*American Economic Review 2020, 110(10): 3100–3138*
*https://doi.org/10.1257/aer.20181243*

# Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession†

*By* Peter Ganong and Pascal Noel*

*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.* (*JEL* E21, G21, G51, R38)

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

* Ganong: Harris School of Public Policy (email: ganong@uchicago.edu); Noel: Booth School of Business (email: pnoel@uchicago.edu). Gita Gopinath was the coeditor for this article. 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, João Cocco, John Coglianese, Marco Di Maggio, Will Dobbie, Jan Eberly, Avi Feller, Xavier Gabaix, John Geanakoplos, Edward Glaeser, Paul Goldsmith-Pinkham, 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 G-2011-6-22 and the National Institute on Aging grant T32-AG000186.
† Go to https://doi.org/10.1257/aer.20181243 to visit the article page for additional materials and author disclosure statements.

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 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

[1] For a review of the academic support for principal reductions, see Zachary Goldfarb, "Economists, Obama Administration at Odds over Role of Mortgage Debt in Recovery," *Washington Post,* November 22, 2012. For example, Goldfarb 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, John Geanakoplos and Susan Koniak argued that a plan to reduce payments and leave principal unchanged "wastes taxpayer money and won't fix the problem" ("Matters of Principal," *New York Times*, March 4, 2009).

[2] See, e.g., Agarwal et al. (2017a, 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).

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 (US 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 \$0.003 per \$1 of principal reduction, with an upper bound of less than \$0.01.

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 1 percent payment reduction reduces default rates by about 1 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,

[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).

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 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 that finds little direct linkage between debt levels and consumption when debt is modeled

[5] A large literature examines the consumption response to house price changes and typically estimates an MPC of around \$0.05 per \$1. 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 (2018), and Leth-Petersen (2010).

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 I describes the data. Sections II and III analyze the effect of principal reduction on default and consumption, respectively. Section IV analyzes the effect of payment reduction on default. Section V provides discussion and interpretation of the empirical results. The final section concludes.

## I. 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.

### A. *Matched HAMP Credit Bureau File*

The US Treasury releases a public data file on the universe of HAMP applicants (US Department of the Treasury 2014b). 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 (2014). 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) documents 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 allow 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 one-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 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 IIIB 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.

### B. *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 (JPMorgan Chase Institute 2020). 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

[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 Section 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 IB, we can isolate purchases and confirm that our results are unchanged.

[8] See online Appendix Section 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 (US Department of the Treasury 2014a).

are owned or have been securitized by Chase. We limit the sample to modifications performed in 2011:IV 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 IV.

## II. 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.

### A. *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 borrowers 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. Panel A of Figure 1 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

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

[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. 2017a, Table 1).

