 Sample Construction and Cleaning

HAMP Public File Analysis – We study permanent HAMP modifications performed between December 2010 and December 2014 for loans that originated prior to 2009 using the HAMP 1st lien file and the HAMP NPV file.

We restrict the sample to identify loans that are candidates for principal reduction. We drop loans which are owned by the GSEs because the GSEs did not allow principal reduction. To construct the running variable V we require both an estimated "payment reduction" NPV and an estimated "payment and principal reduction" NPV. Following Scharlemann and Shore (2016), we drop one servicer whose principal reduction forgiveness allocation method exhibits no discontinuity. Further, for reasons we describe in Section 3.2, we require $V \neq 0$.

We also restrict the sample to loans that can plausibly be matched to the TransUnion credit bureau data. We require non-missing geographic location, pre-modification monthly payment that is non-missing and less than $25,000, and positive non-missing pre-modification unpaid balance.

We then match loans between TransUnion and the public HAMP file. We require exact matches for origination year and geography (MSA when available in the HAMP file, otherwise state), and then take the closest match using normalized Euclidean distance on modification month, pre-modification monthly payment, post-modification monthly payment, and pre-modification principal balance. We keep matches with Euclidean distance less than a threshold of 0.2. Above this threshold, match quality deteriorates.

We clean the matched sample in three steps. First, we measure principal forgiveness using `ln_upb_frgv_amt` in the HAMP 1st Lien file. If it is missing there we use the value reported in the HAMP NPV file.

Second, we adjust the pre-modification mark-to-market-loan-to-value (MTMLTV) ratio in two ways. Servicers were supposed to report the value in percent (e.g. "100" for MTMLTV of 100 percent), but many instead reported the correct value divided by 100 (e.g. "1" for MTMLTV of 100 percent). To correct for this, we multiply reports less than 2 by 100. In addition, we winsorize MTMLTV at 200 percent.

Third, we also winsorize some other variables with substantial outliers: the change in MTMLTV at the 1st percentile of nonzero values, the amount of principal forgiveness at the 99th percentile of nonzero values, monthly payments (both pre-modification and post-modification) at the 99th percentile, pre-modification payment-to-income ratio at the 99th percentile, pre-modification months past due at the 95th percentile, borrower income at the 95th percentile, monthly auto spending at the 99th percentile of nonzero values, and monthly credit card spending at the 99th percentile of nonzero values.

TransUnion Consumption Analysis – We construct a month-to-month measure of credit card expenditure using balance and payments data as reported in TransUnion. We require that the servicer reports month to month payment amounts: we exclude any servicer that never reports payment amounts or for whom over 90 percent of the card-months are associated with payments of zero dollars. This requirement retains 83 percent of the credit cards ever observed in the sample. Among servicers that do regularly report payments, if the payment field is missing, we record this as a payment of zero for the month. We then define credit card expenditure as the balance change between two given months in addition to the payment made that month. Let b_t denote the balance at the end of month t, and p_t be the

payment made in month t. We calculate expenditure in month t as $e_t = b_t - b_{t-1} + p_t$. If this calculation implies negative expenditure on a card, we record the credit card expenditure as zero. We sum the expenditure over all cards associated with a customer for a total measure for that customer.

JPMCI Consumption Analysis — We exclude customers with more than one mortgage and customers with GSE-owned loans. Furthermore, we exclude mortgages that were not underwater prior to modification.

B.1.2 Running Variable: Understanding Variation Arising from the Treasury NPV Model

This section explains what drives variation in the running variable V and accordingly what type of borrowers are likely to have a value of the running variable which is close to the regression discontinuity cutoff.

The Treasury NPV model is designed to value the expected cash flows to investors for a loan with various attributes. The Treasury published a 70-page document describing the model (U.S. Department of the Treasury 2015) and nearly all the model parameters are public. The model makes the simplifying assumption that if a loan defaults, it will default immediately (and be liquidated). If it doesn't default, it will pay on schedule with some probability of prepayment each period. Thus, the NPV of a loan is the weighted average of the NPV in the "pay" state and the NPV in the "default" state, with the weights given by the probability of each state, that is

$$NPV = (1 - p_{default}) \cdot NPV\{Loan\ Pays\} + p_{default} \cdot NPV\{Loan\ Defaults\}, \qquad (5)$$

The equation above is the first equation in Section V of the NPV model documentation.

To allocate borrowers between the two types of HAMP modifications, this model is evaluated twice. First, it is evaluated assuming the mortgage attributes for a standard "payment reduction modification". Second, it is evaluated assuming the mortgage attributes for a "payment and principal reduction modification." For the discussion below, it will be useful to define some notation. The running variable V is the normalized difference between these two model evaluations, $p_{default}^{no\ prin\ red}$ is the probability of default under a standard "payment reduction modification", and Δp is the difference in default rates between the two modification types.

Relative to a standard payment reduction modification, a principal reduction modification has a benefit and a cost. The benefit is that the probability of default is lower and so puts more weight on the high return state of the world (where the loan pays). This benefit is higher when the default reduction Δp is greater and when the difference between the high return state of the world and the low return state of the world is greater.

The cost is that by forgiving principal it reduces $NPV\{Loan\ Pays\}$, the cash flows to investors in this higher return state of the world. This cost is lower when this higher return state of the world is unlikely without principal reduction (i.e., when $1 - p_{default}^{no\ prin\ red}$ is lower). In other words, because the cost of principal reduction is only incurred if the borrower would have repaid this principal, principal reduction is less costly to the investor when this good state of the world was unlikely to occur. Intuitively, an investor who forgives principal gives up more cash flows when borrowers are expected to actually pay this principal than when they are expected to default and never pay it. As a concrete example, suppose that principal reduction always reduces default rates by 10 percentage points ($\Delta p = 0.1$) and that an investor has two groups of loans, where group (a) has $p_{default}^{no\ prin\ red} = 0.1$ and group (b) has

$p_{default}^{no\,prin\,red} = 0.5$. In group (a), she must give up cash flows on nine non-defaulting loans to avoid one default, while in group (b), she must give up cash flows on five non-defaulting loans to avoid one default.

Thus, there are four main forces that could explain why some borrowers have high V and some have low V. First, it could be that the Treasury model assumes that principal reduction is more effective at reducing default for some borrowers than for others (heterogeneity in Δp). Second, it could be that the gains from avoiding default are higher for some borrowers than for others (heterogeneity in $NPV\{LoanDefaults\}/NPV\{LoanPays\}$). Third, it could be that some borrowers are assumed to have higher default rates without principal reduction (heterogeneity in $p_{default}^{no\,prin\,red}$). Fourth, when a borrower is eligible for more principal reduction, the model will have a stronger opinion—positive or negative—about whether principal reduction is valuable to the investor. In Treasury's model, the second, third, and fourth forces are more important than the first.

Before discussing how these forces affect V in detail, we begin by describing how the Treasury model predicts default rates. The model uses a logit function $p_{default} = exp(\ell)/(1+exp(\ell))$ where log-odds $\ell = \beta X$. The model has five X variables:

1. the borrower's credit score,
2. the loan's initial payment-to-income (PTI) ratio,
3. the loan's LTV ratio,
4. the amount of short-term payment relief provided by the modification (as captured by the change in the PTI ratio),
5. and days-past-due (the number of days delinquent at the date of modification).

The model is additively separable with respect to the first four variables. It is specified separately for borrowers who are current, 30 days past due, 60 days past due, and 90+ days past due. These two facts mean that a loan's LTV and days-past-due are sufficient statistics for predicting the impact of principal reduction in log-odds terms. *Within* each days-past-due group, the effect of principal reduction on ℓ is approximately constant. This assumption is based on historical data relating default and LTV, as discussed in Holden et al. (2012). *Across* groups, however, principal reduction is assumed to be more effective for borrowers who are current or 30 days past due at the date of modification.

Empirically, we find that the expected average change in default from principal reduction (Δp, the first force) is approximately constant. The only variable that meaningfully affects Δp is days-past-due, but there is little variation in days-past-due with respect to the running variable.

In contrast, heterogeneity in $p_{default}^{no\,prin\,red}$ is important. One example of what drives heterogeneity in $p_{default}^{no\,prin\,red}$ is that borrowers with lower FICO scores are expected to have higher default probabilities, keeping all other characteristics constant. These borrowers will have higher values of V because principal reduction is most beneficial when $p_{default}^{no\,prin\,red}$ is higher. This pattern is visible in Appendix Figure 4a where higher values of the running variable V are associated with lower mean FICO scores.

The third force explaining variation in V is heterogeneity in the return to the investor in the default and non-default states of the world. In the Treasury NPV model, the ratio $NPV\{LoanDefaults\}/NPV\{LoanPays\}$ depends mostly on macro factors such as expected house price growth and the discount at which foreclosed properties are being sold. This ratio

affects V because the investor return to principal reduction is larger when the bad state that principal reduction is expected to help avoid (default) is costlier to the investor. Thus, some of the variation in V is also coming from differences across loans depending on their location and time period of modification, where high values of V are associated with loans in more distressed local markets (Holden et al. 2012).

The fourth and final force explaining variation in V is how much principal reduction a borrower is eligible for. The "payment and principal reduction" modification wrote down mortgage principal until LTV reached 115, as we discuss in Section 3.1 of the paper. When a borrower is eligible for an especially large amount of principal reduction, the model will have a relatively strong opinion about whether principal reduction is a good idea. This is because the marginal gain (or loss) from any principal reduction is magnified. However, when a borrower is eligible for a moderate amount of principal reduction, the model will be closer to indifferent. The borrowers who are eligible for a moderate amount of principal reduction are those whose pre-modification LTV is closer to the target of 115. This explains why pre-modification LTV exhibits a "V" shape with respect to the running variable as shown in Appendix Figure 4d. Nevertheless, the amount of principal reduction at the RD is still substantial. Borrowers in the treatment group received principal reduction equal to 11 points of LTV, or more than $30,000.

B.1.3 Robustness

Balance Plots – Pre-determined covariates trend smoothly through the cutoff, as shown in Appendix Figure 4. The first five panels show the distribution of pre-modification borrower credit score, monthly income, monthly mortgage payments to monthly income (payment-to-income, or PTI) ratio, LTV ratio, and months past due around the cutoff. In all cases these borrower characteristics trend smoothly. The RD estimates of the discontinuous change in these variables at the cutoff, corresponding to the numerator of equation (2), are reported on the figures. For three variables (credit score, monthly income, and PTI) the sign points to slightly worse-off borrowers to the right of the cutoff, while for two variables (LTV and months past due) the sign points to better-off borrowers to the right of the cutoff. The lack of any systematic correlation supports the validity of the design. The only covariate with a statistically significant jump at the 95 percent level is LTV, and even here the jump is not economically significant.

Lee and Lemieux (2010) note that when there are many covariates, some discontinuities will be significant by random chance. They recommend combining the multiple tests into a single test statistic. We implement a version of this by using all five pre-modification covariates to predict default, and we test whether there is a jump in this pooled predicted default measure at the cutoff. The result is shown in the last panel of Appendix Figure 4. There is no significant change in predicted default at the cutoff.

Density – Another relevant issue in regression discontinuity settings is the possibility that the running variable could be manipulated (McCrary 2008). The usual test is to plot a histogram of the running variable to examine whether there is an unusual increase in mass to the right of the cutoff. We show such a plot in Appendix Figure 5a. While the density is smooth on either side of the cutoff, there is a large bulge exactly at zero.

There are four reasons why we believe the bunching of borrowers at zero is not a challenge for the validity of our research design. First, program officers in charge of the dataset at the U.S. Treasury Department informed us that this bulge is a data artifact. If a borrower was ineligible for principal reduction (e.g. because her home was not underwater), servicers were not supposed to compute the NPV of a "payment and principal reduction" modification which

in our notation is $ENPV(1, X)$. Instead, it appears that several servicers reported their calculation for $ENPV(0, X)$ in the $ENPV(1, X)$ field, such that $ENPV(0, X) = ENPV(1, X)$ which in turn implies $V = 0$. Consistent with this theory, we find that observations with $V = 0$ are disproportionately likely to be above water and disproportionately likely to have "0" as the potential dollar value of principal forgiveness submitted to the NPV model.

Second, the conventional economic environment that would incentivize manipulation is not relevant here. Servicers have no economic incentive to manipulate the running variable because they receive the same compensation regardless of which modification is offered.

Third, even if servicers did have an economic incentive to manipulate, that incentive would not vary discontinuously at this cutoff: principal reduction provision is optional regardless of the outcome of the calculation.

Fourth, were servicers manipulating the running variable to zero in an attempt to rationalize principal reduction, they failed; the share of borrowers receiving principal reduction in this zero group is actually half what it is for borrowers with actual positive values of the running variable.

