

The Effect of Debt on Default and Consumption: Evidence from Housing Policy in the Great Recession

Peter Ganong and Pascal Noel*

December 16, 2017

Abstract

This paper empirically and theoretically analyzes the effect of debt reductions that reduce long-term but not short-term obligations. Isolating the effect of future obligations allows us to test alternative explanations for borrower default decisions and to analyze the consumption response to mortgage principal reduction for underwater borrowers. Our empirical analysis uses regression discontinuity and difference-in-differences research designs on de-identified bank account and credit bureau records from participants in the U.S. government's Home Affordable Modification Program. We find that mortgage principal reductions worth an average of \$70,000 have no impact on default or consumption for borrowers who remain underwater. Our results are sufficiently precise to rule out economically meaningful effects. We develop a quantitative life-cycle model that clarifies that borrowers' short-term constraints govern their response to long-term debt obligations. When defaulting imposes utility costs in the short-term, default is driven by cash-flow shocks such as unemployment rather than by future debt burdens. When principal reductions do not push borrowers sufficiently above water so as to relax collateral constraints, consumption is unaffected because borrowers are unable to monetize increased housing wealth. Collateral constraints drive a wedge between an underwater borrower's marginal propensity to consume out of cash and their marginal propensity to consume out of housing wealth. Our results help explain why policies that lowered current mortgage payments were more effective than principal reductions at stemming foreclosures and increasing demand during the Great Recession.

*ganong@uchicago.edu, pascal.noel@chicagobooth.edu. We thank Alex Bartik, David Berger, Laura Blattner, John Campbell, Gary Chamberlain, Raj Chetty, Gabriel Chodorow-Reich, Joao Cocco, John Coghlanese, Marco Di Maggio, Will Dobbie, Emmanuel Farhi, Avi Feller, Xavier Gabaix, Edward Glaeser, Paul Goldsmith-Pinkham, Adam Guren, Sam Hanson, Nathan Hendren, Kyle Herkenhoff, Larry Katz, Rohan Kekre, David Laibson, Yueran Ma, Laurie Maggiano, Atif Mian, Kurt Mitman, Bill Murphy, Vijay Narasiman, Charles Nathanson, Elizabeth Noel, Christopher Palmer, Jonathan Roth, David Scharfstein, Therese Scharlemann, Andrei Shleifer, Mark Shepard, Jon Spader, Jeremy Stein, Amir Sufi, Larry Summers, Adi Sunderam, Owen Zidar, and Eric Zwick. Thanks to the team at Heterogeneous Agent Resource and toolKit (HARK, <https://github.com/econ-ark/HARK>) for sharing simulation code. Technical support was provided by the Research Technology Consulting team at the Institute for Quantitative Social Science. 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, Ganong and Noel were paid contractors of JPMCI. We gratefully acknowledge funding from the Joint Center for Housing Studies, the Washington Center for Equitable Growth, the Hirtle Callaghan Fund, and the National Bureau of Economic Research through the Alfred P. Sloan Foundation Grant No. G-2011-6-22 and the National Institute on Aging Grant No. T32-AG000186.

1 Introduction

The depressed house prices, 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 has explained that the government’s “biggest debate was whether to try to reduce overall mortgage loans or just monthly payments.”² A wide range of economists have argued that failing to address debt levels by permanently forgiving mortgage principal was a missed opportunity and one of the biggest policy mistakes of the Great Recession.³ Others have argued instead that if borrowers are liquidity constrained, focusing on short-term payment reductions is more cost effective (Eberly and Krishnamurthy 2014).

The normative policy debate hinges on fundamental economic questions about the effect of long-term debt obligations. For default, the underlying question is whether it is primarily driven by cash-flow shocks such as unemployment or by overall debt obligations, sometimes known as “strategic default.” For consumption, the underlying question is whether underwater borrowers have a high marginal propensity to consume (MPC) out of changes in housing wealth.⁴ A broad literature evaluates the effects of changes in debt obligations that simultaneously reduce both short-term and long-term payments and consistently shows that debt obligations matter.⁵ Reducing payments leads to decreased defaults and increased consumption. However, to fully understand the effect of debt on borrower decisions and to inform the debate about principal versus payment reductions, it is essential to distinguish the effect of long-term obligations alone.

We fill this void by exploiting a distinctive policy that allows us to isolate the effect of long-term obligations. Our data show that mortgage principal reductions which do not affect short-term payments but substantially reduce long-term obligations have no significant impact on default or consumption for underwater borrowers. Our estimates are sufficiently precise to rule out economically meaningful effects. We build a partial-equilibrium life-cycle model that shows that borrowers’ short-term constraints govern their response to long-term debt obligations. This can explain why principal reduction is ineffective for underwater borrowers. When defaulting imposes short-term costs, default is driven by immediate cash-flow shocks rather than by future debt burdens. When principal reductions do not push borrowers sufficiently above water so as to relax collateral constraints, consumption is unaffected. We show that collateral constraints drive a wedge between an underwater borrower’s marginal propensity to consume out of cash and their marginal propensity to consume out of housing wealth. Thus, policies to increase housing wealth are unable to stimulate demand when borrowers are so far underwater that home equity gains fail to relax binding

²Geithner (2014)

³See Goldfarb (2012) for a review of the academic support for principal reductions. See also Feldstein (2009), Geanakoplos and Koniak (2009), Romer (2011), and Mian and Sufi (2014b).

⁴An “underwater” borrower owes more on their home than its current market value. This is also referred to as having negative equity.

⁵See Agarwal et al. (2016a); Agarwal et al. (2015); Di Maggio et al. (2016); Ehrlich and Perry (2015); Fuster and Willen (2015); Keys et al. (2014); and Narasiman (2016)

collateral constraints, which was precisely the situation when such policies were implemented.

We study the effect of long-term debt obligations by comparing underwater borrowers who received two types of modifications in the federal government’s Home Affordable Modification Program (HAMP). Both modification types resulted in identical payment reductions for the first five years. However, one group also received \$70,000 in mortgage principal reduction, which translated into increased home equity and substantial long-term payment relief. By comparing borrowers in each of these modification types, we isolate the effects of long-run debt levels holding fixed short-run payments. Throughout the paper we use “principal reduction” to refer to a policy that reduces mortgage principal without affecting short-term payments. An important feature of the policy we study is that borrowers remained underwater even after substantial debt forgiveness.

Our empirical strategy exploits quasi-experimental variation in borrower assignment to the two modification types. Mortgage servicers evaluated underwater applicants for both modification types by calculating the expected gain to investors under either option using a standard government-supplied formula. There is a sharp 41 percentage point jump in the probability that a borrower receives principal reduction when the calculation shows that principal reduction is marginally more beneficial to investors. Our first research design is a regression discontinuity estimator comparing borrowers on either side of this cutoff. Borrower assignment was also influenced by servicer and investor type. Because of legal and regulatory barriers, borrowers whose loans were held in servicers’ own portfolios were more likely to receive principal reduction than those held in securitized pools or guaranteed by the GSEs. This motivates our second research design: a panel difference-in-differences estimator.

We build two new datasets with information on borrower outcomes and HAMP participation. Our first dataset matches administrative data on HAMP participants to monthly consumer credit bureau records. We exploit detailed account-level information to construct a novel measure of consumer spending based on monthly credit card expenditures. We also follow Keys et al. (2014) and Di Maggio et al. (2016) by using new auto loan originations as a proxy for durable consumption. The dataset covers a broad sample population. Furthermore, it crucially includes the calculation of expected investor gain from each modification type, which allows us to implement our regression discontinuity strategy exploiting cutoffs in this calculation. Our second dataset uses de-identified data assembled by the JPMorgan Chase Institute (JPMCI). It includes mortgage, credit card, and checking account information for borrowers who received a HAMP modification from Chase. The JPMCI dataset has a smaller sample size and does not include the investor gain calculation. However, it includes richer outcome measures such as income as derived from checking account transactions.

Using a regression discontinuity empirical design, we find that principal reduction has no effect on default. The analysis exploits the 41 percentage point jump in the probability borrowers receive principal reduction when it is marginally more beneficial for investors. To the right of the cutoff, borrowers receive a marginal \$34,000 of principal forgiveness, which translates to an 11 percentage point reduction in their loan-to-value (LTV) ratio with no change in their short-term monthly

payment. Borrower characteristics trend smoothly through this cutoff, validating the use of a regression discontinuity strategy. Our point estimate is that principal reduction changes default probabilities in the first three years by less than one tenth of a percentage point, and we can rule out a default reduction of more than 5 percentage points. We calculate that even at the upper bound of our confidence interval, the government spent \$800,000 per avoided foreclosure. The program cost is over an order of magnitude greater than estimates of the social cost of foreclosures (Hembre 2014).

Our finding that default is unresponsive to principal reduction is surprising because prior research and policy projections based on historical data had suggested large benefits from principal reduction. For example, Haughwout et al. (2016) use data on modifications performed prior to HAMP and find using cross-sectional variation that an equivalent 11 point reduction in LTV is associated with an 18.0 percentage point reduction in default rates. Additionally, the U.S. Treasury built a redefault model based on the cross-sectional relationships in historical data in order to calculate the benefits to investors of providing modifications. This model predicts a 6.6 percentage point reduction in default rates from principal reduction at the cutoff, an effect size we can rule out. The cross-sectional evidence may have been misleading because borrowers with less equity purchased homes near the height of the credit boom and might have been less creditworthy on other dimensions (Palmer 2015). In addition, prior area-level evidence suggested larger price reductions are associated with increased default rates, though several authors have emphasized the difficulty of inferring a causal link between the degree of negative equity and default rates due to the possibility of correlated omitted economic shocks (Adelino et al. 2016, Scharlemann and Shore 2016).

While we find no causal impact of long-term debt obligations on default decisions, we provide suggestive evidence that default is responsive to short-term cash-flow shocks. We analyze the monthly path of income and mortgage payments around defaults in the JPMCI bank account data. We find that missed mortgage payments coincide with a sharp decline in income. This high-frequency evidence is consistent with other recent empirical findings on the relationship between default and liquidity shocks such as income loss (Gerardi et al. 2015, Hsu et al. 2014) and monthly payment changes (Fuster and Willen 2015, Agarwal et al. 2016a, Ehrlich and Perry 2015).

We examine the effect of principal reduction on consumption in the second part of our empirical analysis. Our preferred empirical strategy is a panel difference-in-differences estimator which is more precise than our regression discontinuity strategy.⁶ We compare borrowers who received principal reductions to other underwater borrowers who received only payment reductions. We show that these two groups of borrowers are ex-ante similar on a broad range of observable characteristics, and that their credit card and auto spending measures were trending similarly in the months

⁶While we also report results from the regression discontinuity strategy, there are two reasons we prefer our panel differences-in-differences estimator. First, our regression discontinuity strategy is under-powered for studying economically meaningful changes in consumption. Second, we have lagged consumption measures, which allows us to adjust for underlying differences between groups within a wider bandwidth than the regression discontinuity strategy, providing a more precise estimate.

before modification, supporting the identification assumption underlying the strategy. We find that \$70,000 in principal reduction has no significant impact on underwater borrowers' credit card or auto expenditure. The result holds both in our matched HAMP-credit bureau data and in the JPMCI bank dataset. Translating our results into an annual MPC for total consumption, our point estimate is that borrowers increased consumption by 0.2 cents per dollar of debt forgiveness, with an upper bound of 0.8 cents.

Our estimated MPC out of housing wealth for underwater borrowers is an order of magnitude smaller than the marginal propensity to consume for average homeowners examined in prior studies. A set of comparison points for our estimates of the impact of principal reduction on consumption comes from the literature examining the impact of price-driven housing wealth changes. The literature typically finds MPCs for the average borrower to be between 4 and 9 cents per dollar.⁷ High-leverage borrowers appear to be even more responsive. Mian et al. (2013) find that that households living in zip codes with average LTV ratios above 90 have MPCs that are twice that of the median household, translating into an MPC of 18 cents. This result has been used by policy advocates arguing in favor of principal reduction for underwater borrowers.⁸ However, our evidence suggests that borrowers who are far underwater have much lower MPCs out of housing wealth, rendering principal reductions for these borrowers ineffective.

We argue that the inability of underwater borrowers to monetize the wealth gains from principal reduction can explain why they are far less sensitive to housing wealth changes than borrowers in other economic conditions. Typically, housing wealth gains expand borrowers' credit access. Mian and Sufi (2014a) present evidence that equity withdrawal through increased borrowing can account for the entire effect of housing wealth on spending between 2002 and 2006. But if homeowners face a collateral constraint, principal reduction that still leaves borrowers underwater or nearly underwater will not immediately relax this constraint, which may explain why principal reduction is ineffective.⁹ On the other hand, forward-looking agents who realize their borrowing constraints will be relaxed in the future could theoretically respond by reducing their savings in the present. We calculate that since borrowers remained underwater and collateral constraints had tightened, it would take eight years before the average principal reduction recipient in HAMP would be able to increase borrowing as a result of these principal reductions. A dynamic model of household optimization is useful for understanding whether such a lengthy delay can indeed explain why borrowers did not increase consumption.

To more formally explore the mechanisms driving borrower insensitivity to principal reductions, we build a partial equilibrium life-cycle model of household consumption and default decisions. The model includes uninsurable income risk, a life-cycle savings motive, and a housing asset financed

⁷See Campbell and Cocco (2007), Mian et al. (2013), Carroll et al. (2011), Case et al. (2013), Attanasio et al. (2009) and Berger et al. (2016). Several of these papers estimate elasticities rather than MPCs. We translate elasticities into MPCs by multiplying them by the ratio of census retail sales to household housing assets from the Flow of Funds in 2012, which is 0.25. We use retail sales to be comparable to Mian et al. (2013).

⁸See discussion in Goldfarb (2012).

⁹Indeed, DeFusco (2016) shows that a significant fraction of the additional borrowing arising from house price gains is due to relaxing collateral constraints.

by a modifiable mortgage. Homeowners are able to borrow against their housing wealth subject to a collateral constraint and have the option to default. To parallel our empirical results which focus on near-term default, we model a one-time default decision. Our model is able to match the empirical estimates of large consumption responses to housing wealth gains for borrowers with positive home equity.

The model clarifies that borrowers' short-term constraints govern their response to long-term obligations. In our model, for a given LTV ratio, there is an income level below which households will find it optimal to default. When income is below this threshold, the gain from increasing disposable income outweighs the utility cost of defaulting. We show that when default imposes utility costs in the short-term, default decisions are insensitive to future debt levels until borrowers are substantially underwater.¹⁰ Instead, borrowers default only when they are hit with substantial income shocks affecting their short-term budget constraints, which is consistent with our empirical evidence.

Default behavior in our model of HAMP recipients mirrors the “double trigger” class of models, whereby agents default when they are underwater and face negative income shocks. In our model this is driven by the combination of a utility cost of default, which prevents most borrowers from defaulting, and an income process with a thick left tail (Güvenen et al. 2014), which pushes some borrowers over the default threshold. Although the canonical double-trigger model does not feature agent optimization, our model contributes to a recent literature in which mortgage default and insensitivity to LTV can co-exist in an optimizing framework (Schelkle 2016, Campbell and Cocco 2015).

The main intuition for the consumption results in our model is that the collateral constraint renders principal reduction below the constraint ineffective. We develop a sufficient statistic formulation that decomposes the effect of debt reductions on consumption into a future wealth effect and a future collateral effect. We show that for low-wealth borrowers, debt reductions that cannot be monetized in the near-term do not relax current constraints and have little precautionary value, hence they do not translate into increased consumption. In fact, the consumption function out of home equity is S-shaped, convex below the collateral constraint and concave above it. Borrowers are insensitive to debt reductions until such reductions bring them close to their constraint, explaining the failure of principal reductions in HAMP to stimulate consumption. We show that for principal reduction to effectively increase underwater borrowers' consumption, collateral constraints would need to be substantially relaxed or households would need to expect rapid home price growth. However, neither of these conditions is likely to hold in the period following a financial crisis, which was when principal reduction was implemented.

The main theoretical result from our model is that the inability to monetize housing wealth

¹⁰The assumption that default is accompanied by short-term moral or social stigma is supported by survey evidence in Guiso et al. (2013), who find that about 80 percent of homeowners consider it morally wrong to default when payments are affordable. The result that borrowers without income shocks do not exercise their default option until substantially underwater is consistent with empirical evidence in Bhutta et al. (2011), who show that the median homeowner without an income shock does not default until their LTV is greater than 167.

creates a wedge between an underwater borrower’s marginal propensity to consume out of cash and their marginal propensity to consume out of housing wealth. As in the empirical results of Mian et al. (2013), borrowers near their collateral constraint have a high MPC out of housing wealth gains. However, borrowers far underwater are unresponsive. One policy implication is that high leverage is a bad “tag” for targeting policies that increase housing wealth, even though it is a good “tag” for targeting policies trying to provide cash to borrowers with high MPCs. Our results help explain why policies to lower current mortgage payments were more effective than principal reductions at stemming foreclosures and increasing demand during the Great Recession.

This paper contributes to several additional strands of the literature. Agarwal et al. (2016a) use a variety of research designs to study the overall effect of HAMP modifications which combine both short-term and long-term payment reductions, finding that the program was associated with reduced foreclosures and an increase in durable spending. We complement their findings by separating the effect of long-term from short-term payment reductions. If the effect of long-term payment reductions in HAMP is zero, as suggested by our estimates, it makes sense to infer that short-term payment reductions are responsible for the default and consumption impacts they estimate. The joint lesson from our papers supports the policy recommendation in Eberly and Krishnamurthy (2014) that government resources should be used to support household liquidity rather than reduce future debt burdens. Scharlemann and Shore (2016) analyze principal reduction in HAMP using a regression kink design and find an equivalent amount of forgiveness reduces delinquency by 1.9 percentage points, a result that is within our confidence interval. Relative to Scharlemann and Shore (2016), our contributions are the introduction of a regression discontinuity design which requires weaker parametric assumptions, an examination of the impact of principal reduction on consumption, and the development of a model which clarifies the mechanisms causing principal reductions to be ineffective.¹¹

Our model also contributes to a new literature finding little direct linkage between debt levels and consumption when debt is modeled as a long-term contract (Kaplan et al. 2016a, Justiniano et al. 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 2016). In our model, as in practice, nothing forces borrowers to immediately delever when they are far underwater, removing a mechanical link between debt levels and consumption present in some of the prior literature and reducing the expected effectiveness of debt reduction policies. Our theoretical results also relate to work by Berger et al. (2016), who develop a sufficient statistic formula showing that the consumption response to housing wealth gains depends on current house values and the marginal propensity to consume out of temporary income. Our formula adds that the response will depend crucially on a borrower’s initial home equity position.