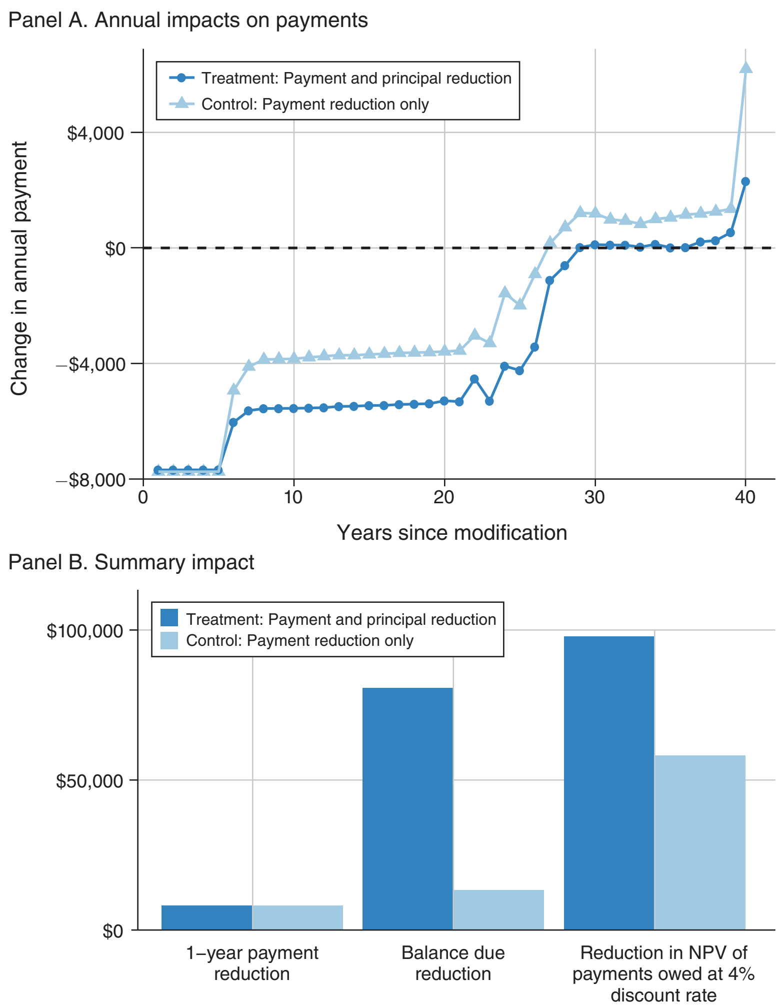


Figure 1. Financial Impact of Modifications with and without Principal Reduction

*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 ninety-fifth percentile; see online Appendix Figure 2 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 4 percent interest rate. See Section IIA for details.

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. A total of 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 toward supporting principal reduction modifications. On average, 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. Panel A of Figure 1 shows that payment reductions are identical for the first five years, after which payments rise more sharply for borrowers with payment reduction modifications. Panel B 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 Section 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 3. Borrowers appear eager to take up modifications. Conditional on being offered a modification, 97 percent of borrowers accepted the offer.

[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.

### B. *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. In this section, we discuss the cutoff rule, which we use in a regression discontinuity to estimate the impact of principal reduction on default. We use variation in servicer and lender type to estimate consumption impacts, and so we defer an explanation of that variation to Section III.

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 2009:IV, just before the peak of the delinquency crisis, which did not begin abating until 2013. Online Appendix Figure 4 plots the delinquency rate for all US 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 US 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 IIA. 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 panel A of Figure 2.

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

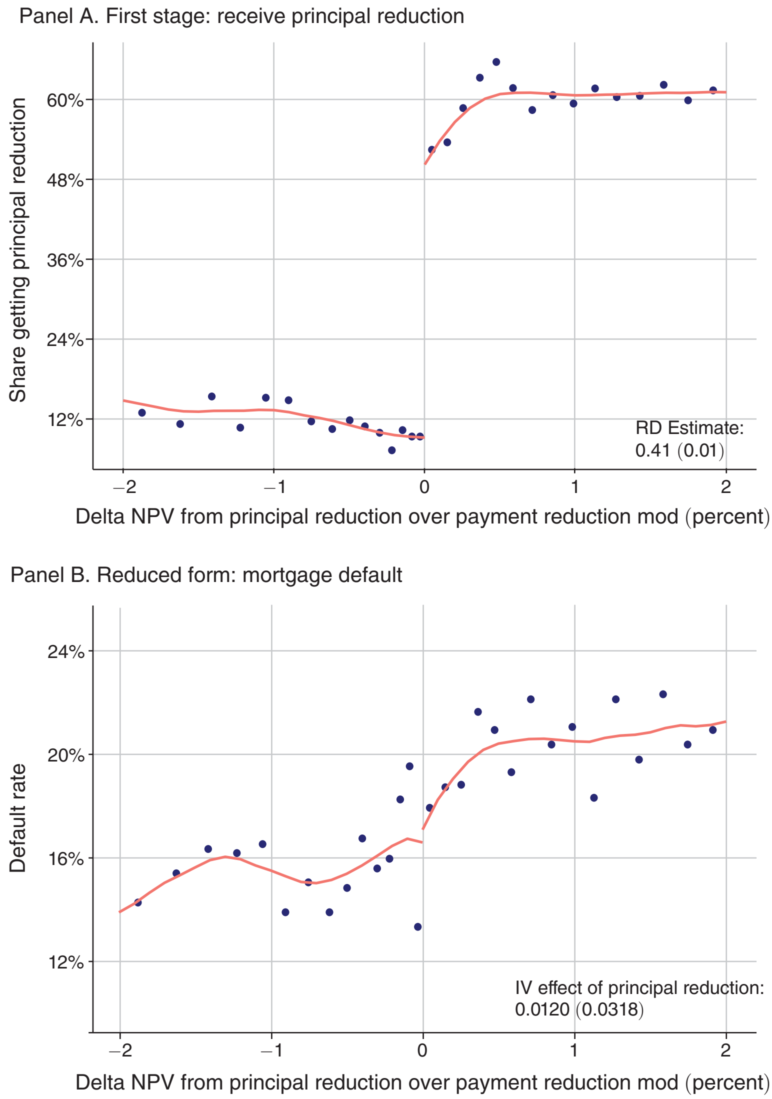


Figure 2. Effect of Principal Reduction on Default

*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 IIB.