We were advised by U.S. Treasury staff to remove these observations as reflecting measurement error. We attribute the bunching of borrowers at zero to data mis-reporting and drop observations exactly at zero. Appendix Figure 5b shows the distribution for the resulting sample, which is our analysis sample. There is no noticeable change in density around the cutoff.

We show in Appendix Figure 5c that borrower take-up rates were high on both sides of the discontinuity. Ninety-seven percent of borrowers who are offered a modification take it up, and this trends smoothly around the cutoff. This provides further evidence against borrower manipulation to obtain one or the other modification type.

Alternative Bandwidths – Appendix Figure 7 tests the sensitivity of our results to the bandwidth chosen for the local linear regression. Our central estimates are constructed using the optimal bandwidth from the Imbens and Kalyanaraman (2012) procedure, which is 0.61. The optimal bandwidth recommended by the Calonico, Cattaneo and Titiunik (2014) procedure is 0.63. The point estimate begins to rise at wider bandwidths. The rise at wider bandwidths is not surprising given the shape of the estimated conditional expectation function for default, which is particularly sloped near the cutoff. Wider bandwidths will lead to specification error when this function is particularly steep near the cutoff. A quadratic specification which can more easily mimic this slope is stable for a wider bandwidth, showing a point estimate around zero up to a bandwidth of 1.5 before rising.[2]

Alternative Outcome (Foreclosure Initiation) – Our evidence suggests that principal reduction is also ineffective at reducing foreclosures. Appendix Figure 26 shows that there is no jump in the foreclosure initiation rate at the cutoff. Due to the lengthy delay between foreclosure initiations and foreclosure completions, foreclosure completions are rarely observed in our sample period. The most optimistic point in the 95 percent confidence interval suggests foreclosure initiations were reduced at most by 3.1 percentage points. It is not surprising that the same pattern would be seen in foreclosures as in defaults. Borrowers who have defaulted and are unable to self-cure are generally unable to sell their home to avoid foreclosure while they are underwater.

Alternative Normalization of Running Variable – In Appendix Figure 22, we show the change in actual default as well as the change in predicted default using baseline (pre-

[2]The optimal bandwidths for a quadratic specification from Imbens and Kalyanaraman (2012) and Calonico, Cattaneo and Titiunik (2014) are 0.8 and 1.0, respectively.

determined) covariates when the running variable is specified as the change in NPV in dollars (rather than the percent change). This specification fails the balance test using baseline covariates and is thus uninformative about the effect of principal reduction on default.

Unmatched Sample – In principle, the regression discontinuity analysis of principal reduction can be done using the public HAMP file alone, without matching to TransUnion. Appendix Figures 23 and 24 show that our findings that default rates are the same on both sides of the RD cutoff and that predicted delinquency is similar on both sides of the cutoff continue to hold in this larger sample that includes unmatched loans.

However, there are two problems with using this larger sample. First, data quality is lower (the HAMP loans which match to TransUnion will tend to be the ones where loan characteristics are accurately measured). Second, this sample includes one large servicer that gave principal reduction to all borrowers (Scharlemann and Shore (2016) also drop this servicer). The public HAMP data do not include a servicer identifier, and so it is impossible to drop this servicer when working only with the public data. Our preferred specification in the paper uses the matched sample because it is able to address these issues.

B.1.4 Representativeness of HAMP Participants Relative to Typical Delinquent Underwater Borrowers

Our empirical analysis of the effect of principal reduction on default focuses on borrowers near the assignment cutoff for receiving principal reduction. To assess the representativeness of our analysis sample, we compare borrowers near the cutoff in the matched HAMP credit bureau file to a sample of delinquent borrowers in the Panel Study of Income Dynamics (PSID) between 2009 and 2011. Summary statistics for borrowers in both samples are shown in Table 1a. Borrowers in our sample are broadly representative of delinquent underwater borrowers during the recent crisis.

The median borrower in our sample has a higher LTV than delinquent borrowers in the PSID (121 compared to 94), but the 90th percentile LTV is similar (168 compared to 166). Since all the borrowers who are evaluated for principal reduction must be underwater, we would expect them to be concentrated in the underwater portion of the delinquent borrower distribution. The fact that borrowers in our 90th percentile are "only" at an LTV of 168, and that the median borrower is substantially less underwater, is important for interpreting our empirical results.

The PSID comparison is also helpful because it allows us to examine the liquid assets of borrowers. Delinquent borrowers in the PSID have very low levels of liquid assets. To be eligible for HAMP, borrowers had to attest that their liquid assets were less than three times their total monthly debt payments. However, the PSID data shows that this screen had little force. Even the delinquent borrower at the 90th percentile of the liquid asset distribution would have passed the HAMP screen.

B.2 Effect of Payment Reduction

B.2.1 Sample Construction and Robustness

We winsorize credit card spending at the 95th percentile of positive values.

Balance Plots – We show the trend in pre-determined covariates through the cutoff in Appendix Figure 11. The first four panels show the borrower monthly income, pre-modification monthly payment, LTV ratio, and months past due around the cutoff. These balance plots differ in two ways from the balance plots for the discontinuity for principal reduction. First, unlike in the matched HAMP credit bureau dataset used for the investor NPV strategy, borrower credit score is not available in the JPMCI bank dataset. Second, we

cannot show balance on PTI because it is the running variable. Instead, we show balance on pre-modification monthly payment.

There is no statistically significant jump in these loan and borrower characteristics at the 95 percent confidence level. In the bottom panel we use these observable borrower characteristics to predict default and show that predicted default is also smooth at the cutoff.

Density – Appendix Figure 12 shows that borrower density is also smooth around the cutoff.

Alternative Bandwidths – Our point estimate of $\hat{\tau}$ from equation (4) is that an extra 1 percent payment reduction reduces default rates in the two years after modification by 0.38 percentage points. Appendix Figure 15 tests the sensitivity of our results to the bandwidth chosen for the local linear regression. Our central estimates are constructed using a bandwidth of 0.06 points of PTI. We test alternative bandwidths between 0.01 and 0.1 and find that the point estimate is stable.

Adjusting for Upper Bound of Potential Principal Forgiveness Impact – If we take the most optimistic point in our 95 percent confidence interval for the impact of principal reduction on default from Section 3.3, and scale it by the amount of relative principal increase received by borrowers just below the 31 percent PTI cutoff, we find that a principal increase of this magnitude would have led to at most a 1.3 percentage point increase in default rates. If the payment reductions had to offset this effect, this would mean that the reduced form jump in default at the cutoff would have been 8.4 percentage points without the principal increase rather than 7.1 percentage points or, alternatively, that each 1 percent reduction in payment reduced default rates by 0.44 percentage points (1.4 percent), similar to our baseline estimate of 0.38 percentage points (1.2 percent).

B.2.2 Impact of Payment Reduction on Consumption

Our payment reduction regression discontinuity empirical strategy is under-powered for studying consumption impacts. In Appendix Figure 27, we plot the reduced form of the 31 percent PTI strategy with the change in mean credit card spending from the year before modification to the year after modification as the outcome variable. The standard error is so large that, using the same procedure for calculating an MPC as described in Section 3.4, we cannot rule out an MPC above 1 or below -1.

Unlike with principal reduction, we are unable to increase the precision of our payment reduction estimates by using a difference-in-differences design. The difference in principal reduction received by borrowers with and without principal reduction remains large when we expand the sample to a wider bandwidth. In contrast, the difference in payment reduction between HAMP and Chase modifications falls when looking at a wider sample (as can be seen by looking at the edges of Figure 5a). This is because the PTI target in HAMP generates larger payment reduction for higher PTI borrowers. Hence, comparing borrowers who received HAMP and Chase modifications at a wider bandwidth results in a shrinking size in the payment reduction treatment. We therefore conclude that our data and available research designs are unsuited for credibly estimating the effect of payment reduction on consumption.

C Net Present Value Calculations

In this section we provide more detail on our NPV calculations. The NPV calculations use borrower-level micro data on loan terms where available and otherwise use projections from the U.S. Treasury NPV model discussed in Appendix B.1.2.

Appendix C.1 discusses the basic setup, which is applicable to our analysis in Sections 3.1, 3.4, and 6.1. Appendix C.2 pertains to our discussion of the cost per avoided foreclosure from using principal reduction in Section 3.3.

Appendices C.3 and C.4 pertain to our analysis in Section 6.1. Appendix C.3 provides more detail for calculating the change in the NPV of expected payments to lenders at the HAMP eligibility cutoff. Appendix C.4 provides more detail for calculating the potential gains from redesigning HAMP modifications.

C.1 Net Present Value of Expected Payments

We use two equations to estimate the NPV of the loan. Equation (6) estimates the value of a mortgage that "cures," meaning that the borrower repays on time or early:

$$NPV\{Loan\ Pays\}(\delta) = \sum_{i=1}^{T} \frac{1}{(1+\delta)^i} \left[(UPB_{i-1} - Prin_i)(s_{i-1} - s_i) + (Prin_i + I_i)s_{i-1}\right] \quad (6)$$

where T is the term of the loan, δ is the investor's discount rate, UPB_i is the unpaid principal payment at time i, $Prin_i$ is the principal payment for period i, I_i is the interest payment for period i, and s_i is the survival probability of loan, which is constructed as $s_i \equiv \prod_{k=1}^{i}(1 - Prepay_k)$ where $Prepay_k$ is the prepayment probability in year k. The time period is annual. We observe UPB_i, $Prin_i$, and I_i for loans in the treatment and control groups in both the HAMP data and the JPMCI data.

We use the Treasury NPV model to estimate annual prepayment rates. This is the same model used by servicers to calculate the expected cash flows to lenders under various HAMP modification types, which we use for identification in Section 3.1, which is documented in U.S. Department of the Treasury (2015). The model uses a logit equation for predicting prepayment rates (Section V of U.S. Department of the Treasury 2015); we use the coefficients for owner-occupied homes reported in Appendix C of U.S. Department of the Treasury (2015) for borrowers that are 90+ days delinquent at modification date.[3]

Our second key equation incorporates default risk into our NPV estimate. We take equation (5) from Appendix B.1.2 and modify it to allow for the fact that not all defaults end up being liquidated:

$$\begin{aligned} NPV = (1 - p_{default}) * NPV\{Loan\ Pays\} + \\ p_{default} * [P(liquidate|D) * NPV\{Liquidate\} + (1 - P(liquidate|D))NPV\{Loan\ Pays\}] \end{aligned} \quad (7)$$

where $p_{default}$ indicates 90-day default. We follow the Treasury NPV model in making a simplifying assumption that borrowers make a one-time decision to default or not default.

To estimate the probability that a default results in a liquidation (with an accompanying loss for the investor), we use HAMP performance data. Among HAMP modifications that are disqualified due to default, 26 percent end up in foreclosure, 14 percent end in a short sale,

[3] Because we do not have access to all the covariates used in the Treasury NPV model, we need to separately estimate the intercept in the logit equation. We choose this intercept to match an annualized prepayment rate of 0.9 percent. This estimate is based on the prepayment rate of HAMP-modified mortgages in the first five years after modification.

18 percent self-cure, 33 percent get a proprietary modification, and 10 percent have delayed action, such as a borrower going through bankruptcy (U.S. Department of the Treasury 2017). Of loans whose status is fully resolved, 45 percent are foreclosed on, 24 percent end in a short sale, and 31 percent self-cure. We assume that loans which get a proprietary modification or delayed action ultimately have the same distribution of final outcomes. (Unfortunately, we do not have data on the outcomes of these proprietary modifications. Our assumption that these modifications have the same distribution of final outcomes is conservative in that it likely overstates the losses on these loans.) We explore alternative assumptions in the robustness analysis below.

Unfortunately, HAMP does not collect data on losses after disqualification so we draw on GSE performance data to estimate the NPV of loans that are "liquidated." The GSEs report losses on loans that are liquidated via either foreclosure or short sale. Goodman and Zhu (2015) document that GSE losses are quite similar on foreclosures and short sales. We use performance data from loans liquidated in 2011 because that was the year in which the GSEs experienced the largest number of liquidations. In that year, the Fannie Mae reported losses at liquidation equal to 41 percent of the unpaid balance on the loan (Fannie Mae 2018). However, this includes reimbursements from third parties (mortgage insurers and mortgage originators) to the GSEs equal to 15 percent of the unpaid balance of the loan. Altogether, investors lost 56 percent of the unpaid balance of the loan at liquidation. We explore alternative assumptions in the robustness analysis below.

C.2 Cost Of Preventing a Foreclosure

What is the cost to lenders and taxpayers of preventing a foreclosure via principal reduction? In this appendix, we answer this question using data from the policy discontinuity in Section 3 combined with the NPV model described above.