¹¹While Scharlemann and Shore (2016) is the only other paper we know of to isolate the effect of future housing payments, Dobbie and Song (2016) analyze future payment reductions for credit card borrowers. They find that reducing payments four years in the future by 8% of the total debt owed leads to a reduction in 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. Dobbie and Song (2016) point out that the cost of defaulting on mortgage debt may be larger than the cost of defaulting on unsecured credit card debt which can be discharged in bankruptcy.

Finally, our theoretical results demonstrate a setting where policy interventions that alter the economic environment far in the future will have a limited impact in the present. This idea has been recently studied by McKay et al. (2016), who show that households facing income risk and building buffer stocks will be unresponsive to interest-rate changes far in the future because responding requires depleting valuable savings in the present.

The remainder of the paper is organized as follows. Section 2 describes the HAMP program and discusses assignment to the two modification types. Section 3 provides details on the data sources. Section 4 discusses the regression discontinuity empirical strategy and presents the results on default. Section 5 describes the difference-in-differences empirical strategy and presents the results on consumption. Section 6 builds the model and uses it to interpret our results and draw broader lessons. The last section concludes.

2 Home Affordable Modification Program (HAMP)

Two features of the government’s HAMP program make it a useful setting for studying the impact of principal reduction. First, borrowers were assigned into two distinct modification types, with both types receiving identical short-term payment reductions but one type receiving additional principal reduction. Second, borrowers were assigned according to rules that generate quasi-experimental variation, which we exploit in our empirical strategies. We discuss each of these features in turn.

2.1 Two Modification Types in HAMP

The government instituted the HAMP program in 2009 as a response to the foreclosure crisis. It provides government incentives to help facilitate mortgage modifications for borrowers struggling to make their payments. In total, 1.6 million borrowers have received modifications since the program began.

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% 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.¹²

The primary goal of HAMP modifications is to provide borrowers with more affordable mortgages. All borrowers who receive modifications have their payment reduced to 31% of their income for at least five years. This results in substantial modifications. The mean payment reduction is \$680 per month, or 38% of the borrower’s prior monthly payment.

Our research design relies on contrasting borrowers assigned to two distinct modification types.

¹²In addition, most modifications require that the mortgages be non-jumbo first liens and that the property be single-family and owner-occupied. The liquid asset test requires that liquid borrower assets are less than three times their total monthly debt payments.

Both modification types result in the same new 31% monthly payment 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. The top panel of Figure 1 shows the average annual payments for borrowers in this modification type relative to their payments under the status quo.¹³ The short-term payment reduction is achieved by twisting the payment schedule. Up to three steps are followed to achieve the 31% target. First, the interest rate is reduced down to a floor of 2% 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 term 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.¹⁴

The second modification type is what we call a “payment and principal reduction” modification. The first step in this modification is to forgive a borrower’s unpaid principal balance until the new monthly payment achieves the 31% of income target, or their LTV ratio hits 115%, whichever comes first.¹⁵ If the borrower’s monthly payment is still above 31% of their income after this principal forgiveness, then the interest rate reduction, term extension, and principal forbearance steps described above are followed as needed. To date, 218,000 borrowers have received these modifications.

The government introduced these principal reduction modifications in 2010 in response to growing concern that debt levels, rather than just debt payments, were responsible for high default rates and depressed consumption. For default, the concern was that a large share of defaults were “strategic” in nature, committed by borrowers who had negative equity rather than by borrowers with affordability problems (Zingales 2010, Winkler 2010, Experian and Wyman 2009). For consumption, an emerging body of evidence suggested household leverage as a primary explanation for the decline in aggregate demand during the recession (Mian and Sufi 2010).

The government, concerned by high debt levels, was willing to devote substantial resources towards supporting principal reduction modifications.¹⁶ It compensated investors for a portion of the principal reduced, with a sliding scale rising up to \$0.63 per dollar of principal reduced. On average, the government paid an additional \$21,000 per modification to support modifications with principal reduction (Scharlemann and Shore 2016).

By comparing borrowers who received 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 borrower payments in the short term, but dramatically

¹³Appendix Figure A.1 shows both the pre and post-modification payment schedule.

¹⁴Servicers had some flexibility regarding the order they applied these adjustments in order to comply with their servicing contracts, but generally followed the order above.

¹⁵In practice, servicers could vary this target ratio. About 85% of servicers used 115%, while 15% of servicers used a 100% LTV target (Scharlemann and Shore 2016)

¹⁶An additional motivation for pursuing principal reductions is that there may have been a political constraint on reducing short term payments below the 31% level, which was the standard for an “affordable” mortgage. Principal reduction is a way to deliver additional assistance to borrowers without needing to reduce payments further in the short-term.

different effects on long-term payments and homeowner equity. The top panel 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. The bottom panel of Figure 1 summarizes the financial impacts of these modifications. Borrowers with principal reduction modifications received an average of \$70,000 more principal reduction,¹⁷ and the reduction in the net-present-value of the payments owed under the new mortgage contract was \$35,000 greater.¹⁸

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 Appendix Figure A.2. Borrowers appear eager to take up modifications. Conditional on being offered a modification, 97% of borrowers accepted the offer.

2.2 Quasi-Experimental Variation in Borrower Assignment to Principal Reduction

Borrower assignment was determined in part by a running variable with a distinct cutoff, and in part by servicer and investor type. This assignment generates quasi-experimental variation in which modifications borrowers received, which we will exploit in our empirical strategies.

The first factor affecting borrower assignment is a calculation examining which modification type is expected to be most beneficial for the investor. Using a model developed by the U.S. Treasury Department, servicers calculate the expected net present value (NPV) of cash flows for investors under the status quo and under each of the two modification options. The NPV model takes into consideration government-provided incentives as well as the expected impact that modifications will have on the likelihood that borrowers repay their debt. Agarwal et al. (2016a) use the NPV calculation in one of their empirical strategies for comparing borrowers who received HAMP modifications to those that did not.

Our regression discontinuity empirical strategy exploits a large jump in the share of borrowers receiving modifications with principal reductions when the NPV model shows it will be more beneficial to investors than the alternative. This jump is shown in Figure 2. The purpose of the calculation is to reduce contracting frictions between investors and servicers. Servicers are bound by their fiduciary duty to the investors to maximize repayment, and as a result are more likely to

¹⁷Some 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% LTV and required the forgiveness to vest over three years.

¹⁸The net-present-value calculation assumes a discount rate of 4%, equal to the average market mortgage interest rate over the program period, which is the rate used to discount cash flows in the model described in Section 2.2. In practice, a portion of the borrowers defaulted before realizing the future payment reductions. Those who end up in foreclosure, short sale, or chapter 7 bankruptcy may never owe these payments. If lifetime 90-day default rates are 57% as predicted by the Treasury default model described in section 4.2, and half of the 90-day defaulters end up in permanent non-payment as found in Treasury Department (2014c), then the expected reduction in the net-present-value of payments owed would be \$26,000 greater for borrowers with principal reduction modifications than for borrowers with payment reduction modifications.

offer the modifications shown to be most beneficial to investors.

The second factor affecting assignment is servicer and investor type. Borrowers are not assigned according to the NPV calculation alone because different investors 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. The contractual frictions are particularly acute with securitized loans. For example, Kruger (2016) shows that 22% 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-investor frictions are mitigated (Agarwal et al. 2011). Conditional on investor 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 (Treasury Department 2014b).

3 Data

To estimate the effect of principal reduction in HAMP, we need a dataset that includes information on borrower characteristics, HAMP participation, and relevant borrower outcomes. We build two separate datasets, each with distinct strengths along the necessary dimensions. Our first dataset matches detailed HAMP participation data to consumer credit bureau records. This dataset covers a broad sample population and allows us to carefully exploit the mechanism assigning borrowers to each modification type, which we exploit in our regression discontinuity analysis. Our second dataset comes from a bank that is also a mortgage servicer. This dataset has a smaller sample size and less information about borrower assignment, but it includes richer outcome measures such as income as measured from checking account transactions.

3.1 Matched HAMP Credit Bureau File

The U.S. Treasury releases a public data file on the universe of HAMP applicants. This loan-level dataset includes information on borrower characteristics and mortgage terms before and after modification. Crucially, it also includes the NPV calculations run by servicers when evaluating borrowers for each modification type.

In order to observe a wide range of borrower outcomes, we use consumer credit bureau records from Transunion. HAMP program rules require servicers to report borrower participation to credit bureaus. We use the universe of records for borrowers flagged as having received HAMP. We have monthly account-level information between January 2010 and December 2014 for each borrower.

We develop proxies for both durable and nondurable consumption based on the credit bureau records. For durable consumption, we follow Keys et al. (2014) and Di Maggio et al. (2016) by using changes in auto loan balances as a measure of car purchases. Di Maggio et al. (2016) document that leveraged car purchases account for 80 percent of total car sales. While prior work relied on

observing jumps in total auto loan balances to infer new loans, our product account-level data allows us to observe new loans directly. The detailed nature of our credit bureau data also allow 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.¹⁹ We are able to construct this measure for 83% of all credit and charge card accounts.²⁰ We find average credit card spending of \$425 per month in our sample, which is 80% of the average credit card spending per adult in 2012 (Federal Reserve 2014), commensurate with the 83% 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. The attributes we use are MSA, modification month, origination year, loan balance, and monthly payment before and after modification. We are able to match half of the records in our sample window, resulting in a panel dataset of about 112,000 underwater borrowers.

Our match rate is less than 100% 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 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 (Treasury Department (2014a)).

The imperfect match rate does not bias our sample along any observed borrower characteristics, and we can replicate our main empirical results in alternative datasets that do not rely on matching. 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 Appendix Figure A.3. In section 5.2 we show that our consumption result is unchanged (though slightly less precise) when we estimate it using the bank dataset described in the following section.

3.2 Bank dataset

Our second dataset uses de-identified data assembled by JPMCI, and complements our first dataset in three ways. First, it includes information on modifications and borrower outcomes from the same source, reducing any measurement error from imperfect matching. Second, it includes

¹⁹Let 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 $e_t = b_t - b_{t-1} + p_t$. 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 3.2 we can isolate purchases and confirm that our results are unchanged.

²⁰The balance of accounts are for cards with servicers who do not report monthly payment data to the credit bureaus.

alternative consumption data. Third, it includes other borrower outcomes such as income, mortgage payments, and monthly credit scores.

The JPMCI dataset includes account-level monthly information on mortgages, credit cards, and checking accounts. The dataset is described in more detail in Ganong and Noel (2016). We focus on two subsamples of HAMP borrowers in the JPMCI dataset. The first includes HAMP borrowers with both a mortgage and a checking account with Chase. The combined dataset is available from October 2012 through August 2016. We build a measure of checking account income as all inflows into a customer checking account in a given month. This combines labor and capital income, government support, electronic transfers, and paper check deposits. We use this sample to examine the path of income for borrowers who default after receiving a modification. We observe 7,224 borrowers in this dataset within one year of redefaulting.

The second sample includes HAMP borrowers with both a mortgage and a credit card with Chase. This combined dataset is available from April 2009 through August 2016. We observe credit card spending and credit scores of 22,924 borrowers one year before and after modification.

4 Default

In this section we analyze the effect of principal reduction on borrower default. Using a regression discontinuity 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. Although we find that default is unresponsive to long-term debt obligations, we provide evidence that it is responsive to short-term cash-flow shocks.

4.1 Representativeness of HAMP Participants Relative to Typical Delinquent Underwater Borrowers

Our empirical analysis, described in the next section, focuses on borrowers near the assignment cutoff for receiving principal reduction. To assess the representativeness of our analysis sample, we compare borrowers near the cutoff in the matched HAMP credit bureau file to a sample of delinquent borrowers in the Panel Study of Income Dynamics (PSID) between 2009 and 2011. Summary statistics for borrowers in both samples are shown in Table 2. Borrowers in our sample are broadly representative of delinquent underwater borrowers during the recent crisis.

The median borrower in our sample has a higher LTV than delinquent borrowers in the PSID (119 compared to 94), but the 90th percentile LTV is similar (163 compared to 166). Since all the borrowers who are evaluated for principal reduction must be underwater, we would expect them to be concentrated in the underwater portion of the delinquent borrower distribution. The fact that borrowers in our 90th percentile are “only” at an LTV of 163, and that the median borrower is substantially less underwater, will be important for interpreting our empirical results. Prior evidence from Bhutta et al. (2011) suggests borrowers most sensitive to debt levels are even further underwater, which will be consistent with our findings.

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

Our sample also covers the period with the most severe delinquency rates in the recent crisis. Appendix Figure A.4 plots the delinquency rate for all U.S. borrowers over time. Our sample of borrowers have their first delinquencies in the fourth quarter of 2009, just before the peak of the delinquency crisis, which did not begin abating until 2013.

4.2 Regression Discontinuity Empirical Strategy

Our strategy to identify the effect of principal reduction on default exploits cutoffs in the expected benefit to investors using a regression discontinuity research design. Let the receipt of principal reduction treatment be denoted by the dummy variable $T \in \{0, 1\}$, where 0 represents receiving a payment reduction modification, and X capture the characteristics of the borrower. The Treasury NPV model calculated the expected net present value to investors $ENPV(T, X)$ under either scenario. Our running variable V is the normalized predicted gain to investors of providing principal reduction to borrowers, i.e.

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

A realization v reflects the expected percent gain to the investor from principal reduction relative to a standard modification. The cutoff that affects assignment to treatment or control is at $v = 0$. We normalize the predicted gain to avoid a high concentration of low-balance mortgages near the cutoff, though our key finding that principal reduction does not affect default is insensitive to this normalization.

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 regression discontinuity (RD) estimand is given by

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

The parameter τ identifies the local average treatment effect of providing principal reduction to borrowers near the cutoff. These are borrowers for whom the increased government incentive payments and the reduction in default expected by the Treasury NPV model generate gains to investors that offset investors' cash-flow loss from forgiving principal. We discuss Treasury's estimate of the expected default reduction further in section 4.5.

We follow the standard advice for Regression Discontinuity Designs from Athey and Imbens

(2016) to estimate $\hat{\tau}$, which is numerically equivalent to an instrumental variables (IV) estimator. We pre-process the data, restricting the analysis sample to the set of borrowers where $|v| < 2$. We use a local linear regression to estimate the limits in equation 2. Our preferred specification uses the recommended bandwidth selection procedure from Imbens and Kalyanaraman (2012), which yields $v = 0.5$. Following Imbens and Kalyanaraman (2012), we use the triangle kernel in our primary specification and a uniform kernel as a robustness check. Our results are robust to using the procedures recommended by Calonico et al. (2014). Our analysis dataset is the matched HAMP credit bureau dataset, which includes the predicted gain to investors of providing principal reduction v .

4.3 Internal Validity of Design

The validity of a fuzzy regression discontinuity research design rests on two assumptions. First, there is a discontinuous jump in treatment probability at the cutoff. Second, all other variables that could affect outcomes are continuous at the cutoff.

The top panel of Figure 2 validates the first assumption. It shows that there is a discontinuous jump of 41 percentage points in the share of borrowers receiving principal reduction at the cutoff. The bottom panel of Figure 2 depicts the treatment in terms of dollars of principal reduction. It shows that the treatment size at the cutoff is \$34,000. This is half as large as the cross-sectional difference in principal reduction between all borrowers considered for both programs.²¹ Appendix Figure A.5 shows that the treatment reduces borrower LTV by 11 percentage points, and amounts to a \$20,000 reduction in the NPV of borrower payments owed over the full mortgage term. The bottom panel of Appendix Figure A.5 shows that 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.

Figure 3 provides evidence supporting the second assumption. The first five panels show the distribution of pre-modification borrower credit score, monthly income, months past due, monthly mortgage payments to monthly income (debt-to-income, or DTI) ratio, and LTV ratio around the cutoff. In all cases these borrower characteristics trend smoothly. The RD estimates of the discontinuous change in these variables at the cutoff, corresponding to the numerator of equation 2, are reported on the figures. For three variables (credit score, monthly income, and DTI) the sign points to slightly worse-off borrowers to the right of the cutoff, while for two variables (LTV and months past due) the sign points to better-off borrowers to the right of the cutoff. The lack of any systematic correlation supports the validity of the design. The only covariate with a marginally statistically significant jump is months past due at application date, and even here the jump is not economically significant.²² Lee and Lemieux (2010) note that when there are many covariates,

²¹This research design focuses on borrowers where the NPV model says that an investor should be nearly indifferent to whether principal reduction is pursued. Indifference is more likely when the quantity of principal reduction prescribed by the program rules is moderate and so treatment at the cutoff is smaller than the average treatment.

²²Pre-modification months past due is hardly predictive of post-modification default. Using the cross-sectional relationship between the two, we find that a jump of 0.5 months in pre-modification months past due is associated with a 0.2 percentage point lower probability of re-default.

some discontinuities will be significant by random chance. They recommend combining the multiple tests into a single test statistic. We implement a version of this by using all five pre-modification covariates to predict default, and we test whether there is a jump in this pooled predicted default measure at the cutoff. The result is shown in the last panel of Figure 3. There is no significant change in predicted default at the cutoff.

Another relevant issue in regression discontinuity settings is the possibility that the running variable could be manipulated (McCrary 2008). The usual test is to plot a histogram of the running variable to examine whether there is an unusual increase in mass to the right of the cutoff. We show such a plot in the top panel of Appendix Figure A.6. While the density is smooth on either side of the cutoff, there is a large bulge exactly at zero.

There are two reasons why we believe the bunching of borrowers at zero is not a challenge for the validity of our research design. First, program officers in charge of the dataset at the U.S. Treasury Department told us that this bulge is a data artifact. They believe several servicers ran only one NPV calculation and reported this single number as the calculation for both payment reduction and principal reduction modifications, meaning that they reported $ENPV(1, X) = ENPV(0, X)$. We were advised by U.S. Treasury staff to remove these observations as reflecting measurement error. Second, the conventional economic environment that would incentivize manipulation is not relevant here. Servicers have no economic incentive to manipulate the running variable because they receive the same compensation regardless of which modification is offered.²³

We attribute the bunching of borrowers at zero to data mis-reporting and drop observations exactly at zero. The bottom panel of Appendix Figure A.6 shows the distribution for the resulting sample, which is our analysis sample. There is no noticeable change in density around the cutoff. We show in Appendix Figure A.7 that borrower takeup rates were high on both sides of the discontinuity. Ninety-seven percent of borrowers who are offered a modification take it up, and this trends smoothly around the cutoff. This provides further evidence against borrower manipulation to obtain one or the other modification type.

4.4 Effect of Principal Reduction on Default

We find that principal reduction has no impact on default. The top panel of Figure 4 shows the reduced form of the fuzzy regression discontinuity specification, plotting the default rate against the running variable. Default is defined as being 90 days delinquent at any point between modification date and March 2015, when our dataset ends, which is an average of three years. There is no jump in default rates at the cutoff. Our point estimate is that \$35,000 of principal reduction changes default probabilities by less than one tenth of a percentage point, and we can rule out a reduction of more than five percentage points.