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 Section 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 Section B.1.2, and further discuss the normalization in online Appendix Section 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 $\tau$ 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 panel A of Table 1, 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 predetermined covariates trend smoothly through the cutoff, as shown in online Appendix Figure 5. Some servicers ran only one NPV calculation and reported this single number as the NPV calculation for both "payment reduction" and "payment and principal reduction" modifications, meaning that they reported $ENPV(1,X) = ENPV(0,X)$. Following the advice of US Treasury staff, we assume that observations exactly at zero reflect misreporting and subsequently drop them from the analysis sample. Online Appendix Figure 6 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 Section B.1.3.

### C. *Results: Effect of Principal Reduction on Default*

Panel A of Figure 2 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 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 7.

[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.

TABLE 1—REPRESENTATIVENESS

| | RD analysis sample | | | | PSID delinquent households | | | |
|---|---|---|---|---|---|---|---|---|
| | Mean | p10 | p50 | p90 | Mean | p10 | p50 | p90 |
| *Panel A. Principal reduction regression discontinuity sample* | | | | | | | | |
| 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 | 128 | 104 | 121 | 167 | 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 |
| Observations | 9,725 | | | | 190 | | | |
| *Panel B. Payment reduction regression discontinuity sample* | | | | | | | | |
| Income | 67,811 | 27,076 | 54,623 | 125,094 | 64,000 | 21,000 | 55,000 | 120,000 |
| Home value | 190,341 | 49,620 | 140,000 | 400,000 | 190,000 | 50,000 | 140,000 | 350,000 |
| Loan to value ratio | 129 | 63 | 107 | 205 | 101 | 52 | 94 | 166 |
| Monthly mortgage payment | 1,327 | 496 | 1,055 | 2,567 | 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 |
| Observations | 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.

We find that principal reduction has no impact on default. Panel B of Figure 2 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 IID. There is no jump in default rates at the cutoff, and we can rule out a reduction of more than 5 percentage points using the 95 percent confidence interval. Online Appendix Figure 8 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 Section 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 (US 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 (US 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 write-down 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 Section C.2 contains additional detail on the calculations in the two 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 write-downs 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.

### D. *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) uses data on modifications performed prior to HAMP and finds, using cross-sectional variation, that borrowers who received principal reductions equivalent to ours saw an 18 percentage point reduction in default. Furthermore,

[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.

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 US 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 re-default model (US 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 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 Scharlemann and Shore (2016)—henceforth, SS—which uses 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 nonlinearities 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]

[15] Outside of mortgages, Dobbie and Song (2019) analyzes 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.

[16] SS explains 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 writes: "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

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 estimates 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 payment reduction on default, to which we turn next.

## III. 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.

### A. *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 underpowered 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 IIIC, even a small change in consumption on the order of $0.05 for each $1 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 allows 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,

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.

[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 pretreatment 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.

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.

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 (US Department of the Treasury 2014b). 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 3, 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

$$(3)\quad y_{i,g,t,s} = \gamma_g + \gamma_t + \gamma_{m(i),s} + \beta\big(PrincipalReduction_g \times Post_t\big) + x'_{it}\delta + \varepsilon_{i,g,t,s},$$

where $i$ denotes borrowers, $g \in \{\text{payment reduction, payment \& principal reduction}\}$ 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

[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.

[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] Panel A of online Appendix Figure 9 normalizes expenditure to zero at modification date in order to more clearly show the parallel pretrends. Panel B plots mean auto expenditure around modification date and similarly demonstrates parallel pretrends.

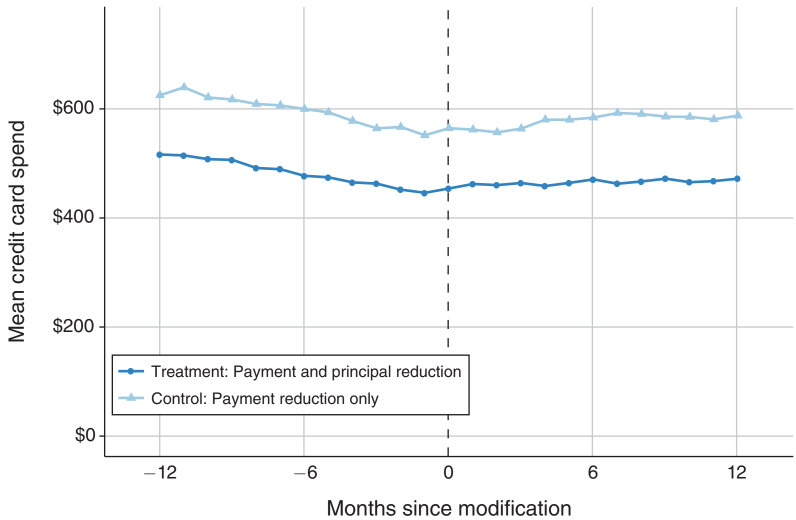


Figure 3. Effect of Principal Reduction on Consumption

*Notes:* This figure empirically evaluates the impact of principal forgiveness on consumption. It 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 III for details.