Our point estimates imply that principal reduction *raises* the probability of default by 1.2 percentage points. This estimate is small and statistically insignificant, with a standard error of 3.18 percentage points. However, we are unable to reject that principal reduction has *any* impact on lowering default rates. To assess whether principal reduction might be *cost-effective*, we consider the extreme case that principal reduction actually reduces default by 5.0 percentage points (1.2 - 1.96*3.18). This is the most optimistic number which is consistent with our 95 percent confidence interval.

We translate this default reduction to a foreclosure reduction using the estimates in Appendix C.1. Specifically, we assume that 45 percent of defaults end in completed foreclosures in our baseline analysis, so principal reduction reduces foreclosure completion by no more than 2.3 percentage points.

An alternative method to compute foreclosure reduction delivers a similar result. Appendix Figure 26 shows that we can rule out a reduction in foreclosure initiations of 2.9 percentage points. It is uncertain what fraction of initiations end up in foreclosures in our sample. Herkenhoff and Ohanian (2019) report that approximately half of borrowers with a foreclosure initiation end up with a completed foreclosure. Applying their estimate to our sample implies that principal reduction reduces foreclosure completion by no more than 1.5 percentage points during the window we study, even less of an impact than we consider in our baseline case.

To assess whether such a foreclosure reduction is cost-effective, we use the lender valuation model described in Appendix C.1. The modification terms correspond to those at the policy discontinuity arising from the Treasury NPV model. The borrower receives an

average of $31,000 in principal forgiveness at the discontinuity. We implement the model for a representative standard HAMP modification that includes only payment reduction and for a HAMP Principal Reduction Alternative modification, which includes both payment and principal reduction.

Relative to a standard payment reduction modification, a principal reduction modification has a benefit and a cost. The benefit is that the probability of default is lower (at least under our optimistic assumptions) and so puts more weight on the high return state of the world (where the loan pays). This benefit is higher when the default reduction is greater and when the difference between the high return state of the world and the low return state of the world is greater. The cost is that by forgiving principal it reduces the cash flows to investors in this higher return state of the world, (i.e., for a borrower that would have paid regardless, the lender is getting less cash flow with principal reduction than without it). This cost is lower when this higher return state of the world is unlikely without principal reduction.

Even under the most optimistic assumption about the effectiveness of principal reduction at reducing defaults, the lender valuation model implies that principal reduction is not a cost-effective way to prevent foreclosures. We calculate that lenders and taxpayers would incur a cost of at least $766,000 to prevent a foreclosure. The intuition for why this number is large is that the benefit is quite small (a 2.3 percentage point reduction in the foreclosure probability), while costs are quite large (reduced cashflows from the vast majority of borrowers who are not defaulting). Further, this estimate is a *lower bound* on the cost of preventing a foreclosure; at our actual point estimate, principal reduction does not prevent any foreclosures.

This cost can be borne by taxpayers, lenders or some combination of the two. Scharlemann and Shore (2016) find that the average government subsidy in the first cohorts of PRA was around $20,000, with an average amount of $74,000 of principal forgiveness. At this subsidy rate, we calculate that the $766,000 estimate in the previous paragraph can be decomposed into a cost per foreclosure prevented of $365,000 to taxpayers and $402,000 to investors.

C.3 Calculations for Pareto Improvement from Maturity Extension

Since the amount of liquidity provision appears to be more important than how that liquidity is provided, our results imply that maturity-extension-financed payment reduction may generate a Pareto improvement, leaving borrowers, lenders, and taxpayers all better off. To understand why this is true, it is useful to revisit the particular structure of mortgage modifications around the HAMP eligibility discontinuity we study in Section 5. A borrower who moves from the right-hand side of the cutoff ("control") to the left-hand side ("treatment") sees deeper immediate payment reductions that are offset by continued payments in the long-term. Intuitively, maturity extension is equivalent to the lender "lending" the borrower their monthly payment reductions at the mortgage interest rate, with repayments on this "new" loan beginning at the end of the original mortgage term and continuing through the end of the new loan term.

For borrowers, a maturity extension which moves funds from the future to the present might be particularly valuable in periods of acute economic distress. In our setting, all borrowers are better off in the treatment group under relatively mild assumptions about monotonicity and revealed preference. Borrowers can be divided into four potential groups: those who default under both contracts, those who switch from defaulting to paying on schedule, those who switch from paying on schedule to defaulting, and those who pay on

schedule under either contract. The first group defaults under both contracts and is therefore no worse off receiving treatment. The second group is better off because they indicate by revealed preference that the modified loan is more attractive than defaulting. If we assume that default is monotonically decreasing with the extent of payment reduction–the canonical assumption from Imbens and Angrist (1994) needed to identify a local average treatment effect–then there is no one in the third group. Finally, the fourth group is better off because their choice set is expanded; they always have the choice to ignore the maturity extension and repay the loan on the original schedule.

One question raised by the conclusion that borrowers are better off from the maturity extension we study is whether this is always the case. Are borrowers always better off from extended maturities, and if so, does this imply that the standard 30-year mortgage is too short? Unfortunately, our evidence does not inform this broader question. We only conclude that in the specific period we study, extending mortgage maturities for already-distressed borrowers appears to have benefited those borrowers. Because borrowers who are experiencing financial distress are likely to value additional liquidity, we conjecture that the benefit of maturity extension for borrowers is likely to extend to other periods of financial distress. However, our findings do not imply that longer maturities are optimal *ex-ante* (only that the option to extend in the face of financial distress may be optimal *ex-post*).

For lenders, a maturity extension will increase the NPV of total payments owed when the mortgage's interest rate is higher than the lender's current discount rate. In this time period the interest rates on new mortgages were a few percentage points below the average interest rate on existing mortgages. Hence, when using current interest rates as a measure of the opportunity cost of capital, extending maturities on existing mortgages would increase the NPV of payments owed to the lender. Although most of the incremental payment reduction at the cutoff we study was achieved from maturity extension, part came from interest rate reduction, which reduces the NPV of payments owed. This works against the maturity extension effect.

In our setting, we find that the NPV of the payments owed under the contract is similar in the treatment and control group. To be specific, Appendix Figure 13 shows that treatment lowers NPV by $2,168 relative to the control group when lenders use a 4.11 percent discount rate, which is the average of the 30-year mortgage rate during our sample period. This is a modest loss, and with a standard error of $2,049 we are unable to statistically reject that there is no change in the NPV to lenders (consistent with the criteria for a Pareto improvement). This estimate—which assumes that all borrowers repay on schedule—has three shortcomings: first, some borrowers default (and treatment reduces default); second, that some borrowers prepay their mortgage; and third, that lenders may discount cash flows after year 30 at a higher rate than short-term cash flows.

To address these limitations, we build an *expected payments* NPV model and find that under plausible assumptions lenders are better off assigning a borrower to treatment. We proceed in three steps. First, we incorporate default by using our causal estimate of the effect of treatment on default. We combine this with prior evidence on the losses incurred by lenders when borrowers default described in Appendix C. Because treatment significantly reduces default rates, this moves the lender to a gain of $9,053. Second, the fact that the payments arrive further in the future means that we need to use a higher discount rate, which decreases the NPV. We estimate a term premium of 32 basis points between 30- and 40-year mortgages by extrapolating from observed mortgage rates (see Appendix Figure 28), which shrinks the gain to the lender to $6,138. Finally, we incorporate realistic prepayment behavior for a final estimate of $6,229.

The finding that lenders are better off from maturity extensions depends on two crucial assumptions. First, we assume that prevailing mortgage interest rates accurately reflect a lender's opportunity cost of capital. But if some lenders are liquidity-constrained (as might be implied by the fall in mortgage originations during this time period), then market rates for those borrowers who can get a mortgage might not reflect a lender's true cost of capital.[4] We calculate that a lender is better off from treatment as long as her discount rate is below 6.13 percent.

Second, we assume that implied mortgage spreads accurately capture the lender's disutility from extending the mortgage term from 30 to 40 years. Because of uncertainty over this spread, we show that our results are robust to using a variety of term premium assumptions in Appendix C.3.1. However, the term premium we use in our baseline case (32 basis points) is already significantly higher than the actual 30- to 40-year spread for swaps and corporate bonds in our sample period (2 and 9 basis points, respectively). Furthermore, the flat swap yield curve during this time period implies that a lender concerned about increased portfolio duration risk from extending mortgage maturities could hedge this risk at low incremental cost.

Finally, taxpayers are also better off from more maturity extensions. The government spent substantial resources subsidizing HAMP modifications above the eligibility cutoff with small payment reductions and high default rates, whereas lenders were willing to provide borrowers below the cutoff private modifications requiring no government assistance which had large payment reductions and low default rates. This suggests that using maturity extensions as the first step in modifying mortgages could have saved substantial taxpayer subsidies.

Our findings for payment reduction—that payment reductions can be structured so as to reduce default rates while leaving all parties better off—contrast sharply with our findings for principal reduction, which was ineffective even while being costly to both lenders and taxpayers. Future private and public modification programs will have a menu of options for restructuring loans. Our findings suggest that among these options, those that maximize immediate payment reduction are likely to be most effective, and that maturity extension is a particular way to achieve large immediate payment reductions at little cost to lenders and taxpayers. We explore the broader implications of these results for modification design in more detail in Appendix C.4.

We first provide more detail on the expected payments NPV calculation as well as describe several robustness checks below.

C.3.1 Robustness of Expected Payments NPV at 31 Percent PTI Discontinuity

Our choice of the discount rate δ for future cash flows depends on the maturity of the mortgage. Recall that assignment to treatment involves an extension in the term of the mortgage and 80 percent of loans in the treatment group last 40 years after modification. Ideally, we would use the interest rate on 40-year mortgages to discount these cash flows, but unfortunately we are unaware of any publicly available data source with prices for 40-year mortgages. Instead, we estimate the price of a 40-year mortgage by using a simple functional form to extrapolate from the price of 15-year and 30-year fixed mortgages sold by Freddie Mac. The JPMCI payment reduction sample includes modifications from October 2011 through January 2014. The average 15-year rate is 3.06 percent during this period and the average 30-year rate is 3.84 percent.[5] We fit an equation $r = \alpha + \beta log(term)$ to these

[4] Appendix Figure 34a shows the time series of mortgage originations during this period.

[5] This is quite similar to the average 30-year rate of 4.11 percent during the time period when modifications

data and estimate a hypothetical 40-year mortgage rate of 4.16 percent.[6] In the robustness analysis below, we explore the sensitivity of our estimates to alternative assumptions about the discount rate and the yield curve.

We estimate the effect of treatment on default rates using our causal estimates from the regression discontinuity design and HAMP performance data. We estimate a 90-day default rate in the two years after modification of 24.8 percent for the treatment group and 32.1 percent for the control group. Among HAMP modifications done in 2010, the default rate is 28.1 percent two years after modification and 45.6 percent five years after modification (U.S. Department of the Treasury 2017), for a ratio of 1.62. We project default rates five years after modification in our data by multiplying our estimated default rates by 1.62. This calculation assumes that payment reduction is equally effective in years three, four, and five. We project the default rate will be 38.6 percent in the treatment group and 46.2 percent in the control group. We explore alternative assumptions for the impact of treatment on the default rate in the robustness analysis below.

Our estimates imply a gain to the investor from assigning a borrower to treatment. Recall that treatment is essentially a loan to the borrower in the form of lower mortgage payments for 22 years which is offset by additional mortgage payments extending beyond the pre-modification term of the loan. The change in the NPV arising from this maturity extension treatment is $6,229, as shown in Appendix Table 4. This is equal to a 3.8 percent increase in the NPV of the loan.

As an alternative to the NPV calculation, we also report the discount rate an investor would need to break even on providing treatment to a group of mortgages. While the prior calculation assumed that the lender discounted future mortgage cash flows at our best estimate of the market interest rate, an alternative approach allows us to be agnostic as to the lender's discount rate. The NPV of a mortgage that cures is a function of the discount rate δ, as shown in equation (6) and the expected NPV of all mortgages in equation (7) relies on this, so we can rewrite NPV in equation (7) as a function $NPV(\delta)$. Then, we can solve for the discount rate that satisfies the lender's indifference condition such that the change in NPV from offering the treatment modification is the same as the change from offering the control modification:

$$\delta^* \text{ such that } \Delta NPV(\delta|T) = \Delta NPV(\delta|C).$$

In our baseline specification, we estimate that a lender that discounts the future annually by 6.13 percent will be indifferent between offering this modification. This implies that a lender with an annual discount rate less than 6.13 percent will be better off offering the treatment modification.

Did lenders in fact discount future cashflows at less than 6.13 percent? It appears that most did. As part of the HAMP NPV test, mortgage servicers chose a discount rate for future cashflows. Their choice set ranged from the market interest rate as a lower bound to 250 basis points above the market interest rate as an upper bound. SIGTARP (2012) report that 96 percent of servicers discounted future cashflows at the market interest rate, which

were performed for our principal reduction sample.