²³Two additional arguments support our claim that the bulge is not a problem. First, even if servicers did have an economic incentive to manipulate, that incentive would not vary discontinuously at this cutoff: principal reduction provision is optional regardless of the outcome of the calculation. Second, were servicers manipulating the running variable to zero in an attempt to rationalize principal reduction, they failed; the share of borrowers receiving principal reduction in this zero group is actually half what it is for borrowers with actual positive values of the running variable.

Our results imply a large government cost per avoided foreclosure. While our data does not follow individual borrowers to foreclosure, the government has published reports showing that 25% of HAMP borrowers who go 90 days delinquent end up in foreclosure (Treasury Department 2014c). Thus, even taking the upper bound of our confidence interval that default was reduced by 5 percentage points, this translates into a 1.25 percentage point reduction in foreclosure. The government spent about \$10,000 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 \$800,000 per avoided foreclosure, more than an order of magnitude larger than common estimates of the social costs of foreclosure (Hembre 2014).

The bottom panel of Figure 4 tests the sensitivity of our results to the bandwidth chosen for the local linear regression. Our central estimates are constructed using the optimal bandwidth from the Imbens and Kalyanaraman (2012) procedure, which is $|v| < 0.5$. The optimal bandwidth recommended by the Calonico et al. (2014) procedure is 0.4. The point estimate is stable out to a bandwidth of 0.7 and then begins to rise. The rise at wider bandwidths is not surprising given the shape of the estimated conditional expectation function for default, which is particularly sloped near the cutoff. Wider bandwidths will lead to specification error when this function is particularly steep near the cutoff. We show in Appendix Figure A.8 that a quadratic specification which can more easily mimic this slope is stable for a wider bandwidth, showing a point estimate around zero up to a bandwidth of 1.3 before rising.²⁴

4.5 Comparison to Prior Empirical Evidence

Our results are inconsistent with prior evidence based on cross-sectional relationships. For example, Haughwout et al. (2016) use data on modifications performed prior to HAMP and find using cross-sectional variation that borrowers who received principal reductions equivalent to ours saw an 18 percentage point reduction in default. Furthermore, there is a strong cross-sectional relationship between the amount of negative equity and default rates across all borrowers (Gerardi et al. 2015).

The U.S. Treasury Department generated its own estimate of the impact of principal reduction on default as part of its model to predict the benefits of modifications to investors. The Treasury redefault model is based on a logistic regression framework, with coefficients estimated to fit default rates using historical data (Holden et al. 2012).

The Treasury model predicted a substantial reduction in default from principal reduction, which is inconsistent with our findings. We implemented the Treasury redefault model in the public Treasury data and calculated the predicted impact of principal reduction at the cutoff. The Treasury model expected a sharp reduction in default of 6.6 percentage points at the cutoff. The black horizontal line in the bottom panel of Figure 4 compares this predicted impact to the results from our various specifications. One nuance of comparing our results to the Treasury model is that the

²⁴The optimal bandwidths for quadratic specifications using the Imbens and Kalyanaraman (2012) and Calonico et al. (2014) procedures are 0.8 and 1.0, respectively.

model is designed to predict lifetime default rates whereas we only observe borrowers an average of three years after modification. If we assume the impact of debt forgiveness should be proportional over time rather than apparent within the first few years, then the appropriate comparison is an odds ratio. Treasury’s predicted odds ratio is 1.11, which is above the 80th percentile of our confidence interval and well above our point estimate of 1.0.

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% of the difference in default rates between the 2003-2004 and the 2006-2007 cohorts.²⁵

In contrast to cross-sectional estimates, Scharlemann and Shore (2016) analyze principal reduction in HAMP using a regression kink design, and find that principal reduction of the size we estimate (reducing the LTV by 11 points) reduces default by 1.9 percentage points, a result that is within our confidence interval. Our empirical results reinforce the conclusions of Scharlemann and Shore (2016) using an alternative estimation strategy that requires weaker parametric assumptions. In Scharlemann and Shore (2016)’s preferred specification, they use a regression kink design and assume that the relationship between the running variable and the outcome variable is globally linear. In contrast, our regression discontinuity identification strategy assumes that unobservable borrower traits are smooth in the local region around the cutoff.

4.6 Sensitivity of Default to Short-Term Cash-Flow Shocks

While we find no causal impact of long-term debt obligations on default decisions, we provide suggestive evidence that default is responsive to cash-flow shocks that tighten short-term budget constraints. For this analysis, we focus on the combined checking and mortgage account dataset from JPMCI which is available for a subset of HAMP borrowers. Figure 5 plots the event study of mortgage payments and checking account income around the date at which borrowers who had received HAMP modifications become 90 days delinquent. Missed mortgage payments coincide with a sharp decline in income, indicating short-term financial stress. The loss in income in the four months before default can account for the entire value of missed mortgage payments. This high-frequency evidence is consistent with the conclusion in Gerardi et al. (2015), who analyze biennial data in the PSID and find that affordability shocks such as unemployment or large income losses are the strongest predictors of mortgage default. Figure 5 also highlights the temporary nature

²⁵Palmer (2015) shows that the remaining variation is attributable to differential exposure to local home price declines. Our empirical and theoretical results are consistent with this finding, in the sense that defaults are higher on average for underwater than above water borrowers, and there is a point at which the degree of negative equity has a causal impact on default rates. The model we present in section 6 shows that default rates jump once borrowers go from positive to negative equity, are flat for moderate levels of underwaterness (consistent with our empirical estimates), and then rise again once borrowers are significantly underwater in a region where defaults occur for strategic reasons even absent negative income shocks.

of the income shocks associated with defaults on average. One year after a 90-day delinquency, average income has recovered to its pre-default mean. This explains why average payments bounce back after the first 90-day delinquency, as some borrowers are able to self-cure.

Our HAMP credit bureau data contain two further pieces of suggestive evidence on the importance of short-term cash flow for default decisions. First, on the intensive margin, the cross-sectional relationship between payment reductions and default rates is strong. The top panel of Appendix Figure A.9 shows that borrowers who receive larger reductions in debt-to-income are substantially less likely to re-default. Second, on the extensive margin, millions of borrowers who are delinquent when they apply to HAMP start making their payments again after receiving payment reductions. The bottom panel of Appendix Figure A.9 shows the event study of default around modification date for borrowers in our sample. Before modification, 70 percent of borrowers are at least 90 days delinquent. However, once borrowers receive lower payments, delinquency rates fall for borrowers receiving both types of modifications, and most continue making their payments throughout the first year.

Other recent empirical evidence also points to the importance of short-term budget constraints. Reductions in borrower mortgage payments are consistently shown to reduce default rates. Fuster and Willen (2015) study payment reductions resulting from adjustable rate mortgage resets, and find that underwater borrowers are highly sensitive to payment reductions. Agarwal et al. (2016a) find that HAMP modifications, which combine both short-term and long-term reductions in mortgage payments, result in substantially lower foreclosure rates. If the effect of long-term payment reductions in HAMP is zero, as suggested by our estimates, it makes sense to infer that short-term payment reductions are responsible for default impact they estimate.

After we develop our model in section 6, we will use it to further analyze the relationship between short-term affordability and long-term debt obligations. We will show that when defaulting imposes utility costs in the short-term, borrower default is driven by cash-flow shocks such as unemployment which tighten borrowers' short-term budget constraints. In this case, default decisions are insensitive to future debt levels until borrowers are substantially underwater. This can explain our empirical findings, since only a small fraction of our borrowers have initial LTVs above 160, the point at which default becomes sensitive to mortgage principal in our model.

5 Consumption

In this section we explore the effect of principal reductions on consumption. We find that the marginal propensity to consume out of housing wealth for underwater borrowers is an order of magnitude smaller than the marginal propensity to consume for average homeowners examined in prior studies.

5.1 Panel Difference-in-Difference Empirical Strategy

Our use of consumption data motivates a change in research design to a panel difference-in-differences strategy. There are two reasons for this change. First, our regression discontinuity strat-

egy is under-powered for studying changes in consumption. Economically meaningful consumption changes can not be ruled out in the regression discontinuity sample given lack of precision. As we discuss in more detail in section 5.3, even a small change in consumption on the order of 5% would be meaningful relative to average marginal propensities to consume out of housing wealth changes studied in other contexts, whereas the predicted impacts on default from the prior literature were much larger. The second reason is that the panel nature of the spending measures from our credit bureau and banking data allow us to exploit an alternative strategy that offers more precision. Lagged spending measures allow us to adjust for underlying differences between borrowers receiving different modification types within a wider bandwidth than with the regression discontinuity. These factors favor a panel difference-in-differences design, though we show that results from the regression discontinuity strategy are consistent with our main findings.²⁶

Our panel difference-in-differences design uses as a control group the set of underwater borrowers who were eligible for principal reductions, but who instead received only payment reduction modifications. This design relies on the fact that borrowers who receive payment reduction modifications experience the same short-term payment reductions as borrowers who receive principal reduction, but receive substantially less generous long-term payment relief. The “treatment” is the effect of long-term debt forgiveness holding short-term payments fixed.

Our identification comes from cross-servicer and cross-investor variation in the propensity to provide principal reductions given observed borrower characteristics. We control for the expected gain to investors of providing principal reduction modifications. Thus, the main difference between our regression discontinuity and difference-in-differences strategies is that the regression discontinuity 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. Intuitively, this research design relies on servicer-investors that were more likely to offer principal reduction not having borrowers whose consumption was trending upward relative to borrowers whose servicer-investors 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 Figures 6 and 7. The top panel of Figure 6 plots mean credit card expenditure around modification date, and the bottom panel normalizes expenditure to zero at modification date in order to more clearly show the parallel pre-trends. Figure 7 plots mean auto expenditure around modification date. Visual inspection of Figures 6 and 7 indicates that principal reduction appears to have little effect, a result we explore in a regression framework.

²⁶We also have lagged measures of default from the credit bureau data. However, a difference-in-differences design is not valid for default because pre-treatment differences in the levels of default are mechanically removed at modification date, at which point all loans become current. This means that the change in default for the control group is not a valid counterfactual for the change in the treatment group. The event study of default around modification is shown in the bottom panel of Figure A.9

Formally, our main specification is the following:

$$y_{i,g,\tau,t} = \gamma_g + \gamma_\tau + \gamma_{m(i),t} + \beta (\text{Principal Reduction}_g \times \text{Post}_\tau) + x'_{i\tau} \delta + \varepsilon_{i,g,\tau,t}, \quad (3)$$

where i denotes borrowers, $g \in \{\text{payment reduction, payment \& principal reduction}\}$ the modification group, τ the number of months since modification, t the calendar month, and m the household's Metropolitan Statistical Area (MSA). Our main outcome variables $y_{i,g,\tau,t}$ are monthly credit card and auto expenditure, which proxy for non-durable and durable spending, respectively. γ_g captures the modification group fixed effect and γ_τ captures a fixed effect for each month relative to modification. *Principal Reduction* $_g$ is a dummy variable equal to 1 for the group receiving modifications with principal reduction while *Post* $_\tau$ is a dummy variable equal to 1 for $\tau \geq 0$. The main coefficient of interest is β , which captures the difference-in-differences effect of principal reduction. One potential concern is that different geographies were experiencing different trends in their house price recoveries, which affected borrower outcomes. To address this concern $\gamma_{m(i),t}$ captures MSA-by-calendar-month fixed effects. x_i is a vector of individual characteristics designed to capture any residual heterogeneity between treatment and control groups. These characteristics x_i are interacted with the *Post* $_\tau$ variable to allow for borrower characteristics to explain changes in underlying trends after modification ($x'_{i\tau} = (x_i \quad x_i \times \text{Post}_\tau)'$). This includes the predicted gain to investors from providing principal reduction, the predicted gain interacted with a dummy variable equal to one when the gain is positive, borrower characteristics (credit score, monthly income, non-housing monthly debt payment), pre-modification loan characteristics (LTV, principal balance, DTI, monthly payment), property value, origination LTV, and monthly payment reduction. Standard errors are clustered at the borrower level.

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. Table 3 reports summary statistics for this sample. Both groups of borrowers are broadly similar on a range of pre-modification characteristics, including LTV, DTI, monthly mortgage payment, income, and credit score.

Table 3 also shows the size of the treatment. While the size of short-term payment reductions are nearly identical for borrowers in both groups, borrowers who receive payment and principal reduction modifications receive on average \$70,000 more principal reduction, reducing the NPV of the payments owed under their mortgage contract by an additional \$35,000.

5.2 The Effect of Principal Reduction on Consumption

We find that neither credit card nor auto expenditures are affected by principal reduction. Our main results are reported in Panels A and B of Table 4. In both panels, column (1) reports the most sparse specification, while columns (2)-(6) add in additional fixed effects and controls. Across all specifications, the treatment effect of principal reduction on both monthly credit card and auto expenditure is small and statistically insignificant. Our preferred estimate using equation 3 is in column (6), which includes MSA by calendar month fixed effects and interacts control variables with

a post-modification dummy. In this specification, our point estimate is that principal reduction of \$70,000 increases borrower monthly credit card expenditure by less than \$1 and auto spending by less than \$5.

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 3.1. The second is any measurement error introduced by our matching procedure. The JPMCI dataset covers only one servicer, but does not suffer from either of these two potential limitations. It includes credit card data but not auto loan data. Appendix Figure A.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 Table 5. 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 regression discontinuity strategy. Our outcome variables are the change in mean credit card and auto spending from the twelve months before modification to the twelve months after modification. The reduced form plots are shown in Appendix Figure A.11. These plots confirm the weakness of this strategy for studying consumption impacts since the strategy suffers from lack of precision. Nevertheless, the results are consistent with no significant change in consumption at the discontinuity.

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.²⁷ The slope of expenditure changes sharply around modification, suggesting that the combination of lower payment obligations and resolved uncertainty help expenditure to recover. Another important factor is that borrower credit scores improve sharply after modification. The path of credit scores around modification is shown in Appendix Figure A.12 for borrowers in the JPMCI dataset, where credit scores are observed at a monthly frequency. Credit scores recover because borrowers who were previously delinquent become current once they make the first payment on the modified mortgage. This may also explain why auto spending recovers more dramatically than credit card spending, since the auto spending we measure depends on opening a new credit line.

The auto spending recovery from short-term payment reductions is consistent with findings in Agarwal et al. (2016a). That paper exploits regional variation in the implementation of HAMP to estimate the effects of HAMP modifications which combine both short-term and long-term payment reductions. They find that the combined modifications are associated with increased auto spending. If the effect of long-term payment reductions in HAMP is zero, as suggested by our estimates, it makes sense to infer that short-term payment reductions are responsible for the consumption impact

²⁷At the time of modification 70 percent of borrowers have missed at least three mortgage payments. Furthermore, evidence from Bernstein (2015) indicates that labor income is temporarily depressed before modification.

they estimate.

5.3 Estimated MPC From Principal Reduction Compared to Housing Wealth MPCs

To help interpret the economic significance of our results, we convert our estimate for the impact on credit card and auto consumption into a marginal propensity to consume out of housing wealth. First, we scale up borrower credit card spending to a measure of household non-auto retail spending to be comparable to Mian et al. (2013). We do this by adjusting for credit card spending on cards not reported in the credit bureau data, inflating individual spending to household level, and then multiplying by the ratio of non-auto consumer retail spending to consumer credit card spending in 2012.²⁸ 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 0.2 cents per dollar of principal reduction, with the upper bound of the 95% confidence interval corresponding to 0.8 cents. A similar procedure using results from the JPMCI bank dataset yields an upper bound of 1.2 cents. If we normalize by the NPV of reduced mortgage payments owed under the new mortgage contract rather than the dollar value of principal reduction, we get an upper bound of 1.5 cents.

Our estimated MPC out of housing wealth for underwater borrowers is an order of magnitude smaller than the marginal propensity to consume for average homeowners examined in prior studies. A set of comparison points for our estimates of the impact of principal reduction on consumption comes from the literature examining the impact of price-driven housing wealth changes. We present several such estimates in Figure 8. Campbell and Cocco (2007) and Mian et al. (2013) find annual MPCs for homeowners around 9 cents per dollar of housing wealth gain. Carroll et al. (2011) find MPCs between 2 and 9 cents. There is also a wide literature estimating the elasticity of consumption to housing wealth, with findings typically between 0.1 and 0.3 (Case et al. (2013), Attanasio et al. (2009), Berger et al. (2016)). These elasticities translate into MPCs between 3 and 8 cents per dollar.²⁹

While average MPCs typically range between 4 and 9 cents per dollar, there is evidence that high leverage borrowers are even more responsive. Mian et al. (2013) find that that households living in zip codes with average LTV ratios above 90 have MPCs that are twice that of the median household, translating into an MPC of 18 cents.³⁰ This result has been used by policy advocates

²⁸Specifically, our adjustment factor is

$$\left(\frac{\text{credit cards in Transunion}}{\text{credit cards with spending reported in Transunion}} \right) \cdot (\text{adults per household}) \cdot \left(\frac{2012 \text{ non-auto consumer retail spend}}{2012 \text{ consumer credit card spend}} \right).$$

The first term uses our data, adults per household and retail spend are from the census, and consumer credit card spending is from Federal Reserve Payment Study (Federal Reserve 2014). This gives an adjustment factor of $(1.2)(1.9)(2.5) = 5.9$.

²⁹We translate elasticities into MPCs by multiplying them by the ratio of census retail sales to household housing assets from the Flow of Funds in 2012, which is 0.25. We use retail sales to be comparable to Mian et al. (2013).

³⁰Similarly, Disney et al. (2010) analyze panel survey data in the U.K. and find that borrowers who begin the

arguing in favor of principal reduction for underwater borrowers.³¹ However, our evidence suggests that the response of borrowers who are far underwater, with LTVs far above 90, have much lower MPCs out of housing wealth, rendering principal reductions for underwater borrowers ineffective.

Why are underwater borrowers receiving principal reduction less sensitive to housing wealth changes than borrowers in other economic conditions? We argue that the key issue is access to liquidity. In the changes examined in prior research, housing wealth gains expanded borrowers' credit access. Mian and Sufi (2014a) show that equity withdrawal through increased borrowing can account for the entire effect of housing wealth on spending between 2002 and 2006. But if homeowners face a collateral constraint rather than a "natural" borrowing limit allowing them to monetize the present value of their minimum expected lifetime net worth, principal reductions that still leave borrowers underwater or nearly underwater will not immediately relax this constraint.³² Indeed, DeFusco (2016) shows that a significant fraction of the additional borrowing arising from house price gains is due to relaxing collateral constraints. If borrowers cannot monetize the wealth gained by debt forgiveness, it is not surprising that they do not respond by increasing consumption.