proxy for nondurable and durable spending, respectively. The term $\gamma_g$ captures the modification group fixed effect, and $\gamma_t$ captures a fixed effect for each month relative to modification. The variable $PrincipalReduction_g$ is a dummy 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 $\beta$, 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. The term $x_i$ is a vector of individual characteristics designed 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 $\left(x'_{it} = (x_i \quad x_i \times Post_t)'\right)$.

## B. *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,

[21] This includes the predicted gain to lenders from providing principal reduction, the predicted gain interacted with a dummy variable equal to 1 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.

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 IA. 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 10 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 11. 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 VB.

*Effect of Payment Reduction on Consumption*.—A natural concern with our 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. (2017a). 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

[22] Translating these estimates to a marginal propensity to consume, as in Section IIIC, our confidence interval ranges all the way from −\$0.15 to \$0.41.

Table 2—Impact of Principal Reduction on Expenditure

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| *Panel A. Credit card expenditure (\$/month)* | | | | | | |
| Treatment (Principal reduction × post) | 0.686 | 0.721 | 0.811 | 2.068 | 0.496 | 2.144 |
| | (3.621) | (3.619) | (3.685) | (3.855) | (3.887) | (3.913) |
| MSA fixed effects | | Yes | | | | |
| Calendar month fixed effects | | | Yes | | | |
| MSA by calendar month fixed effects | | | | Yes | Yes | Yes |
| Controls | | | | | Yes | Yes |
| Controls × 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 |
| *Panel B. Auto expenditure (\$/month)* | | | | | | |
| Treatment (Principal reduction × post) | 13.541 | 13.545 | 13.062 | 14.648 | 15.327 | 10.994 |
| | (8.824) | (8.824) | (8.896) | (9.011) | (9.135) | (9.312) |
| MSA fixed effects | | Yes | | | | |
| Calendar month fixed effects | | | Yes | | | |
| MSA by calendar month fixed effects | | | | Yes | Yes | Yes |
| Controls | | | | | Yes | Yes |
| Controls × 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 III for additional detail on the specification, outcome measures, and sample.

reductions are responsible for the consumption impact they estimate. In online Appendix Section 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 IV, but we conclude that this strategy is underpowered for studying consumption impacts.

## C. *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 \$0.003 per \$1 of principal reduction, with the upper bound of the 95 percent confidence interval corresponding to \$0.009. 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 \$0.007 and an upper bound of \$0.018. 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 of the MPC out of housing wealth increases (which affect both wealth and liquidity). We interpret the lessons from this result in Section VB.

## IV. 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.

### A. *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 IIA 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 IB. 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

[23] Specifically, our adjustment factor is the product of two ratios: (i) the ratio of the number of credit cards in TransUnion to the number of credit cards with spending reported in TransUnion, and (ii) 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 spending is from Census, and consumer credit card spending is from Federal Reserve Payment Study (US 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 spending to the average credit card spending we observe in our sample gives the same MPC point estimate.

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 IIA. 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 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 IIA, 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. Panel A of Figure 4 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

[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) shows that the GSE program effectively 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.

deeper payment reductions for the first 22 years and then to higher payments for the remaining 18 years.

Panel B 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 IIA. 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 II we find that future mortgage payments do not affect default 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.

### B. *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 one-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 one-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

[27] As we discuss above, because we found no independent impact of principal reductions on default in Section IIC, we attribute the causal impact of this treatment on default to the immediate payment reduction portion of the treatment. In Section IVC 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.