[6]When a mortgage term lasts less than 35 years, we use the 30-year rate and when a mortgage term lasts 35 years or more, we use the 40-year rate. Our results would change very little if we instead used different discount rates for every possible mortgage maturity between 30 and 40 years. In the analysis sample, 51 percent of mortgages last exactly 40 years after modification and 40 percent last 30 years or less after modification.

was around 4 percent during this time period. Any servicer that discounted cashflows at this rate would have accepted maturity extension, which was NPV positive for any discount rate less than 6.02 percent.

We explore the robustness of our NPV and discount rate estimates to alternative assumptions on default rates, recovery rates on losses, discounting, and prepayment in Appendix Table 4. Across almost all scenarios, we find that the NPV of the loan to the investor increases from assigning a loan to treatment instead of control. First, we explore the impact of alternative assumptions about the impact of treatment on mortgage default. Using the lower and upper bounds of our 95 percent confidence interval, we estimate the change in NPV ranges from $2,597 to $9,861.

Second, we show that impact of treatment on NPV is sensitive to our assumptions on the recovery rate on defaulted loans, but is always positive or statistically indistinguishable from zero. Our specification with the most optimistic recovery rates assumes that *every* proprietary modification and *every* action pending self cures, meaning that there is a 61 percent self cure rate, and uses the highest possible recovery rate on GSE loans during the crisis, which was a 48 percent loss in 2009. Our specification with the most pessimistic recovery rates assumes that all proprietary modifications and action pending ends in liquidation, meaning that there is an 18 percent self-cure rate, and the lowest possible recovery rate on GSE loans, which was a 61 percent loss in 2014. Treatment in the optimistic scenario causes an NPV loss to the investor of -$1,028, while in the pessimistic scenario it causes an increase of $9,106. Note that -$1,028 is indistinguishable from zero given our standard errors, and therefore the criteria for a Pareto improvement (which is that at least one party is better off and no party is worse off) is still satisfied in this scenario.

Third, we show the impact of using alternative methodologies for estimating the discount rate. Intuitively, the treatment modification defers cash flows from the present to the future and investors require a higher rate of return for deferring these cash flows. Recall that the average interest rate for a 30-year fixed rate mortgage during our sample period is 3.84 percent and in our baseline specification we estimated an additional 32 basis points for a 40-year mortgage. At one extreme, an alternative methodology which relies on a comparison of 30-year and 40-year loans is the swap rate where the yield curve is flatter and the average spread in our sample period is only 2 basis points. At the other extreme, projecting hypothetical spreads using interest rates on debt issued by the U.S. Treasury implies a steeper yield curve with an additional 34 basis points for a 40-year mortgage. Both of these projections are shown in Appendix Figure 28. This flatter yield curve implies a change in NPV of $10,137 and the steeper yield curve implies a change in NPV of $5,917. The figure also shows that, if anything, the log functional form overestimates the term premium at higher maturities. Forty-year maturities are actually observed for swaps and corporate bonds. For these, the "implied" spread between 30 and 40 year maturities using the log functional form assumption are much larger than the actual spreads.

It may be preferable to use the risk-free rate to discount cashflows in our model. The argument for using the risk-free rate here is that lenders offering mortgages charge a premium over the risk-free rate in order to compensate the lender for prepayment risk and default risk. However, our expected payments NPV calculation already takes into account default and prepayment risk. The average rate on 30-year Treasury notes during this time period is 3.17 percent, and we project that that the rate on a 40-year note would be 3.52 percent. The average rate on fixed-for-floating swaps is 3.00 percent for 30 years and 3.02 percent for 40 years. Under these assumptions, we calculate changes in NPV of $9,838 and $15,845 respectively. The value is greater to the investor under this scenario because a maturity

extension delays cashflows, and switching to a lower discount rate makes cashflows far in the future more valuable.

Fourth, we show that prepayment rates have little effect on the change in NPV from treatment. At one extreme, we assume an annual prepayment rate of 0.9 percent (the observed prepayment rate after HAMP modification) for the life of the loan. At the other, we assume an annual prepayment rate 6.8 percent (the observed prepayment rate on all Fannie Mae loans 1999-2017). The change in NPV varies from $6,064 under the low prepayment scenario to $7,014 under the high prepayment scenario..

Finally, we crosswalk our expected payments NPV estimate to the payments owed NPV estimates reported elsewhere in the text. Recall that the investor's return from treatment is a $6,229 gain in terms of expected payments NPV, but a loss of $2,168 when using the payments owed NPV estimate reported in Appendix Figure 13. (To be precise, the figure shows that the investor loses $2,168 more from treatment). This assumes that the loan is repaid on schedule (no default or prepayment) and the investor discounts cashflows at 4.11 percent annually.

C.3.2 Does More Time Underwater Offset the Default-reducing Benefits of Payment Reduction?

Our empirical results show a 23 percent reduction in default rates (from 32.1 percent to 24.8 percent) at the cutoff. One possible downside of maturity extension is that it leaves borrowers underwater for longer. The default-reducing benefits of payment reduction might be offset by a longer period of being underwater during which potential shocks (such as health or job loss) could potentially push a borrower into foreclosure.

The magnitude of this offsetting effect depends on how much longer a borrower who receives maturity extension remains underwater. To quantify this, we analyze two hypothetical mortgages which match the average characteristics of borrowers on each side of the 31 percent PTI discontinuity. These borrowers have an average loan-to-value ratio of 131 percent. The right-hand side borrower receives a 70 basis point reduction in the interest rate and $19,500 in principal forgiveness, which enables a payment reduction of 14 percent (equivalent to payment reduction on the right-hand side) for the remaining 23 years on the loan. The left-hand side borrower receives a payment reduction of 31 percent (equivalent to payment reduction on the left-hand side). This borrower receives an extension of the maturity of the loan to 40 years, a 110 basis point reduction in the interest rate, and $14,000 in principal forgiveness. These parameterization choices are designed to approximate the modification terms at the cutoff; we could instead obtain the same amount of payment reduction on each side if we gave each borrower a 160 point basis point reduction in the interest rate and also gave maturity extension to 40 years to the borrower on the left-hand side. We assume that nominal house prices grow at 5.8 percent per year, which is the average annual growth rate of the FHFA housing price index from 2011-2016.

We find that the typical borrower on the left-hand side of the cutoff spends an additional year underwater. Appendix Figure 29a plots the projected value of the home and the loan's unpaid balance under each type of modification. It shows that a typical borrower on the right-hand side reaches the above water mark in the third year after modification. On the left-hand side, the borrower reaches the above water in the fourth year.[7] Defaults after these

[7]The difference is even smaller in many high-LTV geographies who had bigger price declines during the crisis followed by sharper rebounds. For example, Las Vegas had the biggest price drop of any MSA between 2006 and 2010, and it experienced average annual nominal house price growth of 11 percent between 2011 and 2016. A borrower in Las Vegas on the left-hand side of the cutoff would have only spent an additional five months underwater compared to a borrower on the right-hand side of the cutoff.

points in time are very unlikely to result in foreclosure because the borrower would prefer to sell the house and get back her home equity.

The additional year underwater raise the risk of foreclosure modestly, but by less than the reduction in defaults generated by the additional payment reduction. To project default rates beyond the two-year horizon that we observe in the JPMCI data, we use public tabulations of the performance of HAMP loans. These tabulations include default rates of HAMP recipients from three months to five years after modification. We fit a regression model where $default_t = \beta \log(t)$ and show the projections in Appendix Figure 29b. We project default rates for the left-hand and right-hand sides using $default_t^{\{LHS,RHS\}} = default_{24 \text{ months}}^{\{LHS,RHS\}} \frac{default_t^{perf}}{default_{24 \text{ months}}^{perf}}$ and fit the regression model from above separately for each group. We project a three-year default rate for the right-hand side group of 41.2 percent and a four-year default rate for the left-hand side group of 35.8 percent. Defaults with significant foreclosure risk are 5.4 percentage points (about 13 percent) lower for borrowers on the left-hand side. Thus, maturity extension appears to benefit borrowers even after incorporating how the additional time underwater exacerbates foreclosure risk.

These calculations change little when we incorporate prepayment. The intuition for why prepayment is unimportant is that the prepayment rate on recently modified loans is quite low (less than 1 percent annually).

C.4 Efficient Default-Minimizing Modification Design

In this section, we explain how our results help to answer the question of how to efficiently design mortgage modifications for all borrowers, and we quantify the potential gains of implementing our proposed design relative to the mortgage modifications that were actually pursued. Our analysis focuses on applying the empirical lessons from the episode we study towards designing a uniform mortgage modification for all borrowers. We do not consider additional gains that could potentially be achieved by targeting specific modification offers to borrowers with particular characteristics.

Lessons for Modification Design – Our empirical findings that default is responsive to liquidity but not wealth suggest a simple principle for reducing mortgage defaults: the "best" modification steps are those that achieve immediate payment reduction at the lowest possible cost. The costs of payment reductions must be borne by either lenders or taxpayers. Minimizing costs per dollar of immediate liquidity provision will maximize the amount of payment reduction (and hence default reduction) that the market will find privately optimal to provide on its own and identify the most efficient use of government subsidies. Hence, we call modifications that follow this minimum-cost structure "efficient default-minimizing modifications."

To uncover the efficient default-minimizing structure of mortgage modifications we evaluate the five modification steps that were used in various combinations in the public and private modification programs we are aware of: maturity extension, temporary five-year interest rate reduction, principal forbearance, permanent interest rate reduction, and principal forgiveness. These policies all reduce mortgage payments for at least five years but have very different costs. We calculate these costs for an illustrative mortgage with the characteristics of the average loan at the HAMP eligibility discontinuity.[8]

[8]This loan has a 6.7 percent fixed interest rate, a 23-year remaining term, an unpaid balance of $248,000 and a loan-to-value ratio of 131 percent.

We rank modification steps by their cost-effectiveness in Appendix Figure 20a. We calculate the change in mortgage terms needed to reduce payments by 10 percent and the change in the NPV of payments owed to the lender from this modification step. We find that maturity extension is NPV-*positive* by nearly $20,000 for the lender because the interest rate on the loan is higher than the lender's discount rate (recall that interest rates on new mortgages fell substantially during the crisis, so the spread between the old rate on the mortgage being modified and the available return on new mortgages widened). Temporary interest rate reductions are NPV-negative, costing about $14,000, while principal forbearance, permanent rate reduction, and principal forgiveness are even more NPV-negative, with costs between $22,000 and $32,000. Appendix Figure 20a also shows that the same ranking of policies continues to hold when we examine the change in the NPV of expected payments incorporating the yield curve and the impact of modifications on default and prepayment risk.[9] Across all modification steps we find that the costs to the lender are lower when we incorporate these features and that the ranking from least cost to highest cost is the same.[10] Finally, Appendix Figure 25 shows that this ranking is robust to alternative assumptions about recovery rates and interest rates.

We use this ranking to propose efficient default-minimizing modifications for a range of payment reduction targets. As we note above, by efficient we mean that we use the lowest-cost policies first. We assume there are limits to how much some mortgage terms can plausibly change. For example, maturity extension cannot possibly reduce payments below the interest payments on the unpaid balance. We restrict the set of possible modifications using the quantitative limits implemented by the HAMP, Chase, and GSE modification programs: we allow for maturity extension up to 40 years, we allow temporary interest rates to be reduced to 2 percent, and we allow permanent interest rates to be reduced to the prevailing 30-year mortgage rate. Adopting these limits makes our characterization of the potential gains from more efficiently-designed mortgage modifications conservative.

The efficient default-minimizing modification uses maturity extension, followed by temporary interest rate reduction, followed by principal forbearance. The efficient modification never uses principal forgiveness. In Appendix Figure 20b, we depict the cost of such a modification. Maturity extension can reduce payments by up to 18 percent if the loan term is extended from 26 to 40 years, payments can be reduced by an additional 41 percent by reducing the temporary interest rate to 2 percent, and forbearance can reduce payments even further.[11] The figure also shows the cost of principal forgiveness. Together these two lines

[9] We use the same assumptions to calculate expected payments NPV calculation as in Appendix C.3. We need one new assumption, which is a function that maps payment reduction to default rates for lower levels of payment reduction. We extrapolate from our evidence on default rates and Appendix Figure 30 shows the function.