The possibility that liquidity can explain the lack of response is also consistent with prior research looking at large price declines. Mian et al. (2013) and Kaplan et al. (2016b) both document a non-linearity in the consumption response to house price declines. The MPC is large for small declines, but gets smaller as declines in home value get larger. Similarly, Gabriel et al. (2016) analyze the impact of California's foreclosure prevention laws and find that in the midst of already steep price declines, policies that slightly moderate the decline had no stimulative effect on auto spending. Mian et al. (2013) suggest that the non-linearity they document could be caused by smaller responses once borrowers become underwater. Similarly, our evidence suggests that for borrowers who start substantially underwater, gains in housing wealth do not affect their consumption.

5.4 Difficulty of Accessing Housing Wealth During Recovery

Even if borrowing constraints are not relaxed immediately, it is possible that forward-looking agents could respond if they believed debt forgiveness would relax their constraints in the near future. We explore this possibility theoretically in section 6. The intuition is that an agent who is building up a liquid buffer against future shocks will cut their savings when they expect future liquidity to increase. In this case, the key parameter determining the impact of debt forgiveness on consumption is the expected delay between when debt is forgiven and when households can access this wealth.

Three pieces of evidence suggest that borrowers could expect a lengthy delay before being able to access wealth from principal reductions. First, most borrowers in our sample are still underwater even after receiving principal reductions. The mean LTV ratio after modification is 122, and only 10% of borrowers have LTVs below 95. Furthermore, these leverage ratios only account for first

period underwater and end the period above water, and are thus near the 100 LTV threshold, have high MPCs of about 13 cents.

³¹See discussion in Goldfarb (2012)

³²See Carroll (1992) and Aiyagari (1994) for discussions of natural borrowing limits.

liens, while home equity depends on all liens on a property (i.e. the combined loan-to-value ratio, or CLTV).

Second, the time series of mortgage credit origination shows that credit constraints had tightened during the recovery. The top panel of Appendix Figure A.13 shows mortgage originations by borrower credit score from 2000 to 2015. This covers all mortgages, including second mortgages and home equity lines of credit (HELOCs). It shows that originations dipped sharply after 2007, and for low-credit score borrowers, originations have never recovered. Borrowers receiving HAMP principal reductions had mean FICO scores of 579, with 85% below 660, the cutoff for the red line in the figure. This evidence suggests that even with positive equity, the low credit-score borrowers in our sample may have been unlikely to obtain additional housing-related credit. This is further reinforced by the bottom panel of Appendix Figure A.13, which shows the time series of average CLTV ratios for borrowers able to obtain HELOCs in a given year. The average CLTV ratio fell 20 points between 2004 and 2009, indicating a tightening of underwriting constraints. Mian and Sufi (2014a) argue that tightening credit conditions could explain why the house price recovery from 2011 onward did not contribute significantly to economic activity, since in this case the borrowing channel is restricted. Our results support this hypothesis for underwater borrowers. Furthermore, Agarwal et al. (2016b) show that credit expansions during the recovery were more likely to benefit higher-FICO borrowers, precisely those least likely to respond by increasing borrowing.

Third, home price expectations were depressed relative to the boom years. Home price future contracts indicated a market expectation of 1% real annual home price growth between 2011 and 2016 (Department of Housing And Urban Development 2016). We show in section 6.4.3 that for the average borrower receiving principal reduction, facing a collateral constraint of 80 LTV and an expected annual price growth of 1%, it would take eight years before they would be able to increase borrowing as a result of home equity gains. We explore the implications of such a delay theoretically in the following section.

6 Partial Equilibrium Life-cycle Model with Housing

In this section, we build a partial equilibrium life-cycle model of household consumption and default decisions. We use this model to explore the mechanisms driving borrower insensitivity to principal reductions and draw broader lessons about the effects of long-term debt obligations on borrower outcomes. The model clarifies that a borrower's short-term constraints govern their response to long-term obligations.

Our model contains two key additions relative to the standard life-cycle consumption problem: a home and a default option. First, households own a home 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 the collateral constraint will cease to bind. This means that while debt reductions or house price increases raise housing wealth, households are unable to monetize this wealth in the near term. Conversely, when house prices fall and LTVs are pushed beyond the initial collateral constraint, consistent with real-world mortgage contracts we do not assume households

are forced to immediately repay excess mortgage debt. This contrasts with the treatment of debt as a one-period bond present in some of the prior theoretical literature where forced deleveraging creates a strong link between debt levels and consumption (e.g., Eggertsson and Krugman 2012, Guerrieri and Lorenzoni 2016). Second, households have an option to default on their mortgage. To parallel our empirical results which focus on near-term default, we model a one-time default decision. Borrowers receive a mortgage modification, receive an income shock, and then decide whether to stay in their home until retirement, or default and enter the rental market. Defaulting leads to a lower short-term housing payment, but agents suffer an immediate utility cost. This cost captures the moral and social stigma of default as well as moving costs.

The model clarifies that borrowers’ short-term constraints govern their response to long-term obligations. When default imposes utility costs in the short-term, default decisions are insensitive to future debt levels until borrowers are substantially underwater. Instead, borrowers default only when they are hit with substantial income shocks affecting their short-term budget constraints, consistent with our empirical evidence and with “double trigger” models of default.

The main intuition for our consumption results is that when borrowers face collateral constraints, principal reductions that leave households underwater do not relax current constraints and have little precautionary value, hence they do not translate into increased consumption. We develop a sufficient statistic formulation that decomposes the effect of debt reductions on consumption into a future wealth effect and a future collateral effect. When borrowers face lengthy delays before payments fall or borrowing limits rise, both of these effects are small. Hence the consumption function out of home equity is S-shaped, convex below the collateral constraint and concave above it. Borrowers are insensitive to debt reductions until such reductions bring them close to their constraint, explaining the failure of principal reductions in HAMP to stimulate consumption. The collateral constraint renders principal reduction below the constraint ineffective.

More generally, we show that collateral constraints drive a wedge between an underwater borrower’s marginal propensity to consume out of cash and their marginal propensity to consume out of housing wealth. Borrowers near their collateral constraint have a high MPC out of housing wealth gains. However, borrowers far underwater are unresponsive. One policy implication is that high leverage is a bad “tag” for targeting policies that increase housing wealth, even though it is a good “tag” for targeting policies trying to provide cash to borrowers with high MPCs. Thus, policies to increase housing wealth are unable to stimulate demand when borrowers are so far underwater that home equity gains fail to relax binding collateral constraints, which was precisely the situation when such policies were implemented.

6.1 Setup

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

Households maximize expected utility, have time-separable preferences, and discount utility at

rate β . Per-period utility is

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

where c_t is non-housing consumption, d_i is an indicator variable equal to 1 if the household defaults, and ψ is a utility cost of defaulting. This additive default cost follows the structure in Campbell and Cocco (2015), Hembre (2014), Kaplan et al. (2016a), and Schelkle (2016). It reflects the moral and social stigma associated with defaulting on debt obligations as well as moving costs. We discuss the timing of default at the end of this section.

Agents consume a fixed quantity of housing. We assume housing and non-housing consumption are separable and, since quantity is fixed, follow Campbell and Cocco (2015) who show that under these conditions it is unnecessary to include housing explicitly in household preferences.³³ In the first period, agents are endowed with a home with market price P_{i1} and a 30-year fixed rate mortgage with balance M_{i1} and interest rate r . We assume home prices evolve deterministically according to $\Delta \log P = g$, where g is a constant, though we solve the model under various home price growth expectations. As long as households stay in this home, their housing costs include their mortgage payments (given by the standard annuity formula), property taxes τ_p that are proportional to the current market value of their home, and maintenance costs τ_m that are proportional to the initial value of their home.³⁴ Renters pay the user cost of housing for the equivalent home. Thus, housing payments are given by

$$h_{itj} = \begin{cases} M_{i1} \frac{r(1+r)^{30}}{(1+r)^{30}-1} + \tau_p P_{it} + \tau_m P_{i1}, & j = \text{owner} \\ (r - g + \tau_p) P_{it} + \tau_m P_{i1}, & j = \text{renter} \end{cases} \quad (4)$$

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

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

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

where $(1 - \phi)$ is the fraction of a house's value that can be used as collateral.³⁵ Renters are not

³³Campbell and Cocco (2015) show that these preferences are consistent with preferences over housing and non-housing consumption given by $\frac{c_{it}^{1-\gamma}}{1-\gamma} + \lambda_i \frac{H_{it}^{1-\gamma}}{1-\gamma}$ for $H_{it} = H_i$ fixed and where the parameter λ_i measures the importance of housing relative to non-housing consumption.

³⁴The assumption that maintenance costs are proportional to initial values ensures that maintaining the same home does not become more expensive simply because market home prices rise.

³⁵In the main parameterizations of our model house price growth is positive, such that once borrowers attain positive equity they do not risk falling back underwater. With negative home price growth, the borrowing limit is given by

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

in order to prevent forced deleveraging of liquid assets.

able to borrow.

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

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

where Γ_t reflects deterministic life-cycle growth and θ_{it} is an i.i.d transitory shock with $\mathbb{E}[\theta_{it}] = 1$. During retirement, income is given by a constant social security transfer which is captured in the Γ_t process. Total income, including income from home sales in the first period of retirement, is³⁶

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

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

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

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

We model default as a one-shot decision. Households begin the first period with a given mortgage, home price, and asset level. They then observe their first-period income shock, and decide whether to default or hold the house until retirement. This provides a simple way to analyze the short-term default decisions which we study empirically in section 4.4. In section 6.3 we study how changing the initial conditions by modifying a borrower's mortgage affects their default decision in the model, and compare this to our empirical results.

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

6.2 Parameterization

The main parameter values are summarized in Table 6. We assume that each period corresponds to one year. In our baseline case we assume households start life at age 45 and live with probability 1 until retirement at age 65. Survival probability shrinks every year during retirement, and households are dead with certainty by age 91 as assumed by Cagetti (2003). We solve the model for different first-period ages from 35 to 55 to examine the effect of principal reduction at different ages.

We follow Carroll (2012) who assumes income shocks have a lognormal component as well as an additional chance of a large negative shock. The large negative shock, which we call unemployment, captures the idea that the income process has a thick left tail (Guisen et al. 2014). Formally,

³⁶In all of our parameterizations borrowers have positive equity by retirement.

income shock θ is distributed as follows:

$$\theta_{it} = \begin{cases} b & \text{with probability } p \\ \frac{\delta_{it}(1-b)p}{1-p} & \text{with probability } (1-p) \end{cases} \quad (5)$$

where $\log \delta_{it} \sim \mathcal{N}\left(-\frac{\sigma_\delta^2}{2}, \sigma_\delta^2\right)$, p is the probability of unemployment, and b is the unemployment replacement rate. This ensures that $E[\theta_{it}] = 1$. All income risk, including unemployment, is turned off in retirement. We follow Carroll (1992) and set $\sigma_\delta = 0.14$.³⁷ We use data from Guvenen et al. (2014) to parameterize b and p . They show that the tenth percentile shock between 2008 and 2010 was a reduction in income of 50 percent, so we set p to 0.1 and b to 0.5. This large negative shock is critical to understanding default dynamics, which we explore in more detail in Section 6.3. The life-cycle growth path of permanent income Γ_t is from Carroll (1997).

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

We choose baseline preference values of $\beta = 0.96$ for the discount factor and $\gamma = 4$ for the coefficient of relative risk aversion. Our choice of a relatively high value for γ is not important for our consumption results, but is necessary in order to generate optimizing double-trigger behavior.³⁸

We estimate our final parameter ψ , the utility cost of default, such that the first-period default rate in the model matches the average 10% first-year default rate for moderately underwater borrowers in our data. Since our empirical default results focus on borrowers below 150 LTV, we allow default to rise above 10% for more underwater borrowers. We estimate ψ to equal 5.4 utils.

³⁷Carroll (1992) allows for temporary and permanent income shocks, each with a standard deviation of 0.1. We only have one income shock, whose standard deviation we set to $\sqrt{\sqrt{0.10} + \sqrt{0.10}} = 0.14$.

³⁸Our choice of a high γ ensures that agents default when they are hit with a bad income shock, but do not default under regular economic circumstances. The model exhibits this behavior because when γ is high, the value function for the agent paying her mortgage is much more concave than the value function for the agent who is defaulting, generating a region where default is sensitive to income. In contrast, when γ is low in our model, LTV is the primary determinant of default decisions, which is inconsistent with our empirical findings. We discuss this choice in more detail in Section 6.3.

To translate this into meaningful units, we calculate that this is equivalent to a 10% permanent income loss. This loss is in line with other estimates in the literature that uses structural models with default costs to match observed default rates. Schelkle (2016) builds a model to match the rise in default rates in the U.S. between 2002 and 2010 and estimates a default cost equal to 8% of permanent income. Kaplan et al. (2016a) calibrate a default cost to match the foreclosure rate in the late 1990s and find a cost which is equal to 4% of permanent income for the median household, and approximately 7% for mortgagors. Hembre (2014) studies default behavior for all HAMP modifications and finds that a cost equal to 70% of per-period consumption is necessary to explain observed default rates.

6.3 Default

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

6.3.1 The Effect of Principal Reduction on Default

In forward-looking models with a housing asset and labor income risk, default emerges from two motives: (1) an agent is so far underwater that her house is no longer a good investment and (2) default offers a way to access short-term liquidity when cash-on-hand is low. In our model, the core tradeoff underwater borrowers face when making their default decision is whether the short-term gain from reduced housing payments is worth the utility cost of defaulting and the lost resale value of the house at retirement.³⁹ Both the costs and benefits of default vary with current payment levels, current incomes, and total debt obligations. When borrowers have high current payments or low current incomes, the short-term payment relief is particularly valuable because it allows borrowers to avoid making severe cuts to consumption. Similarly, when total debt levels are high, the costs of default are low because the house is less valuable as an asset.

To show the effect of principal reduction and relate it back to our empirical results, we simulate changes in mortgage principal holding payments constant. We assume homeowners receive modifications at age 45. To match the low assets of delinquent borrowers in the PSID, we set initial assets $a_t = 0.01$ units of permanent income. We set initial LTV equal to 119, the median pre-modification LTV for borrowers in our regression discontinuity analysis (Table 2). We then vary the LTV, holding mortgage payments for households that have not defaulted fixed for five years, after which payments fall according to the annuity formula in equation 4 applied to the new mortgage balance. To highlight the importance of payment levels, we conduct the exercise both

³⁹Because we assume house prices evolve deterministically, our model does not capture the option value of mortgages. With house price uncertainty, paying a mortgage is equivalent to purchasing a call option, giving the borrower the right to “buy” future home equity gains, if realized, at the price of the unpaid balance on the mortgage. Incorporating house price uncertainty would reduce the gain from defaulting and would lead us to estimate a smaller utility cost of defaulting to match the average 10% default rate.

under our baseline parameters and under an alternative assumption where housing payments are an additional 5% of income.

The top panel of Figure 9 shows that for a given current payment level and LTV ratio, there is a cash-on-hand level below which households will find it optimal to default. The more underwater the household, the smaller the income shock necessary to push them to default. For borrowers with high current payments, the income threshold is both high and sensitive to debt levels. Even some borrowers with incomes equal to their permanent income (for a cash-on-hand value of 1 when assets are zero) will find it optimal to default when payments are this high and they are underwater. However, for borrowers in our baseline scenario, the income cutoff for defaulting is both low and relatively insensitive to debt levels. In particular, below LTVs of about 150, low-asset borrowers will only default if their income is less than three-quarters of its permanent level, a shock of about two standard deviations. This means that default is most likely to occur for borrowers who are hit with “unemployment”, the large negative shock in our income process. Figure 5 shows event study evidence that mortgage default is coincident with a negative income shock.

We find that default rates are insensitive to principal reduction for the typical borrower. The bottom panel of Figure 9 plots the default rate in the first period after modification for borrowers with various amounts of principal reduction. In our baseline case, where borrowers already have low-cost mortgages, additional principal reduction is ineffective below an LTV of about 160. For such moderately underwater borrowers with relatively affordable mortgages, the gain from defaulting is not worth the cost unless they are hit by a severe income shock. However, far underwater borrowers have much higher default rates because they default even in the absence of large income shocks.

Our model shows that payment reductions reduce the effectiveness of principal reduction. If borrowers do not already have low-cost mortgages, as is the case depicted in the orange line of Figure 9, default rates for underwater borrowers are both high and sensitive to mortgage principal. For these borrowers, the gain of defaulting is much greater. This result supports the view that payment reductions in HAMP reduced default, but that conditional on those payment reductions, principal reductions that reduced long-term debt obligations were ineffective.

6.3.2 An Optimizing Double Trigger Model of Default

Borrowers in our baseline case exhibit what we call “optimizing double trigger” behavior. In the “double trigger” class of models, agents default when two conditions are triggered: (1) they are underwater and (2) face negative income shocks. In the most basic of these models, agents are not optimizing. While negative equity is necessary for default, the level of negative equity is irrelevant (see description of these models in Gerardi et al. 2015). Agents do not consider the costs and benefits of defaulting, they simply default when they are forced to by an income shock that leaves them without enough funds to pay their mortgage (Guren and McQuade 2015).

In our model, agents are optimizing. At moderate levels of underwaterness, it is only optimal for unemployed agents to default. Default is insensitive to negative equity in this region because the costs of default are high and the gains for an employed agent are low. However, beyond about

160 LTV, their optimizing behavior generates a steep causal relationship between LTV and default. These borrowers are defaulting for what is sometimes referred to as “strategic” reasons, i.e. they default even when their payments are affordable.

The optimizing double trigger behavior, with a small effect of LTV on default at low LTV levels followed by a steep slope at high LTV levels, is consistent with recent dynamic models of mortgage default such as Schelkle (2016) and Campbell and Cocco (2015). Campbell and Cocco (2015) study default decisions in a calibrated model where borrowers are liquidity constrained and face labor income, house price, inflation, and interest rate risk. In their model the kink occurs at about 135 LTV. Below this level, the option value of staying in the mortgage outweighs the gains of defaulting for most borrowers. Our empirical evidence suggests that default is insensitive to LTVs even at slightly higher LTV ratios, which is consistent with adding a utility cost of default to this type of model. The result that borrowers without income shocks do not exercise their default option until substantially underwater is consistent with empirical evidence in Bhutta et al. (2011), who show that the median homeowner without an income shock does not default until their LTV is greater than 167.⁴⁰

In our model, the key force generating our results is that the income cutoff for defaulting is not very sensitive to the size of mortgage debt. This generates a flat, positive-default-rate region followed by a steep slope at high LTV levels. Generating this region, which is consistent with our empirical evidence, relies on three empirically plausible features of our model. First, most underwater borrowers do not default because they would incur a utility cost of default. This is supported by survey evidence in Guiso et al. (2013), who find that about 80 percent of homeowners consider it morally wrong to default when payments are affordable. Second, agents face thick-tailed income shocks (Guisen et al. 2014).⁴¹ Third, households are risk averse and default when hit with a very bad income shock. When we reduce risk aversion to $\gamma = 2$, default rates are either zero or high, with no flat, positive-default-rate region.⁴²

Our empirical evidence favors models like ours over alternatives that generate smooth upward-sloping relationships between LTV and default. Kau et al. (1993) and Stanton and Wallace (1998) build off of the frictionless option model that predicts a single cutoff LTV value above which all borrowers default. Because the cross-sectional relationship between LTV and default is smooth, these authors propose introducing a distribution of additional default costs, which generates a distribution of cutoff values and therefore a smooth relationship between LTV and default. We add a distribution of default costs in our model in Appendix Figure A.14, and show that this does generate a smooth relationship between LTV and default. However, our empirical results,

⁴⁰Similarly, Foote et al. (2008) study homeowners in Massachusetts who were underwater in the early 1990s, and find that fewer than one percent eventually lost their home to foreclosure.