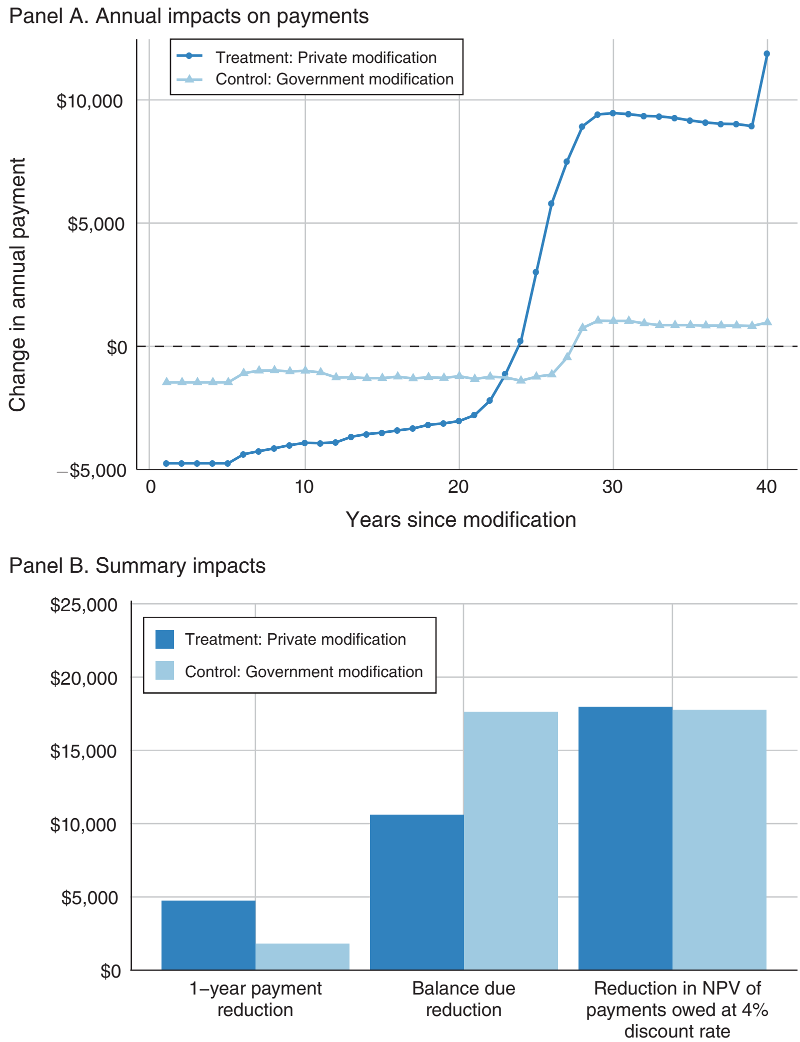


Figure 4. Financial Impacts of Private and Government Modifications

*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 4 percent interest rate. See Section IV for details.

are those failing to meet one of the other eligibility criteria described in Section IIA. 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 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. Panel B of Table 1 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 12 shows that predicted default rates based on predetermined covariates trend smoothly through the cutoff. We provide additional detail on covariate balance in online Appendix Section B.2.1. Finally, online Appendix Figure 13 shows that borrower density is also smooth around the cutoff. The lack of bunching indicates that there is no manipulation of the running variable.

### C. *Results: Effect of Payment Reduction on Default*

Panel A of Figure 5 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

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

We follow the same procedures as in Section IIB 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. Panel B of Figure 5 shows the reduced form, plotting the default rate on the $y$-axis.

[28] Online Appendix Figure 14 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.