[10] One limitation of this ranking is that it assumes that payment reduction over a five-year horizon is the main driver of a policy's effect on a loan's lifetime default rate. Implicitly, we are assuming that payments in the next five years are the relevant variable for the re-default probability of a borrower who is currently in financial distress. This assumption could be wrong in either direction. On one hand, perhaps reducing payments for one year or three years can deliver much of the same reduction in defaults with even less cost. On the other hand, default could rise substantially when payments rise in year six, though empirical evidence in Scharlemann and Shore 2019 shows that this effect is quantitatively small. A payment increase of about 10 percent in year six raises the default rate in the following year from 3.8 percent to 4.6 percent. Unfortunately, our empirical evidence does not speak to this question, and we think that understanding the optimal duration of payment reductions is an important area for future research.

[11] Our estimates imply that a lender would break-even by offering a payment reduction of 65 percent through a maturity extension to 40 years, a temporary interest rate reduction down to 2 percent, and principal forbearance of 18 percent of the loan balance. Appendix Figure 31a shows that when we incorporate the

show the envelope of the least and most costly ways to achieve various amounts of payment reduction and hence various amounts of default reduction.

The modification policies actually used during the crisis occupy an intermediate position within the envelope of the most efficient and least efficient policies. Recall that HAMP targets a PTI ratio of 31 percent. This means that borrowers receive widely-varying amounts of payment reduction on the basis of their initial PTI ratio. As discussed in Section 3.1, HAMP first reduces the permanent interest rate, then the temporary interest rate, then extends the mortgage term, and finally does principal forbearance. Appendix Figure 20b shows that by first reducing the permanent rate, HAMP has a marginal cost for small payment reductions that is very close to the cost of principal forgiveness and much larger than the efficient default-minimizing modification, which uses maturity extension first. HAMP Principal Reduction Alternative (PRA), which offered principal reduction as the first step in a mortgage modification (see Section 3.1 for details) has an even higher cost from payment reductions. The figure shows that HAMP PRA follows the upper envelope of the least efficient policy until a payment reduction of 11 percent is reached. After that, the path of the HAMP PRA line parallels the standard HAMP modification.

In contrast to HAMP modifications, the GSE and private modifications we analyze are much closer to the efficient frontier. Recall from Section 5.1 that the GSEs and Chase offered most borrowers maturity extension, followed by permanent rate reduction, temporary rate reduction (Chase only), and then principal forbearance.[12] Because the GSEs and Chase had a specific payment reduction target, we depict them in Appendix Figure 20b using dots rather than the lines that we used to depict HAMP.

Potential Gains from More Efficient Modification Design – Our results imply that modifications in a future crisis can be redesigned to make borrowers, lenders, and taxpayers better off. To gauge the magnitude of the potential gains, we evaluate our proposed modification structure in comparison to HAMP, where the median borrower received a 38 percent payment reduction. We consider two ways to quantify the potential gains of more efficient modifications.

First, using our "efficient default-minimizing modification" structure we find that the same median payment reduction could have been provided at $67,000 lower cost per modification to lenders and taxpayers. This is equal to 27 percent of the unpaid balance of a typical loan.[13] Aggregating over all 1.8 million HAMP modifications, this implies a potential unnecessary cost of $121 billion. Taxpayers spent around $27 billion subsidizing HAMP modifications (Government Accountability Office 2016). Therefore, our results imply that HAMP could have been designed with no taxpayer subsidies and a much lower cost to investors, while maintaining the same amount of payment reduction.

Alternatively, if we allocate all the gains from redesigning modifications to reducing borrower payments, the same amount of lender and taxpayer cost can be used to achieve substantially more default reduction. Again considering the median HAMP borrower, we find that it was possible to reduce payments by 72 percent (rather than the 38 percent that actually occurred) at the same cost. Extrapolating from our empirical results using

social cost of foreclosure of $51,000 from U.S. Department of Housing and Urban Development (2010), the break-even point for lenders and taxpayers is a payment reduction of 70 percent. Appendix Figure 31b shows that the same broad patterns hold for payments owed NPV as for expected payments NPV.

[12] Chase also offered principal forgiveness to some borrowers with high LTVs.

[13] Although our estimate of the potential value from well-designed modifications is large, it is similar to prior work by Maturana (2017) showing that private modifications raised the NPV of a loan by 16 percent of the unpaid balance.

the function mapping payment reduction to default rates shown in Appendix Figure 30, we find that this quantity of payment reduction would have cut default rates by one-third. Aggregating over all HAMP modifications, this implies that 267,000 defaults could have been avoided at no additional cost to lenders or taxpayers. Looking forward, one caveat to achieving these gains in a future modification program is the extent to which the diminishing returns to payment reduction may be more pronounced in alternative (perhaps less severe) economic environments.

Our proposed approach has benefits beyond improving outcomes for borrowers in a government program. For example, the potential default reduction would be even larger if we considered redesigning *all* the 10 million public and private modifications completed in the Great Recession. Furthermore, when incorporating the $51,000 social cost of foreclosures estimated in U.S. Department of Housing and Urban Development (2010), the default reduction just from redesigning HAMP would have generated $6 billion in social value. Taken together, these results suggest that the gains in a future modification program could be divided such that borrowers, lenders and taxpayers all benefit relative to what was implemented in the Great Recession

Another advantage of the efficient default-minimizing modification is that it is likely to generate less ex-ante moral hazard. Some borrowers who did not have a liquidity problem may have defaulted on their mortgages in order to become eligible for the generous subsidies. Mayer et al. (2014) document increased defaults after Countrywide announced a generous modification program. The Home Affordable Refinance Program was specifically created to allow underwater borrowers to refinance without needing to become delinquent to get HAMP. Compared to HAMP, a modification with little change in the NPV of payments owed may not be attractive to borrowers who are current on their mortgages, and, even if it is, the resulting modifications will not be costly to lenders or taxpayers.

C.4.1 Ex-ante Mortgage Design

Our empirical results about ex-post debt restructuring also help inform the theoretical debate about optimal ex-ante mortgage design. A number of recent papers have analyzed alternative mortgage contracts with built-in features designed to assist households overcome periods of financial distress. Eberly and Krishnamurthy (2014) propose a fixed rate mortgage (FRM) with a one-time option to convert to an adjustable-rate mortgage (ARM), and Guren, Krishnamurthy and McQuade (2018) develop an equilibrium model of the housing market to evaluate this contract. They find that the option to convert to an ARM is more effective than alternative contracts because it front-loads the payment reduction to the borrower, and does so at similar cost to the lender. Our results provide empirical evidence that contracts front-loading payment reductions to households in financial distress will be more effective at preventing defaults than an alternative contract with equal cost to the lender that spreads payment reductions throughout the mortgage term.

Similarly, we also provide empirical support for proposals specifically contemplating the option to extend mortgage maturities. Campbell, Clara and Cocco (2018) compare an ARM with a refinance option to an ARM with the option to temporarily allow for interest-only payments with a corresponding increase in its maturity. They find that this maturity extension option outperforms a refinance option because it provides similar liquidity-provision benefits to borrowers at much lower cost to lenders. Our results show that immediate payment reduction can indeed be effective at reducing defaults even if structured with offsetting payments in the future so as to minimize costs to lenders.

D Partial Equilibrium Life-cycle Model with Housing

We argue that the inability of underwater borrowers to monetize the wealth gains from principal reduction can explain why they are far less sensitive to housing wealth changes than borrowers in other economic conditions. In the changes examined in prior research, housing wealth gains expanded borrowers' credit access. Mian and Sufi (2014) show that equity withdrawal through increased borrowing can account for the entire effect of housing wealth on spending between 2002 and 2006. But if homeowners face a collateral constraint rather than a "natural" borrowing limit, allowing them to monetize the present value of their minimum expected lifetime net worth, principal reductions that still leave borrowers underwater will not immediately relax this constraint.[14] Indeed, Defusco (2017) shows that a significant fraction of the additional borrowing arising from house price gains is due to relaxing collateral constraints. If borrowers cannot immediately monetize the wealth gained by debt forgiveness, it may not be surprising that they do not respond by increasing consumption.[15]

On the other hand, even if borrowing constraints are not relaxed immediately, it is possible that forward-looking agents building up a buffer of assets could respond if they believed principal forgiveness would relax their constraints in the near future. We calculate that since borrowers remained underwater and collateral constraints had tightened, it would take eight years before the average principal reduction recipient in HAMP would expect to be able to increase borrowing as a result of these principal reductions. A dynamic incomplete markets model of household optimization is useful for understanding whether such a lengthy delay can indeed explain why borrowers did not increase consumption.

We describe such a model in the remainder of this appendix. Section D.1 describes the model setup, Section D.2 describes the model's parameterization, Section D.3 describes the model's predictions for consumption (including an extended discussion of Figure 3b in Appendix D.3.5), and Section D.4 describes the model's predictions for default.

D.1 Setup

We consider a partial equilibrium life-cycle model of household consumption and default decisions. Households live for a maximum of T periods. The first $T_y - 1$ periods correspond to working age, the subsequent periods to retirement.

Households maximize expected utility, have time-separable preferences, and discount utility at rate β. Per-period utility is

$$U\left(c_{it}, d_i\right) = \frac{c_{it}^{1-\gamma}}{1-\gamma} - d_i \mathbf{1}(t=0)\psi$$

where c_t is non-housing consumption, d_i is an indicator variable equal to 1 if the household defaults, and ψ is a utility cost of defaulting. This additive default cost follows the structure in Campbell and Cocco (2015), Hembre (2018), Kaplan, Mitman and Violante (2017), and Schelkle (2018). It reflects the moral and social stigma associated with defaulting on debt

[14] Beraja et al. (2019) document a related channel: underwater borrowers are usually not able to refinance their mortgages. See Carroll (1992) and Aiyagari (1994) for discussions of natural borrowing limits.

[15] The possibility that liquidity can explain the lack of response is also consistent with prior research looking at large price declines. Mian, Rao and Sufi (2013) and Kaplan, Mitman and Violante (2016) both document a non-linearity in the consumption response to house price declines, with a large MPC for small declines but a decreasing MPC for large declines. Mian, Rao and Sufi (2013) suggest that the non-linearity they document could be caused by smaller responses once borrowers become underwater. Similarly, our evidence suggests that for borrowers who start substantially underwater, gains in housing wealth do not affect their consumption.

obligations as well as moving costs. We discuss the timing of default at the end of this section.

Agents consume a fixed quantity of housing. We assume housing and non-housing consumption are separable and, since quantity is fixed, follow Campbell and Cocco (2015) who show that under these conditions it is unnecessary to include housing explicitly in household preferences.[16] In the first period, agents are endowed with a home with market price P_{i1} and a 30-year fixed rate mortgage with balance M_{i1} and interest rate r. We assume home prices evolve deterministically according to $\Delta \log P = g$, where g is a constant, though we solve the model under various home price growth expectations. As long as households stay in this home, their housing costs include their mortgage payments (given by the standard annuity formula), property taxes τ_p that are proportional to the current market value of their home, and maintenance costs τ_m that are proportional to the initial value of their home.[17] Renters pay the user cost of housing for the equivalent home. Thus, housing payments are given by

$$h_{itj} = \begin{cases} M_{i1} \frac{r(1+r)^{30}}{(1+r)^{30}-1} + \tau_p P_{it} + \tau_m P_{i1}, & j = owner \\ (r - g + \tau_p) P_{it} + \tau_m P_{i1}, & j = renter \end{cases}. \tag{8}$$

If they have not defaulted, households sell their home at retirement (i.e. at $t = T_y$), enter the rental market, and use the proceeds of the home sale to supplement their income for the remainder of their life.

Households can only borrow out of positive home equity, subject to a collateral constraint. Thus, their liquid assets a_t can never fall below their borrowing limit $\underline{a}_t$ given by

$$a_{it} \geq \underline{a}_{it} = \min \left\{ - \left[(1 - d_i)(1 - \phi) P_{it} - M_{it} \right], 0 \right\},$$

where $(1 - \phi)$ is the fraction of a house's value that can be used as collateral.[18] Renters are not able to borrow.

Households face an exogenous income process. During working age, labor income is given by

$$z_{it} = \Gamma_t \theta_{it},$$

where Γ_t reflects deterministic life-cycle growth and θ_{it} is an i.i.d transitory shock with $\mathbb{E}[\theta_{it}] = 1$. During retirement, income is given by a constant social security transfer which is captured in the Γ_t process. Total income, including income from home sales in the first

[16] Campbell and Cocco (2015) show that these preferences are consistent with preferences over housing and non-housing consumption given by $\frac{c_{it}^{1-\gamma}}{1-\gamma} + \lambda_i \frac{H_{it}^{1-\gamma}}{1-\gamma}$ for $H_{it} = H_i$ fixed and where the parameter λ_i measures the importance of housing relative to non-housing consumption.

[17] The assumption that maintenance costs are proportional to initial values ensures that maintaining the same home does not become more expensive simply because market home prices rise.