⁴¹If we eliminate this feature of our income process, we estimate both a smaller stigma cost in order to match an average 10% default rate, and we find that default is sensitive to LTVs even at low LTV levels, which is inconsistent with our empirical results.

⁴²The short-term liquidity motive for default is most valuable when risk aversion is high. When risk aversion is low, default is largely a function of LTV. As the utility function becomes increasingly linear, the function mapping LTV to default becomes increasingly binary, approaching a rule of thumb where *no* agents default below an LTV cutoff and *all* agents default above the LTV cutoff.

which find that default is insensitive to LTVs for moderate amounts of underwaterness, reject this parameterization of our model.

6.3.3 Policy Implications from Default Results

The finding from our empirics and model that short-term cash-flow shocks, rather than overall debt levels, are the key drivers of default until borrowers are far underwater has two implications for policy. First, from the point of view of reducing defaults in the short-term, principal reduction is costly but ineffective. It is less costly (to the government or the lender) to temporarily reduce payments than to permanently reduce debt levels. But once payments are reduced in the short-term, overall debt levels are relatively unimportant for default decisions except for the most underwater borrowers. The government paid investors for every dollar of principal reduction down to 105 LTV. Our findings suggest the majority of these payments were unnecessary.

Second, evidence that borrowers are unlikely to default absent substantial income shocks unless far underwater implies that the moral hazard cost of the default option as a form of social insurance is low (Bhutta et al. 2011). Mortgage default can be thought of as a social insurance program, since the government enforces laws making it costly for lenders to recover unpaid balances. For example, in 10 U.S. states, mortgages are non-recourse, meaning that borrowers can default on their mortgages with no liability for the remaining mortgage balance (Mitman 2016). Furthermore, deficiency judgments against borrowers are rare, even in recourse states where they are legal (Ghent and Kudlyak 2011). Some states force lenders to go through a lengthy judicial process to initiate a foreclosure. A range of other policy interventions during the Great Recession lengthened foreclosure timelines, allowing delinquent borrowers to remain in their homes for months without making payments (Herkenhoff 2015). These laws pass on the cost of default from the borrower to the lender. For traditional social insurance programs the government is concerned about moral hazard. However, evidence that mortgage default is mainly driven by affordability implies the moral hazard cost of the default option is low.

6.4 Consumption

In this section, we explore the effect of principal reduction on consumption. We show that when borrowers face collateral constraints, principal reductions that leave households underwater do not relax current constraints and have little precautionary value, hence do not translate into increased consumption.

6.4.1 Comparison to Boom-Era Housing MPC Distribution

Our model makes reasonable quantitative predictions about consumption out of housing wealth changes, for which prior empirical papers provide an external benchmark. We focus on replicating estimates corresponding to the pre-2009 period and use Mian et al. (2013) as our external benchmark. We use our model to estimate the MPC out of housing wealth gains for age 45 borrowers with different initial LTVs. We endow each agent with cash-on-hand equal to two years of perma-

ment income, which is the median non-housing wealth for all homeowners in the 2007 Survey of Consumer Finances (SCF).⁴³ We then calculate the MPC for these agents at different LTV values, and weight them according to the distribution of LTV in 2007 reported in Carter (2012).⁴⁴

Table 7 reports the average MPC out of an additional dollar of housing equity for the average borrower as well as for high-leverage (but still above-water) borrowers. We find MPCs of 8 and 15 cents, respectively. These are similar to the average MPC for homeowners of 9 cents reported in Mian et al. (2013), and the 18 cent MPC of homeowners living in counties with average LTV ratios above 90. In our model, high-leverage above-water borrowers have high MPCs because they have low housing wealth and are the most borrowing constrained. Having shown that our model produces quantitatively reasonable predictions for above-water borrowers, we now turn to analyzing its predictions for underwater borrowers.

6.4.2 Sufficient Statistic Expression for Principal Reduction

To build intuition for the effect of principal reductions on consumption, we consider a simplified version of our model without a default option, in which we can develop a straightforward formula for the effect of debt levels on consumption. In this case a homeowner’s problem can be written as a function of four state variables: cash-on-hand (m_{it}), the wealth gain from home sale at retirement ($w_{iT_y} = P_{iT_y} - M_{iT_y}$), and the vectors of housing payments and collateral constraints for the rest of life (\vec{h}_i, \vec{a}_i). We can then decompose the effect of a change in mortgage debt level at date t in the following way:

$$\begin{aligned} \frac{dc_{it} \left(m_{it}, w_{iT_y}, \vec{h}_i, \vec{a}_i \right)}{dM_{it}} &= \frac{\partial c_{it}}{\partial w_{iT_y}} \cdot \frac{\partial w_{iT_y}}{\partial M_{it}} + \sum_{s=t}^T \frac{\partial c_{it}}{\partial h_{is}} \cdot \frac{\partial h_{is}}{\partial M_{it}} + \sum_{s=t}^T \frac{\partial c_{it}}{\partial a_{is}} \cdot \frac{\partial a_{is}}{\partial M_{it}} \\ &= \underbrace{MPC_{t, w_{iT_y}} \frac{\partial w_{iT_y}}{\partial M_{it}} + \sum_{s=t}^T MPC_{t, h_{is}} \cdot \frac{\partial h_{is}}{\partial M_{it}}}_{\text{Future wealth effect}} + \underbrace{\sum_{s=t}^T MPC_{t, a_{is}} \cdot \frac{\partial a_{is}}{\partial M_{it}}}_{\text{Collateral effect}}. \quad (6) \end{aligned}$$

Equation 6 shows that a reduction in debt levels affects today’s consumption through two channels. The first is a future wealth effect. Reducing mortgage debt increases a homeowner’s expected home equity gain when they sell the house and reduces their housing payments every year. These translate into consumption according to the homeowner’s marginal propensity to consume today out of wealth gains in future dates. The second channel is a collateral effect. Reducing debt levels frees up home equity that raises the household’s borrowing limit over time. This change translates into consumption today according to the homeowner’s marginal propensity to consume out of increased collateral in future dates.

This expression clarifies that the key forces determining the consumption response to debt forgiveness are the timeline under which debt reductions translate into higher wealth, lower payments, and increased borrowing capacity, and the borrower’s marginal propensity to consume out of these

⁴³Mian et al. (2013) show that wealth does not vary with LTV, so we assign this median number to all borrowers.

⁴⁴Carter (2012) reports LTV distributions in 2005 and 2009, so we take the average.

future wealth and future collateral gains. Berger et al. (2016) develop a related sufficient statistic formula, which shows that the response to housing wealth gains achieved from increased house prices depends on the marginal propensity to consume out of wealth and current house values. Our model clarifies that the MPC out of wealth gains will depend crucially on a borrower’s initial home equity position. Underwater borrowers receiving a dollar of housing wealth are only able to monetize this gain by borrowing or selling their home once they are above water, which may be far in the future. In this case, the consumption response today will depend on a household’s MPC out of expected wealth or collateral gains far in the future.

Since key drivers of the consumption response to debt forgiveness are the effect of wealth and collateral gains at various future dates on today’s consumption, we study each of these effects in order to gain intuition about the expected impact of debt forgiveness. The top panel of Figure 10 plots consumption as a function of cash-on-hand under various scenarios. The green line shows the baseline scenario, where household wealth only comes from liquid assets. This is the standard concave consumption function delivered by life-cycle models with idiosyncratic income risk. The orange line shows the consumption function when the household is given a wealth grant equal to one year’s worth of income (about the size of the principal forgiveness we study empirically). Consumption is responsive across the cash-on-hand distribution, but especially responsive at low cash-on-hand levels.

Household consumption is unresponsive to wealth grants several years in the future. In blue we plot the consumption function if the same-sized grant is promised one year in the future. In this case, the most constrained borrowers have no response, since the only way they could respond is by cutting savings, but they are already consuming all their assets.⁴⁵ However, for households who have moderately low assets but are saving in the current period as a buffer against future income risk, a wealth grant one year in the future has a large effect. This is because such a grant reduces the precautionary benefit of saving in the current period. The purple line shows the response to a wealth grant six years in the future, which is when payments start falling for the borrowers receiving principal forgiveness in our empirical setting. For low cash-on-hand borrowers, this grant has little precautionary value because it is so far in the future, hence consumption is unaffected. This result is related to recent work by McKay et al. (2016), who show that households facing income risk and building buffer stocks will be less responsive to interest-rate changes farther in the future because responding requires depleting valuable savings in the present. Only high cash-on-hand households respond to wealth grants they can only monetize this far in the future. For the highest-wealth households with cash-on-hand equivalent to three years worth of permanent income, a wealth grant in the future is roughly equivalent to a wealth grant today, since these households act like Permanent Income Hypothesis (PIH) consumers and consume the annuity value of a grant regardless of its timing.

Household consumption is also unresponsive to increased borrowing limits expected several

⁴⁵In models without an income floor, borrowers will never consume 100% of their assets, since this would expose them to potentially zero consumption in the following period. In our model the worst outcome is an income equal to 50% of permanent income, which drives some borrowers to consume all their assets in the current period.

years in the future. The bottom panel of Figure 10 shows the equivalent exercise, but instead of granting the households wealth at different dates, they are granted collateral at different dates. This figure shows that for low-wealth households, the response to increased collateral is the same as the response to increased wealth. Only at very high wealth levels is collateral less beneficial than wealth.

Overall, Figure 10 shows that there are three important regions for understanding the effect of wealth and collateral gains in the future. First, at very low cash-on-hand levels, consumption is only responsive to wealth or collateral that can be used in the current period. Second, for borrowers with moderate cash-on-hand between 0.7 and 1.5 years worth of permanent income, consumption is responsive to gains that can be monetized in the near-term, but not six or more years in the future. Finally, for borrowers with higher cash-on-hand, there is a small response even to gains that can only be monetized six or more years in the future. This means that for low-wealth borrowers, principal reduction will likely be ineffective unless it can be monetized in the near-term.

6.4.3 Consumption Response to Principal Reduction for Underwater Borrowers

In our model, principal reduction similar to that implemented in HAMP does not raise consumption. We consider an average household with first period income $y_t = 0.85$ units of permanent income, based on Bernstein (2015) who finds that borrowers receiving mortgage modifications during the recent crisis had temporarily low incomes. We set initial LTV equal to 150, the median pre-modification LTV for borrowers receiving principal reduction in our difference-in-differences analysis.⁴⁶ For our treatment group, we then reduce their mortgage balance by \$70,000, bringing them to an LTV of 106.⁴⁷ As with our default policy simulation, to mimic the policy implemented in HAMP we keep mortgage payments for households who have not defaulted fixed for five years.

Principal reduction translates into increased borrowing capacity and increased wealth with a considerable delay. In Section 6.4.2 we showed that the consumption response depends on the timing of reduced payments and increased collateral under the policy reform, so we plot these changes in Figure 11. These figures show borrowing limits and housing payments over time under the status-quo and under the principal reduction policy (labeled as “Payment Reduction” and “Payment and Principal Reduction”, respectively, to parallel our empirical setting). The top panel shows that principal reduction eventually increases borrowing limits, but that these increases do not occur for another eight years. This is because even after receiving principal reduction, borrowers remain slightly underwater. Furthermore, to be able to borrow against their home given the collateral constraint they need to get down to an LTV of 80, which takes several years under baseline price growth and mortgage principal pay-down schedules. The bottom panel shows that housing payments decrease substantially, but only starting six years in the future.⁴⁸ Since the

⁴⁶This corresponds to an initial home price equal to \$173,000 (or 3.25 units of permanent income) and an initial mortgage debt of \$259,000 (or 4.88 units of permanent income).

⁴⁷The median LTV post-modification in our data is actually 114, because borrowers’ unpaid mortgage payments are capitalized into the new mortgage balance. We abstract from this in our model, though it would only serve to further reduce the effect of principal reduction.

⁴⁸Housing payments are declining in the first five years because they are plotted relative to permanent income,

lesson from Figure 10 was that household consumption for low-wealth borrowers is unresponsive to wealth and collateral gains that are only monetized several years in the future, these plots suggest that the principal reduction policy we have modeled will not have a large effect on consumption.

Figure 12 shows directly that principal reduction as done in HAMP is ineffective at increasing consumption in our model. The top panel of Figure 12 shows the consumption function out of increasing amounts of principal reduction for borrowers starting at a LTV of 150. It shows that the consumption function out of home equity gains is S-shaped, convex below the collateral constraint and concave above it. Borrowers are insensitive to principal reductions until such reductions bring them close to their constraint. Principal reduction for deeply underwater borrowers does not relax current constraints and has little precautionary value, hence consumption is unaffected.

The model clarifies that the collateral constraint renders principal reduction below this constraint ineffective. Because borrowers are far from the boundary where they can monetize increased housing wealth, they are insensitive to wealth increases. The intuition is similar to the result in Berger and Vavra (2015), who show that in recessions households are farther from their adjustment boundaries for durable consumption, and hence are less sensitive to wealth shocks. Furthermore, our result is consistent with Kaplan et al. (2016a), who find in a general equilibrium heterogeneous-agent model that a policy forgiving all excess mortgage debt above an LTV of 95 has only a trivial effect on consumption.⁴⁹

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 (2016), debt is modeled as a one-period bond. In this setting, when a credit crunch reduces the borrowing limit, borrowers who find themselves beyond the borrowing constraint are forced to immediately cut consumption in order to delever. Applying this assumption to the mortgage setting implies that when housing prices fall such that the loan-to-value ratio becomes greater than 100, borrowers need to immediately repay their outstanding debt until they are abovewater. Under this hypothetical scenario, borrowers receiving principal reduction would see immediate reductions in the amount of forced repayment. Principal reduction would increase consumption by reducing debt overhang. But in practice, as in our model, mortgages in particular are long-term loans. Nothing forces borrowers to immediately delever when they are far underwater. Modeling housing debt as a long-term contract – unlike the short-term contracts in Eggertsson and Krugman (2012) and Guerrieri and Lorenzoni (2016)– removes a mechanical link between debt levels and consumption present in some of the prior literature, and reduces the expected effectiveness of mortgage debt reduction policies. Other recent papers to consider the effect of debt and housing wealth in settings with long-term contracts include Berger et al. (2016), Chen et al. (2013), Kaplan et al. (2016a), and

which is rising at these ages according to the standard life-cycle pattern.

⁴⁹While the results are consistent, the mechanism is different. The principal reduction policy modeled in Kaplan et al. (2016a) reduces both short-term and long-term payments. They find that borrowers who would have made payments anyways increase consumption as a result of payment reductions. However, borrowers who would have defaulted without principal reduction, but no longer default under the new policy, see mortgage payments rise relative to the status-quo and therefore reduce their consumption. Averaging over these two groups, they find little impact on aggregate consumption.

Justiniano et al. (2015).

6.4.4 Consumption Response to Principal Reduction Under Alternative Parameterizations

In our model, principal reduction is ineffective under a variety of alternative parameterizations. Table 8 reports the MPC for the principal reduction policy experiment described above under various alternative assumptions. The baseline MPC is 0.3 cents per dollar of mortgage principal reduced. This is similar to the 0.2 cent estimate in our empirical results. Changing borrower age, discount rate, and risk aversion has little impact on the MPC.

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

Generating a large consumption response requires an alternative unrealized economic environment (relaxed collateral constraints and optimistic home price growth) or an alternative policy of more generous writedowns. Setting the collateral constraint to zero such that homeowners can lever up to 100 LTV generates a moderate MPC of 4.8 cents. Even though borrowers remain underwater after principal reduction, allowing them to monetize wealth starting at 100 LTV would have some immediate precautionary value. Similarly, if households expected permanent real annual house price growth of 5% (equal to realized growth rates from 2000 to 2005), the MPC would be 6.2 cents because borrowers would expect to be able to monetize their housing wealth more quickly. Combining both of these assumptions about the economic environment generates a large MPC of 24.0 cents. However, the period when principal reduction was implemented is exactly when neither of these conditions was likely to hold. In the aftermath of the crisis, home price growth expectations were tepid and credit supply was tight.⁵⁰

Although more generous principal reductions might have led to some observed consumption response in the data, they would have been an inefficient way to raise consumption. For example, a writedown to 90 LTV with a relaxed collateral constraint raises consumption in our model. However, Figure 12 shows that all of the consumption increase is coming from the region near the collateral constraint. This means that a policy of targeting deeply underwater borrowers with more generous writedowns will expend substantial resources in a region with no stimulative effect.

⁵⁰See Department of Housing And Urban Development (2016) for house price expectations data and Corelogic (2016) for evidence of tightening collateral constraints.

6.4.5 Collateral Constraints Drive a Wedge Between MPC Out of Housing Wealth and MPC out of Cash

The inability to monetize housing wealth drives a wedge between an underwater borrower’s marginal propensity to consume out of cash and their marginal propensity to consume out of housing wealth. This can be observed visually in the bottom panel of Figure 12. In this figure we take low-cash-on-hand borrowers at various LTV levels and plot the MPC out of \$1 of cash or \$1 of housing wealth gained by principal reduction.⁵¹ As in the empirical results in Mian et al. (2013), borrowers near their collateral constraint have a high MPC out of housing wealth gains. However, borrowers far underwater are unresponsive to housing wealth changes even though they are highly responsive to cash transfers. This highlights one way that housing wealth is special. Because it can only be monetized when borrowers have positive home equity above a collateral constraint, borrowers respond less to housing wealth gains than they do to cash.

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

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

⁵¹In this experiment, housing payments fall immediately when debt is reduced, unlike in Section 6.4.3 where we delayed payment relief in order to mimic principal reduction in HAMP. We consider the average household from our policy experiment, so we set cash-on-hand to 0.85.