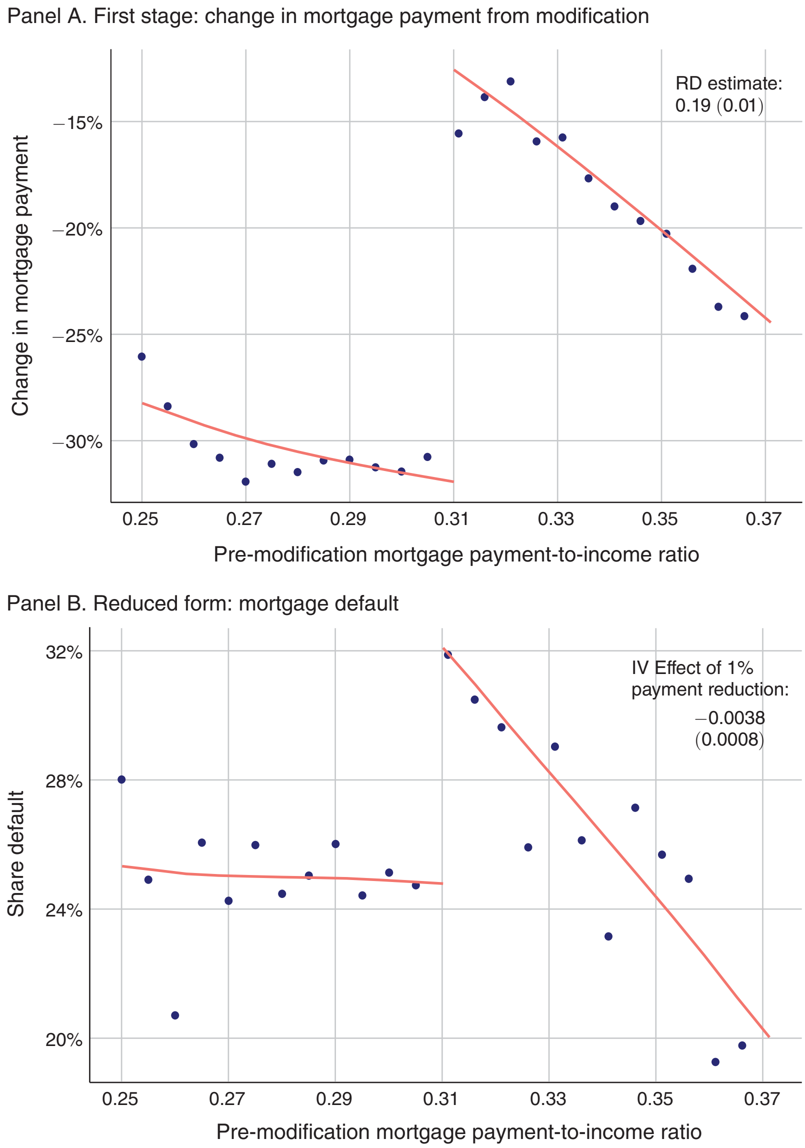


Figure 5. Effect of Payment Reduction on Default

*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 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 IVB.

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 panel B mirrors the slope of the payment reductions in panel A; default rates are approximately constant on the left-hand side of the cutoff, consistent with a constant amount of payment reduction, and 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 15, 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 1 percent payment reduction reduces default rates in the two years after modification by 0.38 percentage points, 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 15 shows that when we remove all three of the sample restrictions described above we find that an extra 1 percent payment reduction reduces default rates by 0.25 percentage points. 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 16 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 IIC to adjust for any potential effect of the relative increase in mortgage principal at the cutoff. Online Appendix Section 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 a private modification that was very similar to the Chase

[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) discusses, 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) analyzes the effect of sustained payment increases for non-distressed borrowers in a setting where these payment increases are predetermined 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), which shows that liquidity provided by unemployment insurance benefits also helps reduce mortgage default.

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 1 percent reduction in payments reduces default by 0.26 percentage points. 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 20. 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.

## V. Discussion and Interpretation

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

### A. *Default*

In our sample immediate liquidity substantially affects default, but reduction in mortgage principal does not. From a positive perspective, this is consistent with

[31] The sample of GSE modifications covers the same 2011 through 2016 time period as the sample of Chase modifications. Online Appendix Figures 17–20 replicate our 31 percent PTI cutoff analysis for the GSE-backed loans sample. Recall that in online Appendix Figure 14 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 19. There is little change in predicted delinquency at the cutoff. Second, the marginal borrowers on the right-hand side are more credit-worthy. This change in sample composition will lower average default rates on the right-hand side, leading us to understate the effectiveness of payment reduction.

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) estimates 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 IIC, 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) finds 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 (US Department of Housing and Urban Development 2010). In addition, we estimate in Section IIC 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 IV, 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. 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

[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).

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 Section 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 redesigning modification programs requires extrapolating 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

[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 Section 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. (2017a) examines the extensive margin in terms of the quantity of modifications and finds little evidence of crowd out.

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 15. 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 panel B of online Appendix Figure 7. Yet panel B of Figure 2 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 21.

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 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.

[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 Section C.4. In this Appendix we also compare the cost of private modifications to the cost of our proposed efficient modification.

## B. *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 \$0.05 per \$1.[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 *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]

[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 (2018), and Leth-Petersen (2010). See also Berger, Turner, and Zwick (2016); Argyle, Nadauld, and Palmer (2019); and Fadlon, Ramnath, and Tong (2019) which 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.

[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.

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 Section 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 6. 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 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 6 suggests that providing direct liquidity to low-wealth,

[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 Section D.2.