[18] In the main parameterizations of our model house price growth is positive, such that once borrowers attain positive equity they do not risk falling back underwater. With negative home price growth, the borrowing limit is given by

$$a_{it} \geq \underline{a}_{it} = \min \left\{ \min \left\{ - \left[(1 - d_i)(1 - \phi) P_{it} - M_{it} \right], 0 \right\}, a_{it-1} \right\}$$

in order to prevent forced deleveraging of liquid assets.

period of retirement, is[19]

$$y_{it} = \begin{cases} \Gamma_t \theta_{it} & t < T_y \\ \Gamma_t + (1 - d_i)(P_{it} - M_{it}) & t = T_y \\ \Gamma_t & t > T_y \end{cases}.$$

Households can invest in a liquid asset earning a rate of return r. End of period assets evolve according to

$$a_{it} = (1 + r)a_{i,t-1} + y_{it} - c_{it} - h_{itj}.$$

We will often discuss our results in terms of cash-on-hand $m_{it} = (1 + r)a_{i,t-1} + y_{it}$.

We model default as a one-shot decision. Households begin the first period with a given mortgage, home price, and asset level. They then observe their first-period income shock, and decide whether to default or hold the house until retirement. This provides a simple way to analyze the short-term default decisions which we study empirically in Section 3.3. In Section D.4 we study how changing the initial conditions by modifying a borrower's mortgage affects their default decision in the model and compare this to our empirical results.

We solve the household problem recursively using the method of endogenous gridpoints suggested in Carroll (2006). This generates optimal consumption paths and the initial default decision.

D.2 Parameterization

The main parameter values are summarized in Appendix Table 5. We assume that each period corresponds to one year. In our baseline case we assume households start life at age 45 and live with probability 1 until retirement at age 65. Survival probability shrinks every year during retirement, and households are dead with certainty by age 91, as assumed by Cagetti (2003). We solve the model for different first-period ages from 35 to 55 to examine the effect of principal reduction at different ages.

We follow Carroll (2012) who assumes income shocks have a lognormal component as well as an additional chance of a large negative shock. The large negative shock, which we call unemployment, captures the idea that the income process has a thick left tail (Guvenen, Ozkan and Song 2014). Formally, income shock θ is distributed as follows:

$$\theta_{it} = \begin{cases} b & \text{with probability } p \\ \frac{\delta_{it}(1 - b \cdot p)}{1 - p} & \text{with probability } (1 - p) \end{cases} \tag{9}$$

where $\log \delta_{it} \sim \mathcal{N}\left(-\frac{\sigma_\delta^2}{2}, \sigma_\delta^2\right)$, p is the probability of unemployment, and b is the unemployment replacement rate. This ensures that $E[\theta_{it}] = 1$. All income risk, including unemployment, is turned off in retirement. We follow Carroll (1992) and set $\sigma_\delta = 0.14$.[20] We use data from Guvenen, Ozkan and Song (2014) to parameterize b and p. They show that the tenth percentile shock between 2008 and 2010 was a reduction in income of 50 percent, so we set p to 0.1 and b to 0.5. This large negative shock is critical to understanding default dynamics, which we explore in more detail in Section D.4. The life-cycle growth path of permanent income Γ_t is from Carroll (1997).

[19] In all of our parameterizations borrowers have positive equity by retirement.

[20] Carroll (1992) allows for temporary and permanent income shocks, each with a standard deviation of 0.1. We only have one income shock, whose standard deviation we set to $\sqrt{\sqrt{0.10} + \sqrt{0.10}} = 0.14$.

All parameters in our model are real, so we set the interest rate r to 2 percent. This matches the average 30-year mortgage rate from the Freddie Mac Conforming Loan Survey for the period 2010-2014 (4.1 percent) minus the average expected inflation on 30-year Treasury bonds over the same period (2.1 percent). We assume a collateral constraint ϕ of 0.2, such that homeowners can only borrow up to 80 percent of the value of their home. This matches the caps for cash-out refinancing from Fannie Mae and Freddie Mac, as well as evidence from Corelogic (2016) that average CLTVs on new HELOC originations fell 20 points from their peak in 2004 when CLTVs of 100 were possible. In our baseline model we set real annual house price growth g at 0.9 percent, which is the average from FHFA's national index between 1991 and 2010, as well as the expected annual price growth from home price futures in 2011, though we test the sensitivity of our results to alternative house price growth rate paths. We follow Himmelberg, Mayer and Sinai (2005) and set the property tax rate to 1.5 percent and the maintenance cost to 2.5 percent. These parameters generate a first-period user cost of housing of 5.1 percent, similar to the empirical estimates in Diaz and Luengo-Prado (2008) and Poterba and Sinai (2008), who find 5.3 percent and 6 percent, respectively.

We choose baseline preference values of $\beta = 0.96$ for the discount factor and $\gamma = 4$ for the coefficient of relative risk aversion. Our choice of a relatively high value for γ is not important for our consumption results, but is necessary in order to generate optimizing double-trigger behavior.[21]

We estimate our final parameter ψ, the utility cost of default, such that the first-period default rate in the model matches the 10 percent first-year default rate for moderately underwater borrowers in our data. Since our empirical default results focus on borrowers below 150 LTV, we allow default to rise above 10 percent for more underwater borrowers. We estimate ψ to equal 5.4 utils. To translate this into meaningful units, we calculate that this is equivalent to a 10 percent permanent income loss. This loss is in line with other estimates in the literature that uses structural models with default costs to match observed default rates. Schelkle (2018) builds a model to match the rise in default rates in the U.S. between 2002 and 2010 and estimates a default cost equal to 8 percent of permanent income. Kaplan, Mitman and Violante (2017) calibrate a default cost to match the foreclosure rate in the late 1990s and find a cost which is equal to 4 percent of permanent income for the median household, and approximately 7 percent for mortgagors. Hembre (2018) studies default behavior for all HAMP modifications and finds that a cost equal to 70 percent of per-period consumption is necessary to explain observed default rates.

D.3 Consumption

D.3.1 Consumption Response to Principal Reduction in Model

A reduction in mortgage debt levels affects today's consumption through two channels. The first is a future cash-on-hand effect. Reducing mortgage debt reduces a borrower's housing payments over time and increases a homeowner's expected home equity gain when they sell the house. These translate into consumption according to the homeowner's MPC out of cash-on-hand gains at future dates. The second channel is a collateral effect. Reducing debt levels frees up home equity that raises the household's borrowing limit over time.

[21] Our choice of a high γ ensures that agents default when they are hit with a bad income shock but do not default under regular economic circumstances. The model exhibits this behavior because when γ is high, the value function for the agent paying her mortgage is much more concave than the value function for the agent who is defaulting, generating a region where default is sensitive to income. In contrast, when γ is low in our model, LTV is the primary determinant of default decisions, which is inconsistent with our empirical findings. We discuss this choice in more detail in Section D.4.

This change translates into consumption today according to the homeowner's MPC out of increased collateral in future dates. We show this decomposition formally in Appendix D.3.3.

This clarifies that the key forces determining the consumption response to debt forgiveness are the timeline under which debt reductions translate into higher cash-on-hand (through lower payments or through home sale) and increased borrowing capacity, and the borrower's MPC out of these future cash-on-hand and future collateral gains. Berger et al. (2018) show that the response to housing wealth gains achieved from increased house prices depends on current home values and the marginal propensity to consume out of wealth. Our analysis adds that the MPC out of wealth gains will depend crucially on a borrower's initial home equity position. Underwater borrowers receiving a dollar of housing wealth are only able to monetize this gain by borrowing or selling their home once they are above water, which may be far in the future. In this case, the consumption response today will depend on a household's MPC out of expected cash or collateral gains far in the future.

In our model, households are unresponsive to cash or collateral gains far in the future. To explain why they are unresponsive, it is helpful to divide households into three categories based on their cash-on-hand relative to permanent income. First, households with low cash-on-hand consume all their assets each period. These households are only responsive to cash or collateral they can access today. Second, households with moderate levels of cash-on-hand are building up a buffer of assets. For these households, near-term cash or collateral reduces the precautionary value of saving in the current period and increases consumption. However, cash or collateral grants several years in the future have no precautionary value and do not affect spending. Third, high cash-on-hand households consume only the annuity value of cash grants regardless of their timing. We show this visually by plotting the consumption response to future cash and future collateral gains in Appendix Figure 32.

The lack of response to future cash and collateral gains can explain why HAMP-type principal reduction failed to increase consumption. We explore this in the model by considering the consumption response to principal forgiveness for a household matching the typical HAMP borrower. By design, principal reduction in HAMP had no incremental effect on payments relative to alternative modifications until year six, and we find that it would be eight years before principal forgiveness translated into increased borrowing capacity.[22] Since no group of borrowers has strong consumption responses to cash and collateral gains so far in the future, we find that this can explain why HAMP did not lead to increased spending in the short term.

This can be seen visually in Appendix Figure 33, which shows the consumption function out of increasing amounts of principal reduction for borrowers starting at an LTV of 150 (the median LTV for HAMP principal reduction recipients). It shows that the consumption function out of home equity gains is S-shaped, convex in a small region below the collateral constraint, and concave above it. Borrowers are insensitive to principal reductions until such reductions bring them close to their constraint. Principal reduction for deeply underwater borrowers does not relax current constraints and has little precautionary value, hence consumption is unaffected. In Appendix Section D.3.6 we show quantitatively that the MPC out of HAMP-like principal reduction is close to zero under a variety of alternative parameterizations.

[22]This eight years estimate is based on the HAMP mortgage contract, which left borrowers underwater with a median LTV ratio of 114 after modification, an assumption that homeowners can only borrow up to 80% of the value of their home in this time period, and an assumption of 1% real annual house price growth based on contemporaneous futures contracts. See Appendix D.3.4 for details supporting the second and third assumption.

Our result contrasts with debt overhang models in which forced deleveraging leads to depressed consumption. For example, in Eggertsson and Krugman (2012) and Guerrieri and Lorenzoni (2017), debt is modeled as a one-period bond. In this setting, when a credit crunch reduces the borrowing limit, borrowers who find themselves beyond the borrowing constraint are forced to immediately cut consumption in order to delever. Applying this assumption to the mortgage setting implies that when housing prices fall such that the LTV ratio becomes greater than 100, borrowers need to immediately repay their outstanding debt until they are above water. Under this hypothetical scenario, borrowers receiving principal reduction would see immediate decreases in the amount of forced repayment. Principal reduction would increase consumption by reducing debt overhang.

But in practice, as in our model, mortgages in particular are long-term loans. Nothing forces borrowers to immediately delever when they are far underwater. Modeling housing debt as a long-term contract removes a mechanical link between debt levels and consumption present in some of the prior literature, and reduces the expected effectiveness of mortgage debt reduction policies. Other recent papers to consider the effect of debt and housing wealth in settings with long-term contracts include Berger et al. (2018), Chen, Michaux and Roussanov (2019), Kaplan, Mitman and Violante (2017), and Justiniano, Primiceri and Tambalotti (2015).

D.3.2 Comparison to Boom-Era Housing MPC Distribution

Our model makes reasonable quantitative predictions about consumption out of housing wealth changes, for which prior empirical papers provide an external benchmark. We focus on replicating estimates corresponding to the pre-2009 period and use Mian, Rao and Sufi (2013) as our external benchmark. We use our model to estimate the MPC out of housing wealth gains for age 45 borrowers with different initial LTVs. We endow each agent with cash-on-hand equal to two years of permanent income, which is the median non-housing wealth for all homeowners in the 2007 Survey of Consumer Finances (SCF).[23] We then calculate the MPC for these agents at different LTV values, and weight them according to the distribution of LTV in 2007 reported in Carter (2012).[24] Thus, heterogeneity in household leverage is a source of MPC heterogeneity as in Auclert (2019).

Appendix Table 6 reports the average MPC out of an additional dollar of housing equity for the average borrower as well as for high-leverage (but still above-water) borrowers. We find MPCs of 8 and 15 cents, respectively. These are similar to the average MPC for homeowners of 9 cents reported in Mian, Rao and Sufi (2013), and the 18 cent MPC of homeowners living in counties with average LTV ratios above 90. In our model, high-leverage above-water borrowers have high MPCs because they have low housing wealth and are the most borrowing constrained.

D.3.3 Sufficient Statistic Expression for Principal Reduction

To build intuition for the effect of principal reductions on consumption, we consider a simplified version of our model without a default option, in which we can develop a straightforward formula for the effect of debt levels on consumption. In this case, a homeowner's problem can be written as a function of four state variables: cash-on-hand (m_{it}), the wealth gain from home sale at retirement ($w_{iT_y} = P_{iT_y} - M_{iT_y}$), and the vectors of housing payments and collateral constraints for the rest of life $\left(\vec{h}_i, \vec{\underline{a}}_i\right)$. We can then decompose the effect of a

[23] Mian, Rao and Sufi (2013) show that wealth does not vary with LTV, so we assign this median number to all borrowers.