7 Conclusion

This paper empirically and theoretically evaluates the impact of substantial mortgage principal reductions. Our empirical analysis exploits a distinctive government policy that allows us to isolate the effects of reducing long-term payments and increasing home equity while controlling for short-term payments. Using regression discontinuity and difference-in-differences research designs, we find that principal reductions have no short-term impact on default, durable consumption, or nondurable consumption for borrowers who remain underwater. We build a quantitative life-cycle model with a housing asset and default to help interpret these empirical findings. The model clarifies that a borrower's short-term constraints govern their response to long-term debt obligations. When defaulting imposes short-term utility costs, default is more responsive to cash-flow shocks than to future debt burdens. When principal reductions do not push borrowers sufficiently above water so as to relax collateral constraints, consumption is unaffected because borrowers are unable to monetize increased housing wealth.

Although we find that principal reductions have no impact in the short-term, an interesting avenue for future research will be to estimate the long-term impact. Eventually, borrowers who received more substantial reductions will regain positive home equity. When faced with an income or expenditure shock, instead of defaulting and facing foreclosure, borrowers will be able to sell their home. Similarly, borrowers who received principal reductions will more quickly gain enough home equity to overcome their collateral constraint, leading to increases in their borrowing capacity. In future work, we hope to study both of these effects.

A second natural experiment that will emerge over time comes from expected payment increases. Although payments for both modification groups are fixed for five years, they begin rising in year six as the temporary interest rate reduction is phased out. This will generate a completely expected negative shock to disposable income. Prior work on mortgage payments has focused on payment reductions which increase disposable income. In these cases, failure to increase consumption before payments fall can be attributed either to liquidity constraints or to bounded rationality. But in settings with predictable reductions in disposable income, even liquidity constrained households should cut spending preemptively. In prior work we have found that unemployed agents fail to cut spending ahead of predictable exhaustion of their unemployment benefits (Ganong and Noel 2016). The expected mortgage payment increases that will eventually happen for HAMP borrowers will provide a new way to test between the explanations for consumer sensitivity to expected changes in disposable income.

References

- Adelino, M., Schoar, A., and Severino, F. (2016). Loan Origination and Defaults in the Mortgage Crisis: The Role of the Middle Class. *Review of Financial Studies*, 29(7):1635–1670.
- Agarwal, S., Amromin, G., Ben-David, I., Chomsisengphet, S., Piskorski, T., and Seru, A. (2016a). Policy Intervention in Debt Renegotiation: Evidence from the Home Affordable Modification Program. Research Paper 3459, Stanford University, Graduate School of Business.
- Agarwal, S., Amromin, G., Ben-David, I., Chomsisengphet, S., and Evanoff, D. D. (2011). The role of securitization in mortgage renegotiation. *Journal of Financial Economics*, 102(3):559–578.
- Agarwal, S., Amromin, G., Chomsisengphet, S., Piskorski, T., Seru, A., and Yao, V. W. (2015). Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinancing Program. SSRN Scholarly Paper ID 2662906, Social Science Research Network, Rochester, NY.
- Agarwal, S., Chomsisengphet, S., Mahoney, N., and Stroebel, J. (2016b). Do Banks Pass Through Credit Expansions? Asymmetric Information and the Marginal Profitability of Consumer Lending During the Great Recession. *SSRN Electronic Journal*.
- Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, 109(3):659–84.
- Athey, S. and Imbens, G. (2016). The State of Applied Econometrics - Causality and Policy Evaluation. *arXiv:1607.00699 [stat]*.
- Attanasio, O. P., Blow, L., Hamilton, R., and Leicester, A. (2009). Booms and Busts: Consumption, House Prices and Expectations. *Economica*, 76(301):20–50.
- Berger, D., Guerrieri, V., Lorenzoni, G., and Vavra, J. (2016). House Prices and Consumer Spending. Working Paper, University of Chicago and Northwestern University.
- Berger, D. and Vavra, J. (2015). Consumption Dynamics During Recessions. *Econometrica*, 83(1):101–154.
- Bernstein, A. (2015). Household Debt Overhang and Labor Supply. SSRN Scholarly Paper ID 2700781, Social Science Research Network, Rochester, NY.
- Bhutta, N., Dokko, J., and Shan, H. (2011). Consumer Ruthlessness and Mortgage Default During the 2007-2009 Housing Bust. SSRN Scholarly Paper ID 1626969, Social Science Research Network, Rochester, NY.
- Cagetti, M. (2003). Wealth Accumulation over the Life Cycle and Precautionary Savings. *Journal of Business & Economic Statistics*, 21(3):339–353.
- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, 82(6):2295–2326.
- Campbell, J. Y. and Cocco, J. F. (2007). How do house prices affect consumption? Evidence from micro data. *Journal of Monetary Economics*, 54(3):591–621.
- Campbell, J. Y. and Cocco, J. F. (2015). A Model of Mortgage Default. *Journal of Finance*, 70(4):1495–1554.

- Carroll, C. (2012). Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems. Mimeo, Johns Hopkins University.
- Carroll, C. D. (1992). The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence. *Brookings Papers on Economic Activity*, 1992(2):61–156.
- Carroll, C. D. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *The Quarterly Journal of Economics*, 112(1):1–55.
- Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320.
- Carroll, C. D., Otsuka, M., and Slacalek, J. (2011). How Large Are Housing and Financial Wealth Effects? A New Approach. *Journal of Money, Credit and Banking*, 43(1):55–79.
- Carter, G. (2012). Housing Units With Negative Equity, 1997 to 2009. *Cityscape*, 14(1):149–165.
- Case, K. E., Quigley, J. M., and Shiller, R. J. (2013). Wealth Effects Revisited: 1975-2012. Working Paper 18667, National Bureau of Economic Research.
- Chen, H., Michaux, M., and Roussanov, N. L. (2013). Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty. SSRN Scholarly Paper ID 2024392, Social Science Research Network, Rochester, NY.
- Clara, N. and Cocco, J. F. (2016). An Analysis of Consumer Debt Restructuring Policies. Technical Report 480, Society for Economic Dynamics.
- Corelogic (2016). Home Equity Lending Landscape. White Paper.
- Cowley, S. (2017). Foreclosure Prevention Returns to the Unknown. The New York Times, <https://www.nytimes.com/2017/01/25/business/dealbook/foreclosure-prevention-mortgage-lending.html>, January 25, 2017.
- DeFusco, A. A. (2016). Homeowner Borrowing and Housing Collateral: New Evidence from Expiring Price Controls. SSRN Scholarly Paper ID 2696677, Social Science Research Network, Rochester, NY.
- Department of Housing And Urban Development, U. (2016). Housing Scorecard. HUD Report, September 2016, Available at: http://portal.hud.gov/hudportal/HUD?src=/initiatives/Housing_Scorecard.
- Diaz, A. and Luengo-Prado, M. J. (2008). On the user cost and homeownership. *Review of Economic Dynamics*, 11(3):584–613.
- DiMaggio, M., Kermani, A., Keys, B., Piskorski, T., Ramchaan, R., Seru, A., and Yao, V. (2017). Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging. *American Economic Review*, 107(11):3550–3588.
- Disney, R., Gathergood, J., and Henley, A. (2010). House Price Shocks, Negative Equity, And Household Consumption In The United Kingdom 4. *Journal of the European Economic Association*, 8(6):1179–1207.

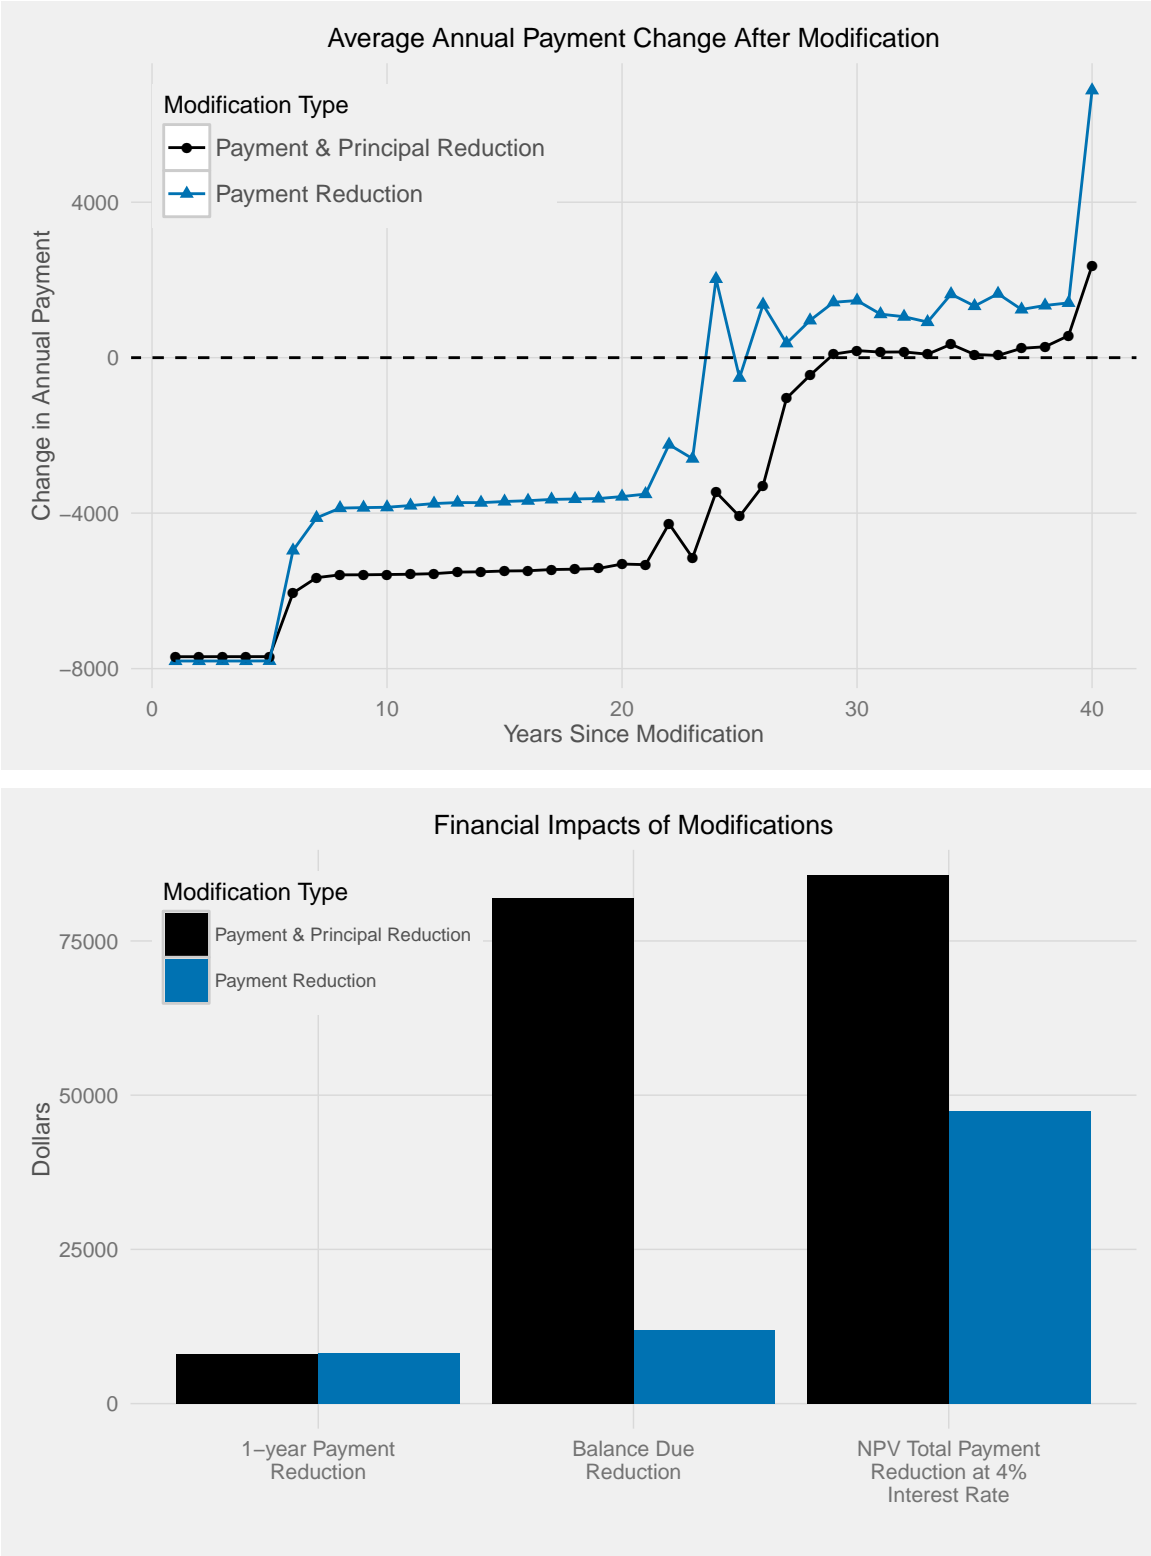
- Dobbie, W. and Song, J. (2016). Debt Relief or Debt Restructuring? Evidence From an Experiment with Distressed Credit Card Borrowers. Working Paper, Princeton University and Social Security Administration.
- Eberly, J. and Krishnamurthy, A. (2014). Efficient Credit Policies in a Housing Debt Crisis. *Brookings Papers on Economic Activity*, 49(2 (Fall)):73–136.
- Eggertsson, G. B. and Krugman, P. (2012). Debt, deleveraging, and the liquidity trap: a Fisher-Minsky-Koo approach. *The quarterly journal of economics*, 127(3):1469–1518.
- Ehrlich, G. and Perry, J. (2015). Do Large-Scale Refinancing Programs Reduce Mortgage Defaults? Evidence From a Regression Discontinuity Design: Working Paper 2015-06. Working Paper 50871, Congressional Budget Office.
- Experian and Wyman, O. (2009). Understanding Strategic Default in Mortgages Part 1. White Paper.
- Federal Reserve, N. Y. (2015). Quarterly Report on Household Debt and Credit. Federal Reserve Bank of New York, Research and Statistics Group, November 2015.
- Federal Reserve, U. S. (2014). *The 2013 Federal Reserve Payments Study*. The Federal Reserve System.
- Feldstein, M. (2009). How to Save an 'Underwater' mortgage. Wall Street Journal, August 7, 2009.
- Fisher, I. (1933). The Debt-Deflation Theory of Great Depressions. *Econometrica*, pages 337–357.
- Foote, C. L., Gerardi, K., and Willen, P. S. (2008). Negative equity and foreclosure: Theory and evidence. *Journal of Urban Economics*, 64(2):234–245.
- Fuster, A. and Willen, P. (2015). Payment Size, Negative Equity, and Mortgage Default. *Federal Reserve Bank of New York Staff Reports*, No. 582.
- Fuster, A. and Willen, P. (2017). Payment Size, Negative Equity, and Mortgage Default. *American Economic Journal: Economic Policy*, 9(4):167–191.
- Gabriel, S. A., Iacoviello, M. M., and Lutz, C. (2016). A Crisis of Missed Opportunities? Foreclosure Costs and Mortgage Modification During the Great Recession. SSRN Scholarly Paper ID 2831830, Social Science Research Network, Rochester, NY.
- Ganong, P. and Noel, P. (2017). Consumer Spending During Unemployment: Positive and Normative Implications.
- Geanakoplos, J. and Koniak, S. (2009). Matters of Principal. The New York Times, March 4, 2009.
- Geithner, T. (2014). *Stress Test*. Broadway Books.
- Gerardi, K., Herkenhoff, K. F., Ohanian, L. E., and Willen, P. S. (2015). Can't Pay or Won't Pay? Unemployment, Negative Equity, and Strategic Default. NBER Working Paper 21630, National Bureau of Economic Research, Inc.
- Ghent, A. C. and Kudlyak, M. (2011). Recourse and Residential Mortgage Default: Evidence from US States. *The Review of Financial Studies*, 24(9):3139–3186.

- Goldfarb, Z. (2012). Economists, Obama Administration at Odds Over Role of Mortgage Debt in Recovery. *The Washington Post*, November 22, 2012.
- Guerrieri, V. and Lorenzoni, G. (2016). Credit Crises, Precautionary Savings, and the Liquidity Trap. Working Paper, University of Chicago and Northwestern University.
- Guiso, L., Sapienza, P., and Zingales, L. (2013). The Determinants of Attitudes toward Strategic Default on Mortgages. *Journal of Finance*, 68(4):1473–1515.
- Guren, A. and McQuade, T. (2015). How do Foreclosures Exacerbate Housing Downturns? Working Paper, Boston University and Stanford University.
- Guren, A. M., Krishnamurthy, A., and McQuade, T. J. (2017). Mortgage Design in an Equilibrium Model of the Housing Market. *working paper*.
- Guvener, F., Ozkan, S., and Song, J. (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy*, 122(3):621–660.
- Haughwout, A., Okah, E., and Tracy, J. (2016). Second Chances: Subprime Mortgage Modification and Redefault. *Journal of Money, Credit and Banking*, 48(4):771–793.
- Hembre, E. (2014). HAMP, Home Attachment, and Mortgage Default. Working Paper, University of Wisconsin.
- Herkenhoff, K. F. and Ohanian, L. E. (2015). The Impact of Foreclosure Delay on U.S. Employment. Working Paper 21532, National Bureau of Economic Research.
- Himmelberg, C., Mayer, C., and Sinai, T. (2005). Assessing High House Prices: Bubbles, Fundamentals and Misperceptions. *Journal of Economic Perspectives*, 19(4):67–92.
- Holden, S., Kelly, A., McManus, D. A., Scharlemann, T. C., Singer, R., and Worth, J. (2012). The HAMP NPV Model: Development and Early Performance. SSRN Scholarly Paper ID 2194279, Social Science Research Network, Rochester, NY.
- Hsu, J. W., Matsa, D. A., and Melzer, B. (2014). Positive Externalities of Social Insurance: Unemployment Insurance and Consumer Credit. SSRN Scholarly Paper ID 2185198, Social Science Research Network, Rochester, NY.
- Imbens, G. and Kalyanaraman, K. (2012). Optimal Bandwidth Choice for the Regression Discontinuity Estimator. *Review of Economic Studies*, 79(3):933–959.
- Justiniano, A., Primiceri, G. E., and Tambalotti, A. (2015). Household leveraging and deleveraging. *Review of Economic Dynamics*, 18(1):3–20.
- Kaplan, G., Mitman, K., and Violante, G. L. (2016a). Consumption and House Prices in the Great Recession. Working Paper.
- Kaplan, G., Mitman, K., and Violante, G. L. (2016b). Non-Durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data. SSRN Scholarly Paper ID 2777320, Social Science Research Network, Rochester, NY.
- Kau, J. B., Keenan, D. C., and Taewon Kim (1993). Transaction costs, suboptimal termination and default probabilities. *Journal of the American Real Estate & Urban Economics Association*, 21(3):247.