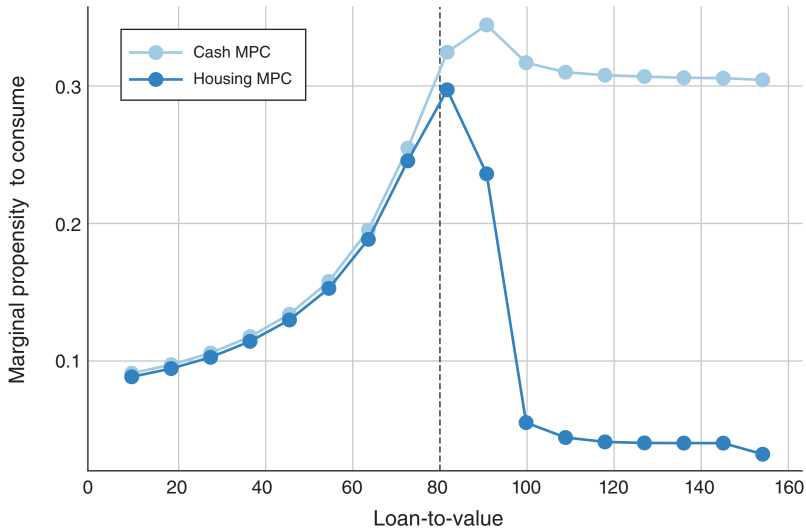


Figure 6. Marginal Propensity to Consume (MPC) by Home Equity Position in Model

*Notes:* This figure theoretically evaluates the impact of principal forgiveness on consumption. It 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. We assume a collateral constraint of 80 percent of LTV and cash on hand (assets + annual income) equal to 85 percent of permanent income. See Section VB for further discussion.

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 6 shows that all of the consumption increase would be coming from the region near the collateral constraint. This means that a policy of targeting

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

deeply underwater borrowers with more generous write-downs will expend substantial resources in a region with no stimulative effect.[46]

## VI. 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, i.e., wealth without liquidity, has no detectable effect on default or consumption. In sharp contrast, we find that short-term payment reduction with no change in the net present value of payments owed, i.e., 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 one-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 Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class." *Review of Financial Studies* 29 (7): 1635–70.

**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): 559–78.

**Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru.** 2017a. "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, Tomasz Piskorski, Amit Seru, and Vincent Yao.** 2017b. "Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinancing Program." NBER Working Paper 21512.

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

[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.

**Aladangady, Aditya.** 2017. "Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata." *American Economic Review* 107 (11): 3415–46.

**Albanesi, Stefania, Giacomo De Giorgi, and Jaromir Nosal.** 2017. "Credit Growth and the Financial Crisis: A New Narrative." NBER Working Paper 23740.

**Argyle, Bronson, Taylor D. Nadauld, and Christopher Palmer.** 2019. "Monthly Payment Targeting and the Demand for Maturity." NBER 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, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra.** 2018. "House Prices and Consumer Spending." *Review of Economic Studies* 85 (3): 1502–42.

**Berger, David, Nicholas Turner, and Eric Zwick.** 2016. "Stimulating Housing Markets." NBER Working Paper 22903.

**Bhutta, Neil, Jane Dokko, and Hui Shan.** 2017. "Consumer Ruthlessness and Mortgage Default during the 2007 to 2009 Housing Bust." *Journal of Finance* 72 (6): 2433–66.

**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): 220–56.

**Campbell, John Y., Nuno Clara, and João F. Cocco.** 2018. "Structuring Mortgages for Macroeconomic Stability." Unpublished.

**Campbell, John Y., and João 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." In *NBER Macroeconomics Annual 1989,* Vol. 4, edited by Olivier Jean Blanchard and Stanley Fischer, 185–246. Chicago: University of Chicago Press.

**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.

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

**Chen, Hui, Michael Michaux, and Nikolai L. Roussanov.** 2013. "Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty." SSRN Paper 2024392.

**Cloyne, James, Kilian Huber, Ethan Ilzetzki, and Henrik Kleven.** 2019. "The Effect of House Prices on Household Borrowing: A New Approach." *American Economic Review* 109 (6): 2104–36.

**DeFusco, Anthony A.** 2018. "Homeowner Borrowing and Housing Collateral: New Evidence from Expiring Price Controls." *Journal of Finance* 73 (2): 523–73.

**DeFusco, Anthony A., Charles G. Nathanson, and Eric Zwick.** 2017. "Speculative Dynamics of Prices and Volume." NBER Working Paper 23449.

**Di Maggio, Marco, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, and Amit Seru.** 2017. "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging." *American Economic Review* 107 (11): 3550–88.

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

**Eberly, Janice, and Arvind Krishnamurthy.** 2014. "Efficient Credit Policies in a Housing Debt Crisis." *Brookings Papers on Economic Activity*: 73–118.