[24] Carter (2012) reports LTV distributions in 2005 and 2009, so we take the average.

change in mortgage debt level at date t in the following way:

$$\begin{aligned}\frac{dc_{it}\left(m_{it}, w_{iT_y}, \vec{h}_i, \vec{\underline{a}_i}\right)}{dM_{it}} &= \frac{\partial c_{it}}{\partial w_{iT_y}} \cdot \frac{\partial w_{iT_y}}{\partial M_{it}} + \sum_{s=t}^{T} \frac{\partial c_{it}}{\partial h_{is}} \cdot \frac{\partial h_{is}}{\partial M_{it}} + \sum_{s=t}^{T} \frac{\partial c_{it}}{\partial \underline{a}_{is}} \cdot \frac{\partial \underline{a}_{is}}{\partial M_{it}} \\ &= \underbrace{MPC_{t,w_{iT_y}} \frac{\partial w_{iT_y}}{\partial M_{it}} + \sum_{s=t}^{T} MPC_{t,h_{is}} \cdot \frac{\partial h_{is}}{\partial M_{it}}}_{\text{Future cash-on-hand effect}} + \underbrace{\sum_{s=t}^{T} MPC_{t,\underline{a}_{is}} \cdot \frac{\partial \underline{a}_{is}}{\partial M_{it}}}_{\text{Collateral effect}} \quad (10)\end{aligned}$$

Equation (10) shows that a reduction in debt levels affects today's consumption through two channels. The first is a future cash-on-hand effect. Reducing mortgage debt increases a homeowner's expected home equity gain when they sell the house and reduces their housing payments every year. These translate into consumption according to the homeowner's marginal propensity to consume today out of wealth gains in future dates. The second channel is a collateral effect. Reducing debt levels frees up home equity that raises the household's borrowing limit over time. This change translates into consumption today, according to the homeowner's marginal propensity to consume out of increased collateral in future dates.

D.3.4 Difficulty of Accessing Housing Wealth During Recovery

Three pieces of evidence suggest that borrowers could expect a lengthy delay before being able to access wealth from principal reductions. First, borrowers in our sample are still underwater even after receiving principal reductions, with a median LTV ratio after modification of 114. Furthermore, these leverage ratios only account for first liens, while home equity depends on all liens on a property (i.e., the combined loan-to-value ratio, or CLTV).

Second, the time series of mortgage credit origination shows that credit constraints had tightened during the recovery. Appendix Figure 34a shows mortgage originations by borrower credit score from 2000 to 2015. This covers all mortgages, including second mortgages and home equity lines of credit (HELOCs). It shows that originations dipped sharply after 2007, and for low-credit score borrowers, originations have never recovered. Borrowers receiving HAMP principal reductions had mean FICO scores of 579, with 85 percent below 660, the cutoff for the red line in the figure. This evidence suggests that even with positive equity, the low credit-score borrowers in our sample may have been unlikely to obtain additional housing-related credit. This is further reinforced by Appendix Figure 34b, which shows the time series of average CLTV ratios for borrowers able to obtain HELOCs in a given year. The average CLTV ratio fell 20 points between 2004 and 2009, indicating a tightening of underwriting constraints. Mian and Sufi (2014) argue that tightening credit conditions could explain why the house price recovery from 2011 onward did not contribute significantly to economic activity, since in this case the borrowing channel is restricted. Our results support this hypothesis for underwater borrowers. Furthermore, Agarwal et al. (2018) show that credit expansions during the recovery were more likely to benefit higher-FICO borrowers, precisely those least likely to respond by increasing borrowing.

Third, home price expectations were depressed relative to the boom years. Home price future contracts indicated a market expectation of 1 percent real annual home price growth between 2011 and 2016 (U.S. Department of Housing and Urban Development 2016).

Appendix Figure 35 shows the evolution of borrowing limits and mortgage payments around principal reduction for the average borrower according to our model and using the assumptions described above. We consider an average household with first period income $y_t = 0.85$ units of permanent income, based on Bernstein (2017) who finds that borrowers

receiving mortgage modifications during the recent crisis had temporarily low incomes. We set initial LTV equal to 150, the median pre-modification LTV for borrowers receiving principal reduction in our difference-in-differences analysis.[25] For our treatment group, we then reduce their mortgage balance by $70,000, bringing them to an LTV of 106.[26] To mimic the policy implemented in HAMP we keep mortgage payments for households who have not defaulted fixed for five years.

Principal reduction translates into increased borrowing capacity and increased wealth with a considerable delay. Principal reduction eventually increases borrowing limits, but these increases do not occur for another eight years. This is because even after receiving principal reduction, borrowers remain slightly underwater. Furthermore, to be able to borrow against their home given the collateral constraint they need to get down to an LTV of 80, which takes several years under baseline price growth and mortgage principal pay-down schedules. The bottom panel shows that housing payments decrease substantially, but only starting six years in the future.

D.3.5 Collateral Constraints Drive a Wedge Between MPC Out of Housing Wealth and MPC out of Cash

The inability to monetize housing wealth drives a wedge between an underwater borrower's marginal propensity to consume out of cash and their marginal propensity to consume out of housing wealth. Figure 3b (in the paper, not the appendix) demonstrates this visually. In this figure we take low-cash-on-hand borrowers at various LTV levels and plot the MPC out of $1 of cash or $1 of housing wealth gained by principal reduction.[27] As in the empirical results in Mian, Rao and Sufi (2013), borrowers near their collateral constraint have a high MPC out of housing wealth gains. However, borrowers far underwater are unresponsive to housing wealth changes even though they are highly responsive to cash transfers. This highlights one way that housing wealth is special. Because it can only be monetized when borrowers have positive home equity above a collateral constraint, borrowers respond less to housing wealth gains than they do to cash.

One implication of the wedge between cash and housing MPCs is that in a period where many borrowers are underwater and collateral constraints have tightened, high leverage is a bad "tag" for targeting policies that increase housing wealth, even though it is a good "tag" for targeting policies trying to provide cash to borrowers with high MPCs. Our model suggests that low-wealth, underwater borrowers would have an MPC out of cash around 30 cents. The government spent an average of $0.30 to subsidize each dollar of principal forgiveness in HAMP, for a total government cost of $4.6 billion. Our model suggests that if the same amount of money had been spent on direct transfers to borrowers, the partial equilibrium spending increase would have been $1.4 billion, ten times more than even the upper bound of our estimates for the consumption response to principal forgiveness would suggest. Policies seeking to raise aggregate demand by increasing the housing wealth of leveraged borrowers will be ineffective precisely when policymakers might otherwise want to use them.

[25] This corresponds to an initial home price equal to $173,000 (or 3.25 units of permanent income) and an initial mortgage debt of $259,000 (or 4.88 units of permanent income).

[26] The median LTV post-modification in our data is actually 114, because borrowers' unpaid mortgage payments are capitalized into the new mortgage balance. We abstract from this in our model, though it would only serve to further reduce the effect of principal reduction.

[27] In this experiment, housing payments fall immediately when debt is reduced, unlike in Section D.3.1 where we delayed payment relief in order to mimic principal reduction in HAMP. We consider the average household from our policy experiment, so we set cash-on-hand to 0.85.

The low MPC out of housing wealth for underwater borrowers can help explain the sluggish response to house price gains during the recent recovery (Mian and Sufi 2014). The borrowers ordinarily most responsive to wealth gains may have found themselves unable to translate increased housing wealth into disposable wealth. This also points to a limitation of one of Fisher's policy recommendations for reversing "debt deflations" (Fisher 1933). He suggests reflating asset prices. Our results suggest that this may be ineffective at increasing demand for those who are underwater unless pursued aggressively enough to bring them into positive equity.

D.3.6 Consumption Response to Principal Reduction Under Alternative Parameterizations

In our model, principal reduction is ineffective under a variety of alternative parameterizations. Appendix Table 7 reports the MPC for the principal reduction policy experiment described above under various alternative assumptions. The baseline MPC is 0.3 cents per dollar of mortgage principal reduced. This is similar to our empirical results. Changing borrower age, discount rate, and risk aversion has little impact on the MPC.

Principal reduction remains ineffective even when borrowers have modest access to liquidity. To show this, we calculated the effect of principal reduction assuming households had access to an unused HELOC line worth $20,000, which is twice the amount available to the average household with a HELOC in the 2015 New York Fed Consumer Credit Panel (Federal Reserve Bank of New York 2015). The MPC for this household is still only 0.9 cents. The reason is that households that have access to liquidity are optimizing incorporating this liquid buffer. Principal reduction does not increase their buffer in the near term, so has little effect on the value of maintaining this buffer. This explains why even borrowers who are actively saving or deleveraging, and therefore not literally at their liquidity constraint, are unresponsive to principal reduction. Even when borrowers are saving for precautionary reasons, the increase in housing wealth gained from principal reduction is of little precautionary value because it cannot be monetized for several years.

Generating a large consumption response requires an alternative, unrealized economic environment (relaxed collateral constraints and optimistic home price growth) or an alternative policy of more generous writedowns. Setting the collateral constraint to zero such that homeowners can lever up to 100 LTV generates a moderate MPC of 4.8 cents. Even though borrowers remain underwater after principal reduction, allowing them to monetize wealth starting at 100 LTV would have some immediate precautionary value. Similarly, if households expected permanent real annual house price growth of 5 percent (equal to realized growth rates from 2000 to 2005), the MPC would be 6.2 cents because borrowers would expect to be able to monetize their housing wealth more quickly. Combining both of these assumptions about the economic environment generates a large MPC of 24.2 cents. However, the period when principal reduction was implemented is exactly when neither of these conditions was likely to hold. In the aftermath of the crisis, home price growth expectations were tepid and credit supply was tight.[28]

D.4 Default

In this section, we explore the effect of principal reduction on default. We show that when defaulting imposes utility costs in the short-term, most households only default when they face a large negative income shock. This means that default is relatively insensitive to mortgage balance until borrowers are substantially underwater.

[28]See U.S. Department of Housing and Urban Development (2016) for house price expectations data and Corelogic (2016) for evidence of tightening collateral constraints.

D.4.1 The Effect of Principal Reduction on Default

In forward-looking models with a housing asset and labor income risk, default emerges from two motives: (1) an agent is so far underwater that her house is no longer a good investment and (2) default offers a way to access short-term liquidity when cash-on-hand is low. In our model, the core tradeoff underwater borrowers face when making their default decision is whether the short-term gain from reduced housing payments is worth the utility cost of defaulting and the lost resale value of the house at retirement.[29] Both the costs and benefits of default vary with current payment levels, current incomes, and total debt obligations. When borrowers have high current payments or low current incomes, the short-term payment relief is particularly valuable because it allows borrowers to avoid making severe cuts to consumption. Similarly, when total debt levels are high, the costs of default are low because the house is less valuable as an asset.

To show the effect of principal reduction and relate it back to our empirical results, we simulate changes in mortgage principal holding payments constant. We assume homeowners receive modifications at age 45. To match the low assets of delinquent borrowers in the PSID, we set initial assets $a_t = 0.01$ units of permanent income. We set initial LTV equal to 119, the median pre-modification LTV for borrowers in our regression discontinuity analysis (Table 1a). We then vary the LTV, holding mortgage payments for households that have not defaulted fixed for five years, after which payments fall according to the annuity formula in equation (8) applied to the new mortgage balance.

Appendix Figure 36a shows that for a given current payment level and LTV ratio, there is a cash-on-hand level below which households will find it optimal to default. The more underwater the household, the smaller the income shock necessary to push them to default. For borrowers in our baseline scenario, the income cutoff for defaulting is both low and relatively insensitive to debt levels. In particular, below LTVs of about 150, low-asset borrowers will only default if their income is less than three-quarters of its permanent level, a shock of about two standard deviations. This means that default is most likely to occur for borrowers who are hit with "unemployment," the large liquidity shock in our income process.

We find that default rates are insensitive to principal reduction for the typical borrower. Appendix Figure 36b plots the default rate in the first period after modification for borrowers with various amounts of principal reduction. In our baseline case, additional principal reduction is ineffective below an LTV of about 160. For such moderately underwater borrowers, the gain from defaulting is not worth the cost unless they are hit by a liquidity shock. However, far underwater borrowers have much higher default rates because they default even in the absence of liquidity shocks.