- Keys, B. J., Piskorski, T., Seru, A., and Yao, V. W. (2014). Mortgage Rates, Household Balance Sheets, and the Real Economy. SSRN Scholarly Paper ID 2505691, Social Science Research Network, Rochester, NY.
- Kruger, S. A. (2016). The Effect of Mortgage Securitization on Foreclosure and Modification. SSRN Scholarly Paper ID 2794122, Social Science Research Network, Rochester, NY.
- Lee, D. S. and Lemieux, T. (2010). Regression Discontinuity Designs in Economics. *Journal of Economic Literature*, 48(2):281–355.
- McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2):698–714.
- Mckay, A., Nakamura, E., and Steinsson, J. (2016). The Power of Forward Guidance Revisited. *American Economic Review*, 106(10):3133–3158.
- Mian, A., Rao, K., and Sufi, A. (2013). Household Balance Sheets, Consumption, and the Economic Slump. *The Quarterly Journal of Economics*, 128(4):1687–1726.
- Mian, A. and Sufi, A. (2010). Household Leverage and the Recession of 2007-09. *IMF Economic Review*, 58(1):74–117.
- Mian, A. and Sufi, A. (2014a). House Price Gains and U.S. Household Spending from 2002 to 2006. Working Paper 20152, National Bureau of Economic Research.
- Mian, A. and Sufi, A. (2014b). Why Tim Geithner is wrong on homeowner debt relief. The Washington Post, May 14, 2014.
- Mitman, K. (2016). Macroeconomic Effects of Bankruptcy and Foreclosure Policies. *American Economic Review*, 106(8):2219–2255.
- Narasiman, V. (2016). What do Adjustable-Rate Mortgage Resets Say About Household Consumption Behavior and the Income Channel of Monetary Policy? Working Paper, Harvard University.
- Palmer, C. (2015). Why Did So Many Subprime Borrowers Default During the Crisis: Loose Credit or Plummeting Prices? Working Paper, University of California at Berkeley.
- Poterba, J. and Sinai, T. (2008). Tax Expenditures for Owner-Occupied Housing: Deductions for Property Taxes and Mortgage Interest and the Exclusion of Imputed Rental Income. *American Economic Review*, 98(2):84–89.
- Romer, C. (2011). Policy Responses to the Great Recession: The Interaction of Leadership and Economic Ideas. Speech, Ohio State University, December 1, 2011.
- Scharlemann, T. and Shore, S. H. (2017). The Effect of Mortgage Payment Size on Default and Prepayment: Evidence from HAMP Resets. *working paper*.
- Scharlemann, T. C. and Shore, S. H. (2016). Does Reducing Negative Equity Prevent Mortgage Default? Evidence from HAMP’s Principal Reduction Alternative. *Review of Financial Studies*, forthcoming.
- Schelkle, T. (2016). Mortgage Default During the U.S. Mortgage Crisis. Working Paper, University of Cologne.

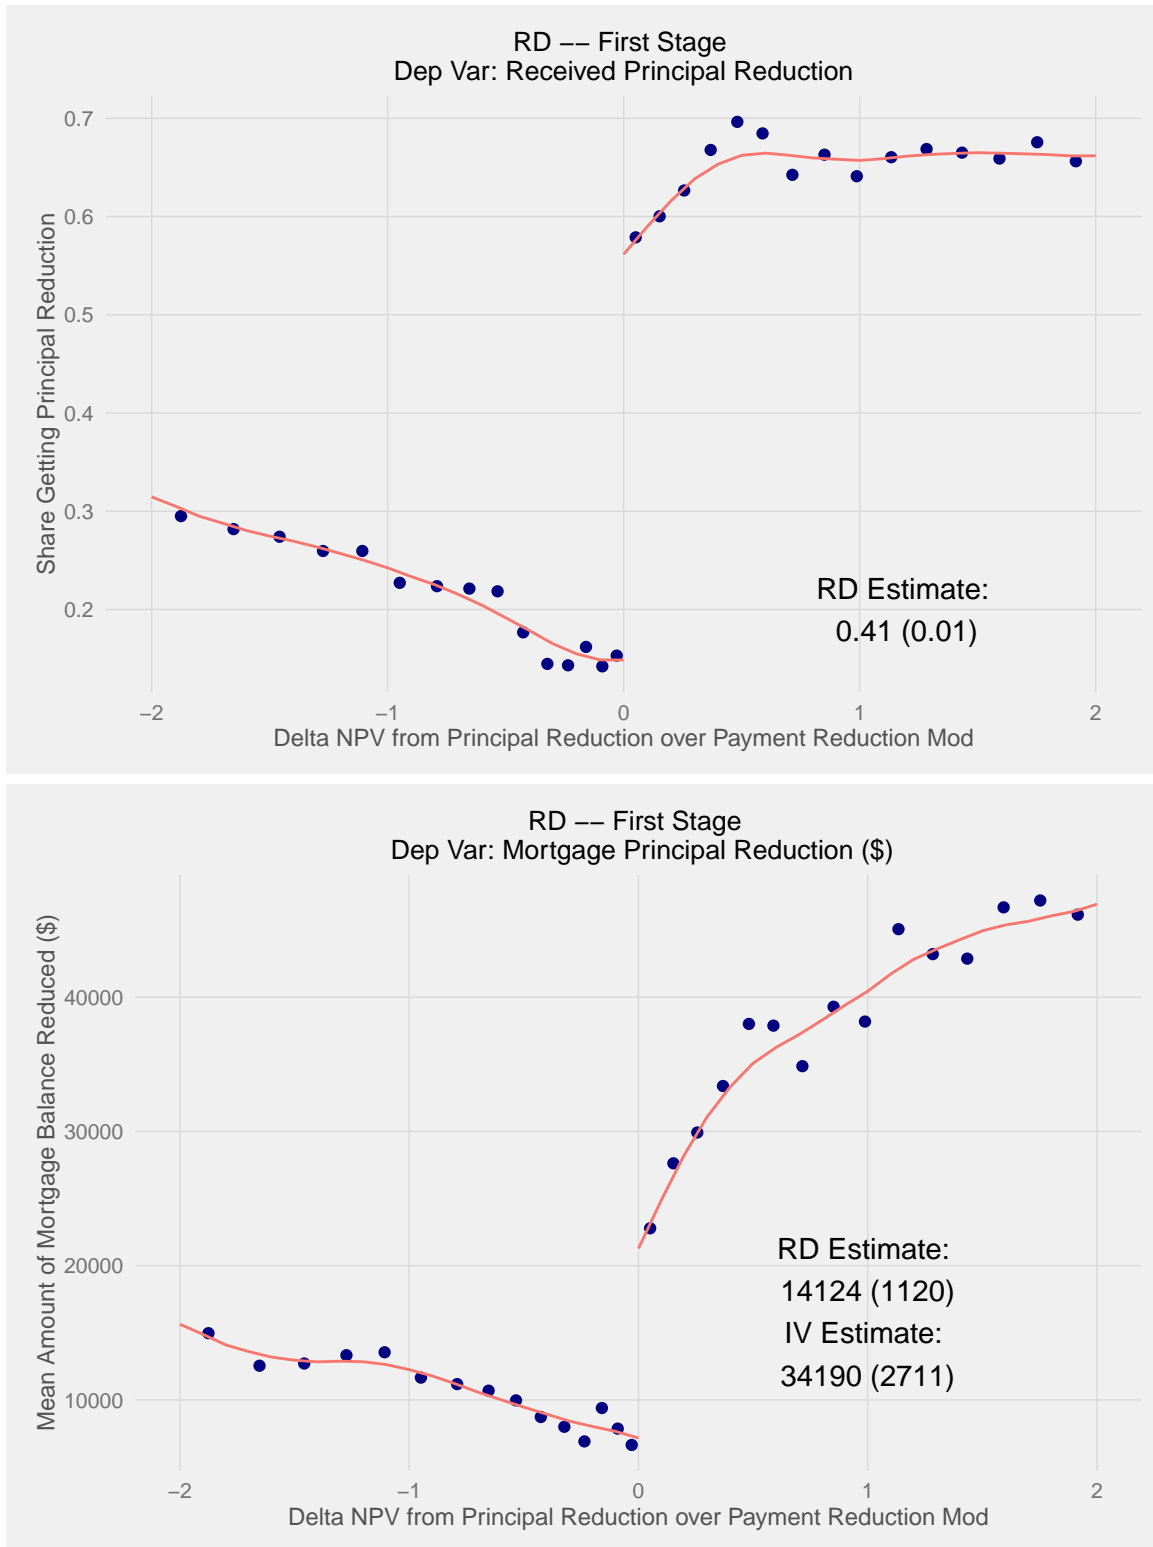
- Schwartz, M. and Wilson, E. (2008). Who Can Afford To Live in a Home? A Look at Data from the 2006 American Community Survey. Technical report, US Census Bureau 2008.
- Stanton, R. and Wallace, N. (1998). Mortgage Choice: What's the Point? *Real Estate Economics*, 26(2):173–205.
- Tracy, J. and Wright, J. (2012). Payment Changes and Default Risk: The Impact of Refinancing on Expected Credit Losses. *Federal Reserve Bank of New York Staff Reports*.
- Treasury Department, U. (2014a). Making Home Affordable Program Performance Report Through The Second Quarter of 2014. Treasury Department Report.
- Treasury Department, U. (2014b). MHA Data File User Guide. Treasury White Paper.
- Treasury Department, U. S. (2014c). *Making Home Affordable Program Handbook for Servicers of Non-GSE Mortgages*. Version 4.4.
- Winkler, R. (2010). Report Shows Strategic Defaults Increasing. Reuters, March 26, 2010.
- Zingales, L. (2010). The Menace of Strategic Default. City Journal, Spring 2010.

Figure 1: Comparing Payment Reduction and Principal Reduction Modifications



Notes: The top panel 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. Modifications consisted of temporarily reducing interest rates for five years, extending the mortgage term up to 40 years, turning a portion of the principal balance into a non-interest-bearing balloon payment, and for modifications with principal reduction, permanently forgiving a portion of the unpaid principal balance. The bottom panel summarizes the financial impacts of modifications along various dimensions.

Figure 2: Principal Forgiveness Around Regression Discontinuity Cutoff




Notes: This figure shows the first stage at the regression discontinuity cutoff in the matched HAMP credit bureau dataset. The top panel plots the share of borrowers receiving principal reduction and the bottom panel plots the average amount of principal reduction per borrower. The horizontal axis shows the normalized predicted gain to investors of providing principal reduction to borrowers from equation (1). The blue dots are conditional means for 15 bins on each side of the cutoff. The red line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (2). Construction of the IV estimate $\hat{\tau}$ in the bottom panel, as well as further details about these specifications, are described in section 4.2.

Figure 3: Initial Borrower and Mortgage Characteristics Around the Cutoff




Notes: This figure shows average pre-treatment characteristics around the regression discontinuity cutoff in the matched HAMP credit bureau dataset. The horizontal axis shows the normalized predicted gain to investors of providing principal reduction to borrowers from equation (1). The vertical axis in the first five panels shows borrower credit score, monthly income, the ratio of monthly mortgage payments to monthly income (debt to income ratio), the ratio of unpaid principal balance to the market value of the house (mark-to-market loan-to-value ratio), and the number of monthly mortgage payments the borrower is past due at application date. The final panel shows predicted default rates from a linear regression of default on the first five borrower characteristics. The red line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (2). See section 4.2 for details.

Figure 4: Impact of Principal Reduction on Default





Notes: This figure shows the estimated impact of principal reduction on default using the regression discontinuity design in the matched HAMP credit bureau dataset. The top panel plots the default rate on the vertical axis and the normalized predicted gain to investors of providing principal reduction on the horizontal axis. Default is defined as being 90 days delinquent at any point between modification date and March 2015, when our dataset ends, which is an average of three years. The blue dots are conditional means for 15 bins on each side of the cutoff. The red line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. Construction of the IV estimate $\hat{\tau}$ is described in section 4.2. The bottom panel plots the IV estimate and associated 95% confidence intervals for various alternative bandwidths. The Imbens and Kalyanaraman (2012) optimal bandwidth of 0.5 is plotted in yellow. The black horizontal line is the predicted impact of principal reduction on default from Treasury's redefault model.

Figure 5: Income Around Default in Bank Account Data



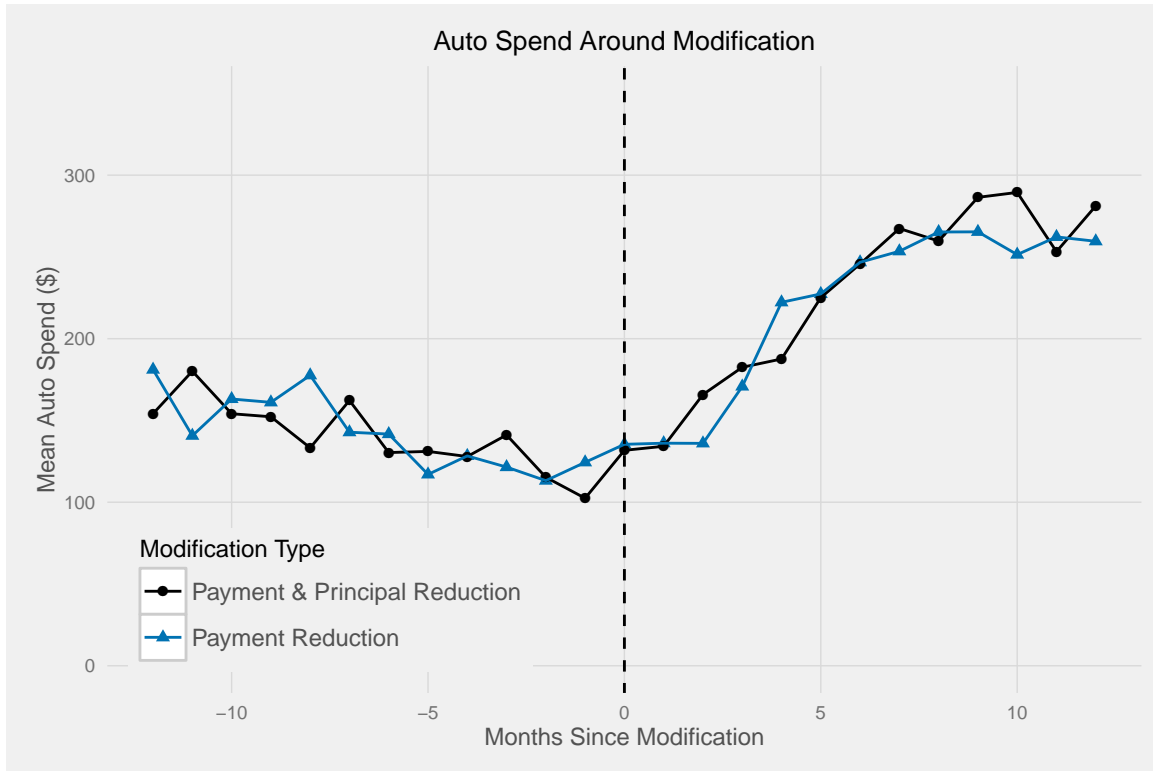
Notes: This figure plots average monthly income and mortgage payments around borrower default in the JPMCI bank account dataset. Income is defined as all inflows into a borrower's checking account in a given month. Default date is defined as the first month a loan is 90-days delinquent after a HAMP modification. The sample consists of 7,224 borrowers with a prior HAMP modification, a mortgage, and a checking account with Chase. For more details see section 3.2.

Figure 6: Credit Card Expenditure Around Modification



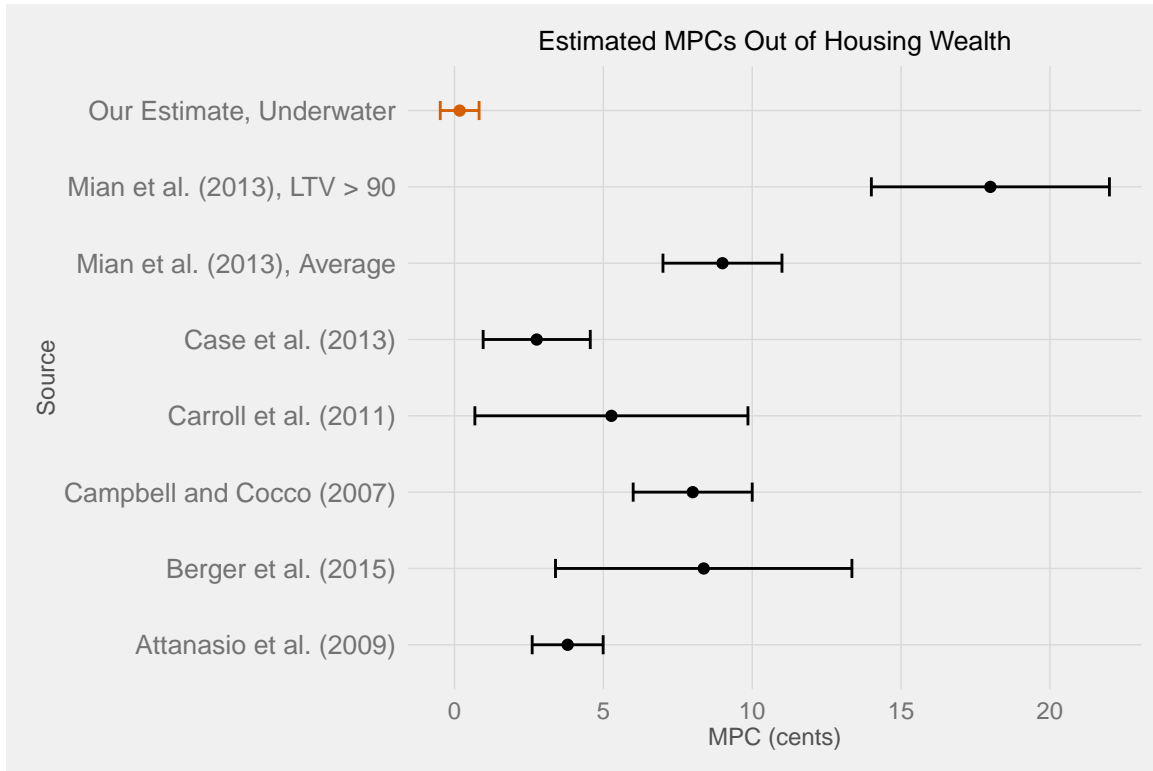
Notes: This figure shows the event study of monthly credit card expenditure around modification for borrowers receiving each type of modification in the matched HAMP credit bureau dataset. The top panel plots credit card expenditure in dollars as measured from credit bureau records (discussed in Section 3.1). The bottom panel plots spending relative to the month of modification (month zero). See Table 3 for sample summary statistics.

Figure 7: Auto Spending Around Modification



Notes: This figure shows the event study of monthly auto spending around modification for borrowers receiving each type of modification in the matched HAMP credit bureau dataset. Auto spending is measured from new auto loans, as described in Section 3.1. See Table 3 for sample summary statistics.