**Eggertsson, Gauti B., and Paul Krugman.** 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach." *Quarterly Journal of Economics* 127 (3): 1469–1513.

**Ehrlich, Gabriel, and Jeffrey Perry.** 2015. "Do Large-Scale Refinancing Programs Reduce Mortgage Defaults? Evidence From a Regression Discontinuity Design." Congressional Budget Office Working Paper 2015-06.

**Fadlon, Itzik, Shanthi P. Ramnath, and Patricia K. Tong.** 2019. "Household Responses to Transfers and Liquidity: Evidence from Social Security's Survivors Benefits." NBER 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." Unpublished.

**Fuster, Andreas, and Paul S. Willen.** 2017. "Payment Size, Negative Equity, and Mortgage Default." *American Economic Journal: Economic Policy* 9 (4): 167–91.

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

**Ganong, Peter, and Pascal Noel.** 2020. "Replication Data for: Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E118367V1.

**Geithner, Timothy.** 2014. *Stress Test: Reflections on Financial Crises.* New York: Crown Publishers.

**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." 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." *Review of Financial Studies* 31 (3): 1098–1131.

**Gorea, Denis, and Virgiliu Midrigan.** 2017. "Liquidity Constraints in the U.S. Housing Market." NBER Working Paper 23345.

**Greenwald, Daniel, Tim Landvoigt, and Stijn Van Nieuwerburgh.** 2018. "Financial Fragility with SAM?" MIT Sloan Research Paper 5261-17.

**Guerrieri, Veronica, and Guido Lorenzoni.** 2017. "Credit Crises, Precautionary Savings, and the Liquidity Trap." *Quarterly Journal of Economics* 132 (3): 1427–67.

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

**Guren, Adam M., Alisdair McKay, Emi Nakamura, and Jon Steinsson.** 2018. "Housing Wealth Effects: The Long View." NBER Working Paper 24729.

**Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw.** 2001. "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design." *Econometrica* 69 (1): 201–9.

**Haughwout, Andrew, Donghoon Lee, Joseph S. Tracy, and Wilbert van der Klaauw.** 2011. "Real Estate Investors, the Leverage Cycle, and the Housing Market Crisis." Federal Reserve Bank of New York Staff Report 514.

**Haughwout, Andrew, Ebiere Okah, and Joseph Tracy.** 2016. "Second Chances: Subprime Mortgage Modification and Redefault." *Journal of Money, Credit, and Banking* 48 (4): 771–93.

**Hedlund, Aaron.** 2015. "Failure to Launch: Housing, Debt Overhang, and the Inflation Option during the Great Recession." Unpublished.

**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): 933–59.

**Ji, Yan.** 2018. "Job Search under Debt: Aggregate Implications of Student Loans." SSRN Paper 2976040.

**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.

**JPMorgan Chase Institute.** 2020. "Mortgage Modification and Servicing Data."

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

**Mian, Atif, and Amir Sufi.** 2011. "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis." *American Economic Review* 101 (5): 2132–56.

**Mian, Atif R., and Amir Sufi.** 2018. "Credit Supply and Housing Speculation." SSRN Paper 209564.

**Mulligan, Casey B.** 2009. "Means-Tested Mortgage Modification: Homes Saved or Income Destroyed?" NBER Working Paper 15281.

**Nathanson, Charles G., and Eric Zwick.** 2018. "Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market." *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?" Unpublished.

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

**Scharlemann, Therese, and Stephen H. Shore.** 2016. "The Effect of Negative Equity on Mortgage Default: Evidence from HAMP's Principal Reduction Alternative." *Review of Financial Studies* 29 (10): 2850–83.

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

**Sinai, Todd, and Nicholas S. Souleles.** 2005. "Owner-Occupied Housing as a Hedge against Rent Risk." *Quarterly Journal of Economics* 120 (2): 763–89.

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

**TransUnion.** 2014. "CreditVision." TransUnion. Provided through agreements with Harvard University and Booth School of Business.

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

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

**US Department of the Treasury.** 2014b. "Making Home Affordable Data Files." https://www.treasury.gov/initiatives/financial-stability/reports/Pages/mha_publicfile.aspx (accessed March 3, 2015).

**US Department of the Treasury.** 2014c. "Making Home Affordable Program Performance Report through the Second Quarter of 2014." Washington, DC: United States Department of the Treasury.

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

**US Department of the Treasury.** 2017. "Making Home Affordable: Program Performance Report Through the First Quarter of 2017." Washington, DC: United States Department of the Treasury.

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

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