D.4.2 An Optimizing Double Trigger Model of Default

Borrowers in our baseline case exhibit what we call "optimizing double trigger" behavior. In the "double trigger" class of models, agents default when two conditions are triggered: (1) they are underwater and (2) face negative income shocks. In the most basic of these models, agents are not optimizing. While negative equity is necessary for default, the level of negative equity is irrelevant (see description of these models in Gerardi et al. 2018). Agents do not consider the costs and benefits of defaulting, they simply default when they are forced

[29] Because we assume house prices evolve deterministically, our model does not capture the option value of mortgages. With house price uncertainty, paying a mortgage is equivalent to purchasing a call option, giving the borrower the right to "buy" future home equity gains, if realized, at the price of the unpaid balance on the mortgage. Incorporating house price uncertainty would reduce the gain from defaulting and would lead us to estimate a smaller utility cost of defaulting to match the average 10 percent default rate.

to by an income shock that leaves them without enough funds to pay their mortgage (Guren and McQuade 2018).

In our model, agents are optimizing. At moderate levels of underwaterness, it is only optimal for agents with large liquidity shocks (i.e., unemployment in our model) to default. Default is insensitive to negative equity in this region because the costs of default are high and the gains for an employed agent are low. However, beyond about 160 LTV, their optimizing behavior generates a steep causal relationship between LTV and default. These borrowers are defaulting for what is sometimes referred to as "strategic" reasons, that is they default even when their payments are affordable.

The optimizing double trigger behavior, with a small effect of LTV on default at low LTV levels followed by a steep slope at high LTV levels, is consistent with recent dynamic models of mortgage default such as Schelkle (2018) and Campbell and Cocco (2015).[30] Campbell and Cocco (2015) study default decisions in a calibrated model where borrowers are liquidity constrained and face labor income, house price, inflation, and interest rate risk. In their model the kink occurs at about 135 LTV. Below this level, the option value of staying in the mortgage outweighs the gains of defaulting for most borrowers. Our empirical evidence suggests that default is insensitive to LTVs even at slightly higher LTV ratios, which is consistent with adding a utility cost of default to this type of model. The result that borrowers without income shocks do not exercise their default option until substantially underwater is consistent with empirical evidence in Bhutta, Dokko and Shan (2017) , who show that the median homeowner without an income shock does not default until their LTV is greater than 174.[31]

In our model, the key force generating our results is that the income cutoff for defaulting is not very sensitive to the size of mortgage debt. This generates a flat, positive-default-rate region followed by a steep slope at high LTV levels. Generating this region, which is consistent with our empirical evidence, relies on three empirically plausible features of our model. First, most underwater borrowers do not default because they would incur a utility cost of default. This is supported by survey evidence in Guiso, Sapienza and Zingales (2013), who find that about 80 percent of homeowners consider it morally wrong to default when payments are affordable. Second, agents face thick-tailed income shocks (Guvenen, Ozkan and Song 2014).[32] Third, households are risk averse and default when hit with a very bad income shock. When we reduce risk aversion to $\gamma = 2$, default rates are either zero or high, with no flat, positive-default-rate region.[33]

Our empirical evidence favors models like ours over alternatives that generate smooth upward-sloping relationships between LTV and default. Kau, Keenan and Taewon Kim (1993) and Stanton and Wallace (1998) build off of the frictionless option model that predicts a single cutoff LTV value above which all borrowers default. Because the cross-sectional relationship between LTV and default is smooth, these authors propose introducing a distribution of additional default costs, which generates a distribution of cutoff values and therefore

[30] See also Li, White and Zhu (2011).

[31] Similarly, Foote, Gerardi and Willen (2008) study homeowners in Massachusetts who were underwater in the early 1990s and find that fewer than one percent eventually lost their home to foreclosure.

[32] If we eliminate this feature of our income process, we estimate both a smaller stigma cost in order to match an average 10 percent default rate, and we find that default is sensitive to LTVs even at low LTV levels, which is inconsistent with our empirical results.

[33] The short-term liquidity motive for default is most valuable when risk aversion is high. When risk aversion is low, default is largely a function of LTV. As the utility function becomes increasingly linear, the function mapping LTV to default becomes increasingly binary, approaching a rule of thumb where *no* agents default below an LTV cutoff and *all* agents default above the LTV cutoff.

a smooth relationship between LTV and default. We add a distribution of default costs in our model in Appendix Figure 37, and show that this does generate a smooth relationship between LTV and default. However, our empirical results, which find that default is insensitive to LTVs for moderate amounts of underwaterness, reject this parameterization of our model.

References

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. 2018. "Do Banks Pass through Credit Expansions to Consumers Who want to Borrow?" *The Quarterly Journal of Economics*, 133(1): 129–190.

Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *The Quarterly Journal of Economics*, 109(3): 659–84.

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review*, Forthcoming.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra. 2019. "Regional Heterogeneity and the Refinancing Channel of Monetary Policy." *The Quarterly Journal of Economics*, 134(1): 109–183.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. 2018. "House Prices and Consumer Spending." *The Review of Economic Studies*.

Bernstein, Asaf. 2017. "Negative Equity, Household Debt Overhang, and Labor Supply." Working Paper.

Bhutta, Neil, Jane Dokko, and Hui Shan. 2017. "Consumer Ruthlessness and Mortgage Default during the 2007 to 2009 Housing Bust." *The Journal of Finance*, 72(6): 2433–2466.

Cagetti, Marco. 2003. "Wealth Accumulation over the Life Cycle and Precautionary Savings." *Journal of Business & Economic Statistics*, 21(3): 339–353.

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik. 2014. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica*, 82(6): 2295–2326.

Campbell, John Y., and JoÃ£o F. Cocco. 2015. "A Model of Mortgage Default." *The Journal of Finance*, 70(4): 1495–1554.

Campbell, John Y., Nuno Clara, and Joao F. Cocco. 2018. "Structuring Mortgages for Macroeconomic Stability." Working Paper.

Carroll, Christopher. 2012. "Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems." Johns Hopkins Universtiy Mimeo.

Carroll, Christopher D. 1992. "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity*, 1992(2): 61–156.

Carroll, Christopher D. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *The Quarterly Journal of Economics*, 112(1): 1–55.

Carroll, Christopher D. 2006. "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems." *Economics Letters*, 91(3): 312–320.

Carter, George. 2012. "Housing Units With Negative Equity, 1997 to 2009." *Cityscape*, 14(1): 149–165.

Chen, Hui, Michael Michaux, and Nikolai L. Roussanov. 2019. "Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty." *The Journal of Finance*, Forthcoming.

Corelogic. 2016. "Home Equity Lending Landscape." White Paper.

Defusco, Anthony A. 2017. "Homeowner Borrowing and Housing Collateral: New Evidence from Expiring Price Controls." *The Journal of Finance*, 73(2): 523–573.

Diaz, Antonia, and MarÃa JosÃ© Luengo-Prado. 2008. "On the User Cost and Homeownership." *Review of Economic Dynamics*, 11(3): 584–613.

Eberly, Janice, and Arvind Krishnamurthy. 2014. "Efficient Credit Policies in a Housing Debt Crisis." *Brookings Papers on Economic Activity*, 49(2): 73–136.

Eggertsson, Gauti B., and Paul Krugman. 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach." *The Quarterly Journal of Economics*, 127(3): 1469–1518.

Fannie Mae. 2018. "Statistical Summary Tables."

Federal Reserve Bank of New York. 2015. "Quarterly Report on Household Debt and Credit." Federal Reserve Bank of New York, Research and Statistics Group, November 2015.

Fisher, Irving. 1933. "The Debt-Deflation Theory of Great Depressions." *Econometrica*, 1(4): 337–357.

Foote, Christopher L., Kristopher Gerardi, and Paul S. Willen. 2008. "Negative Equity and Foreclosure: Theory and Evidence." *Journal of Urban Economics*, 64(2): 234–245.

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

Goodman, Laurie S., and Jun Zhu. 2015. "Loss Severity on Residential Mortgages: Evidence from Freddie Mac's Newest Data." *The Journal of Fixed Income*, 25(2): 48–57.

Government Accountability Office. 2016. "Treasury Should Estimate Future Expenditures for the Making Home Affordable Program." GAO-16-351.

Guerrieri, Veronica, and Guido Lorenzoni. 2017. "Credit Crises, Precautionary Savings, and the Liquidity Trap." *The Quarterly Journal of Economics*, 132(3): 1427–1467.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2013. "The Determinants of Attitudes toward Strategic Default on Mortgages." *Journal of Finance*, 68(4): 1473–1515.

Guren, Adam, and Timothy McQuade. 2018. "How do Forecolosures Exacerbate Housing Downturns?" Working Paper.

Guren, Adam M, Arvind Krishnamurthy, and Timothy J McQuade. 2018. "Mortgage Design in an Equilibrium Model of the Housing Market." National Bureau of Economic Research Working Paper 24446.

Guvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy*, 122(3): 621–660.

Hembre, Erik. 2018. "HAMP, Home Attachment, and Mortgage Default." Working Paper.

Herkenhoff, Kyle F., and Lee E. Ohanian. 2019. "The impact of foreclosure delay on U.S. employment." *Review of Economic Dynamics*, 31: 63–83.

Himmelberg, Charles, Christopher Mayer, and Todd Sinai. 2005. "Assessing High House Prices: Bubbles, Fundamentals and Misperceptions." *Journal of Economic Perspectives*, 19(4): 67–92.

Holden, Steve, Austin Kelly, Douglas McManus, Therese Scharlemann, Ryan Singer, and John D. Worth. 2012. "The HAMP NPV Model: Development and Early Performance." *Real Estate Economics*, 40(s1): S32–S64.

Imbens, Guido, and Karthik Kalyanaraman. 2012. "Optimal Bandwidth Choice for the Regression Discontinuity Estimator." *Review of Economic Studies*, 79(3): 933–959.

Imbens, Guido W., and Joshua D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica*, 62(2): 467–475.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2015. "Household Leveraging and Deleveraging." *Review of Economic Dynamics*, 18(1): 3–20.

Kaplan, Greg, Kurt Mitman, and Giovanni L Violante. 2016. "Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data." National Bureau of Economic Research Working Paper 22232.

Kaplan, Greg, Kurt Mitman, and Giovanni L Violante. 2017. "The Housing Boom and Bust: Model Meets Evidence." National Bureau of Economic Research Working Paper 23694.

Kau, James B., Donald C. Keenan, and Taewon Kim. 1993. "Transaction Costs, Suboptimal Termination and Default Probabilities." *Journal of the American Real Estate & Urban Economics Association*, 21(3): 247.

Lee, David S., and Thomas Lemieux. 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature*, 48(2): 281–355.

Li, Wenli, Michelle J. White, and Ning Zhu. 2011. "Did Bankruptcy Reform Cause Mortgage Defaults to Rise?" *American Economic Journal: Economic Policy*, 3(4): 123–147.

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

Mayer, Christopher, Edward Morrison, Tomasz Piskorski, and Arpit Gupta. 2014. "Mortgage Modification and Strategic Behavior: Evidence from a Legal Settlement with Countrywide." *American Economic Review*, 104(9): 2830–2857.

McCrary, Justin. 2008. "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test." *Journal of Econometrics*, 142(2): 698–714.

Mian, Atif, and Amir Sufi. 2014. "House Price Gains and U.S. Household Spending from 2002 to 2006." National Bureau of Economic Research Working Paper 20152.

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *The Quarterly Journal of Economics*, 128(4): 1687–1726.

Poterba, James, and Todd Sinai. 2008. "Tax Expenditures for Owner-Occupied Housing: Deductions for Property Taxes and Mortgage Interest and the Exclusion of Imputed Rental Income." *American Economic Review*, 98(2): 84–89.

Scharlemann, Therese, and Stephen H. Shore. 2019. "The Effect of Mortgage Payment Size on Default and Prepayment: Evidence from HAMP Resets." *Working Paper*.

Scharlemann, Therese C., and Stephen H. Shore. 2016. "The Effect of Negative Equity on Mortgage Default: Evidence From HAMP's Principal Reduction Alternative." *The Review of Financial Studies*, 29(10): 2850–2883.

Schelkle, Thomas. 2018. "Mortgage Default during the U.S. Mortgage Crisis." *Journal of Money, Credit and Banking*, 50(6): 1101–1137.

SIGTARP. 2012. "The Net Present Value Test's Impact on the Home Affordable Modification Program."

Stanton, Richard, and Nancy Wallace. 1998. "Mortgage Choice: What's the Point?" *Real Estate Economics*, 26(2): 173–205.

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

United States Department of Housing and Urban Development. 2016. "Housing Scorecard." United States Department of Housing and Urban Development.

United States Department of the Treasury. 2015. "Home Affordable Modification Program Base Net Present Value (NPV) Model v7.0 Model Documentation." United States Department of the Treasury.

United States Department of the Treasury. 2017. "Making Home Affordable: Program Performance Report Through the First Quarter of 2017." United States Department of the Treasury.