Figure 8: MPC Estimate Relative to Literature




Notes: This figure plots the point estimate and 95% confidence interval for the MPC out of housing wealth from various studies. All studies report the MPC for the average borrower unless otherwise indicated. Several studies which report elasticities have been converted into MPCs by scaling the elasticity by the ratio of Census retail sales to household housing assets from the Federal Reserve's Flow of Funds in 2012, which is 0.25. Case et al. (2013) do not report a central estimate, so the range presented is the range of estimates in Tables 3 through 6, and the point estimate is the mid-point of this range.

Figure 9: Default and Mortgage Debt Levels



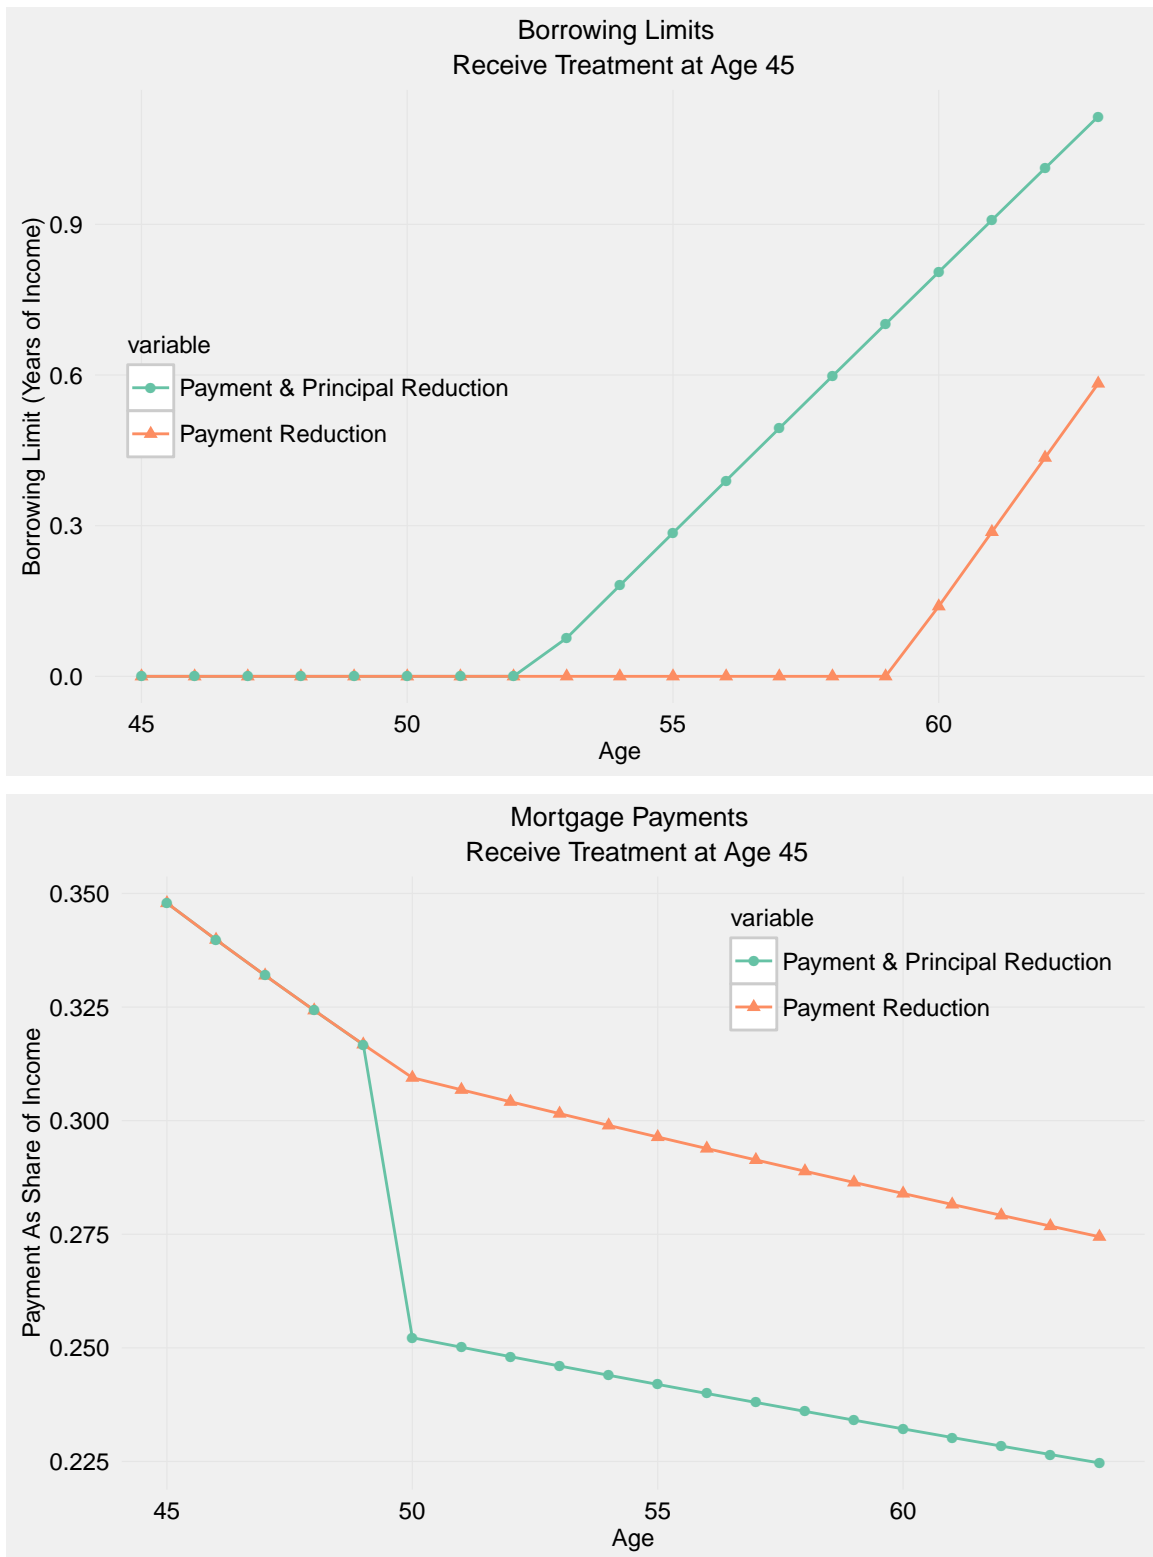
Notes: The top panel plots the cutoff thresholds for borrower default decisions under alternative baseline mortgage payment assumptions. The vertical axis is relative to permanent income. The green line shows the baseline assumptions as described in equation (4). For borrowers with a given LTV ratio, the line shows the cash-on-hand (income plus assets) threshold below which borrowers decide to exercise their default option. The red line shows the same threshold under an alternative assumption where mortgage payments are an extra 5% of income. The bottom panel plots default rates by LTV ratio. Default rates are calculated by taking the default thresholds shown in the top panel and integrating over the distribution of income shocks described in equation (5).

Figure 10: Consumption Functions With Wealth and Collateral Grants at Various Dates




Notes: The top panel plots the consumption function out of cash-on-hand under various alternative scenarios from the model described in Section 6. Both the horizontal and vertical axes are measured relative to permanent income. The baseline case considers a household with no home equity (and hence no current borrowing capacity). The orange, blue, and pink lines show the consumption functions in the current period when the household is granted one year's worth of permanent income in the current period, in one year, and in six years, respectively. The bottom panel shows the equivalent consumption functions for the case when the household is granted collateral, rather than wealth, at various dates.

Figure 11: Effect of Modeled Principal Reduction on Borrowing Limits and Mortgage Payments



Notes: This figure shows the effect of the modeled principal reduction policy on borrowing limits and mortgage payments. We assume homeowners receive modifications at age 45. We set initial LTV equal to 150. For our treatment group, we then reduce their mortgage balance by \$70,000, bringing them to an LTV of 106 in the first year. To mimic our empirical setting, mortgage payments for households who have not defaulted are fixed for five years, after which payments fall according to the new mortgage balance. This is discussed in more detail in Section 6.4.3.

Figure 12: Consumption and Mortgage Debt Levels



Notes: The top panel plots the consumption function out of principal reduction. We begin borrowers at 150 LTV and give increasing amounts of principal reduction as necessary to hit a given LTV ratio. To mimic our empirical setting, mortgage payments for households who have not defaulted are fixed for five years, after which payments fall according to the new mortgage balance. The red arrow shows the treatment for the average borrower in the government program. The dashed green line indicates the collateral constraint. The bottom panel plots the marginal propensity to consume out of an additional dollar of cash or an additional dollar of housing wealth generated by mortgage debt forgiveness. In this panel mortgage debt forgiveness is assumed to translate immediately into reduced mortgage payments. Both panels are plotted for homeowners age 45 with initial assets of at $a_t = 0.01$ units of permanent income and income of $y_t = 0.85$ units of permanent income.

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

	Pre-Match		Post-Match		Normalized Difference
	Mean	SD	Mean	SD	
Income	54564	24605	51432	23517	-0.13
Home Value	196992	123918	178825	114387	-0.15
Loan to Value Ratio	150	35	151	35	0.02
Monthly Mortgage Payment	1716	875	1558	785	-0.18
Monthly Payment to Income Ratio	0.48	0.12	0.47	0.12	-0.11
Mortgage Interest Rate	0.063	0.020	0.063	0.020	-0.00
Mortgage Term Remaining	25.9	4.6	25.9	4.6	-0.00
ARM (d)	0.49	0.50	0.46	0.50	-0.05
Months Past Due	11.4	12.9	9.9	11.7	-0.12
Credit Score	584	74	582	75	-0.03
Male (d)	0.57	0.50	0.56	0.50	-0.02
Age	48.8	10.8	48.6	10.9	-0.02
Monthly Payment Reduction (\$)	737	544	643	480	-0.17
Monthly Payment Reduction (%)	42	20	41	20	-0.07
Principal Forgiveness Amount	53072	70385	46221	62155	-0.10
Received Principal Forgiveness (d)	0.59	0.49	0.59	0.49	-0.01
Post Modification LTV	134	34	135	35	0.03
Post Modification DTI	0.30	0.04	0.30	0.04	-0.05
Post Modification Default (d)	0.201	0.401	0.203	0.402	0.00
N	222695		112176		

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

Table 2: Representativeness of Regression Discontinuity Sample

	RD Analysis Sample				PSID Delinquent Households			
	Mean	p10	p50	p90	Mean	p10	p50	p90
Income	58666	28424	56004	97071	64000	21000	55000	120000
Home Value	257703	100000	240000	440000	190000	50000	140000	350000
Loan to Value Ratio	126	101	119	163	101	52	94	166
Monthly Mortgage Payment	1807	800	1700	3000	1349	459	1100	2528
Mortgage Interest Rate	0.058	0.030	0.060	0.080	0.058	0	6.000	9.000
Mortgage Term Remaining	25.9	22.5	25.0	34.5	23.1	10	25.0	30.0
Months Past Due	8.5	0.0	6.0	21.0	5.0	2	3.0	11.5
Male (d)	0.60	0.00	1.00	1.00	0.68	0	1.00	1.00
Age	48.6	36.0	46.0	66.0	43.2	31	42.5	57.0
Value of Liquid Assets					3238	0	250	5000
N	9783				190			

Notes: This table compares borrowers in our regression discontinuity sample in the matched HAMP credit bureau dataset to delinquent borrowers in the 2009 and 2011 PSID Supplements on Housing, Mortgage Distress, and Wealth Data as reported in Gerardi et al. (2015). The regression discontinuity sample includes borrowers with $|v| < 0.5$ of the cutoff (see Section 4.2 for details). All values are before modification. 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. Liquid assets include checking and savings account balances, money market funds, certificates of deposit, Treasury securities, and other government saving bonds. (d) indicates a dummy variable.

Table 3: Summary Statistics For Difference-in-Differences Analysis

	Payment Reduction		Payment and Principal Reduction	
	Mean	SD	Mean	SD
Principal Forgiveness Amount	11895	35148	81699	64969
NPV Payment Reduction	51231	53506	86512	67758
Monthly Payment Reduction (\$)	680	477	674	502
Monthly Payment Reduction (%)	38.4	18.2	41.7	21.2
Loan to Value Ratio	151	33	153	36
Post Modification LTV	149	34	122	29
Monthly Payment to Income Ratio	0.47	0.11	0.47	0.12
Post Modification DTI	0.31	0.03	0.30	0.05
Income	56055	23751	53122	23592
Credit Score	599	83	579	72
Home Value	207123	119230	176254	111238
Monthly Mortgage Payment	1739	802	1598	777
Mortgage Interest Rate	0.061	0.018	0.064	0.019
Mortgage Term Remaining	26.0	4.3	26.0	4.4
Male (d)	0.57	0.49	0.55	0.50
Age	48.2	11.2	48.7	10.8
N	35842		34250	

Notes: This table shows summary statistics for the matched HAMP credit bureau sample analyzed in the panel difference-in-differences research design discussed in Section 5. The sample includes underwater borrowers who are observed in the credit bureau records one year before and after modification and report positive credit card expenditure in at least one month during this window. All variables are before-modification unless otherwise noted. (d) indicates a dummy variable.

Table 4: Impact of Principal Reduction on Expenditure

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 investors 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 ($n_{borrower} = 70,092$). Asterisks denote significance levels (***=1%, **=5%, *=10%). See the text for additional detail on the specification, outcome measures, and sample.

Panel A: Credit Card Expenditure (\$/month)

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (Principal Reduction x Post)	0.011 (2.783)	0.102 (2.783)	-0.502 (2.826)	0.582 (2.931)	-0.767 (2.957)	0.799 (2.991)
MSA Fixed Effects		Yes				
Calendar Month Fixed Effects			Yes			
MSA by Calendar Month Fixed Effects				Yes	Yes	Yes
Controls					Yes	Yes
Controls x Post Interactions						Yes
Dependent Variable Mean	390.13	390.13	390.13	390.13	391.28	391.28
Observations	1,704,543	1,704,543	1,704,543	1,704,543	1,667,159	1,667,159
Adjusted R ²	0.002	0.014	0.005	0.012	0.059	0.059

Panel B: Auto Expenditure (\$/month)

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (Principal Reduction x Post)	8.350 (6.661)	8.364 (6.661)	7.131 (6.721)	8.072 (6.798)	7.987 (6.895)	4.737 (7.039)
MSA Fixed Effects		Yes				
Calendar Month Fixed Effects			Yes			
MSA by Calendar Month Fixed Effects				Yes	Yes	Yes
Controls					Yes	Yes
Controls x Post Interactions						Yes
Dependent Variable Mean	140.43	140.43	140.43	140.43	141.29	141.29
Observations	1,704,543	1,704,543	1,704,543	1,704,543	1,667,159	1,667,159
Adjusted R ²	0.001	0.001	0.001	0.002	0.003	0.003

Table 5: Impact of Principal Reduction on Credit Card Expenditure Using Bank Data

This table reports difference-in-differences estimates of the effect of principal reduction on credit card expenditure in the JPMCI bank dataset. The coefficient of interest, *Treatment*, is the estimated change in the difference between outcomes of mortgages receiving modifications with and without principal reduction during the year after modification. All specifications include fixed effects for modification type and months since modification. Controls include credit score, pre-modification loan characteristics (LTV, principal balance), property value, and LTV at origination. The sample includes all HAMP borrowers with a mortgage and a credit card with Chase who are observed one year before and after modification. The dependent variable mean is reported for borrowers receiving principal reduction modifications in the year before modification. Standard errors, in parentheses, are clustered at the borrower level ($n_{borrower} = 22,924$). Asterisks denote significance levels (***=1%, **=5%, *=10%). See the text for additional detail on the specification, outcome measures, and sample.

<i>Credit Card Expenditure (\$/month)</i>						
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (Principal Reduction x Post)	-1.156 (3.762)	-0.993 (3.865)	-3.083 (3.866)	-5.718 (4.258)	-3.565 (4.534)	-3.299 (4.632)
MSA Fixed Effects		Yes				
Calendar Month Fixed Effects			Yes			
MSA by Calendar Month Fixed Effects				Yes	Yes	Yes
Controls					Yes	Yes
Controls x Post Interactions						Yes
Dependent Variable Mean	172	171.49	172	171.49	180.06	180.06
Observations	549,398	516,988	549,398	516,988	476,086	476,086
Adjusted R ²	0.0004	0.018	0.002	0.003	0.039	0.04

Table 6: Baseline Model Parameter Values

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

Table 7: Housing Wealth MPC in Model and External Benchmarks

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

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

Table 8: MPC out of Principal Reduction in the Model


Scenario	MPC (cents)
Data	0.2
Model Parameterizations with Small Response	
Baseline Model	0.3
Low Cash-on-Hand	0.0
Age At Mod = 35	0.9
High discount rate ($\beta = 0.9$)	0.8
Low risk aversion ($\gamma = 2$)	0.9
Unused HELOCs	0.9
Model Parameterizations with Larger Response	
High Cash-on-Hand (PIH)	3.4
Collateral Constraint $\phi = 0$	4.8
Expected 5% House Price Growth	6.2
Expected 5% House Price Growth and $\phi = 0$	24.2
Alternative Policy Simulations	
Write Down to 90% LTV	1.0
Write Down to 90% LTV and $\phi = 0$	14.1

Notes: This table compares the MPC out of principal reduction in the model under alternative parameterizations to the MPC calculated in our data (discussed in Section 5.2). The “Baseline Model” corresponds to the parameterization shown in Table 6 and the modeling of principal reduction policy discussed in Section 6.4.3. “Low Cash-on-Hand” corresponds to initial cash-on-hand $m_t = 0.5$ units of permanent income. The “Unused HELOCs” row corresponds to an experiment where the household is given a credit line worth \$20,000 (or 0.38 units of permanent income), and then given principal reduction. The “High Cash-on-Hand (PIH)” row corresponds to initial cash-on-hand $m_t = 3.0$ units of permanent income. The “Expected 5% House Price Growth” row corresponds to an expected permanent annual real house price growth of 5%.

Appendix

A Additional Figures

Figure A.1: Structure of Mortgage Modifications



Notes: This figure plots the average annual payments owed under the status quo and under the modified mortgage contracts for borrowers receiving “payment reduction” modifications in the matched HAMP credit bureau dataset. This is shown for the sample in borrowers in our consumption analysis, described in Table 3.


Figure A.2: Modification Terms Summary

C. Interest at the rate of 2.0 % will begin to accrue on the New Principal Balance as of 1/1/2012 and the first new monthly payment on the New Principal Balance will be due on 1/15/2012. My payment schedule for the modified Loan is as follows:

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


Notes: This figure shows the modified payment terms as explained to borrowers in the modification agreement which they are required to sign. Example terms are shown for a mortgage with a post-modification principal balance of \$300,00, temporary interest rate of 2%, mortgage term of 35 years, and escrow payments equal to 1.5% of the property value (\$250,000).

Figure A.3: Match Rate Around Regression Discontinuity Cutoff



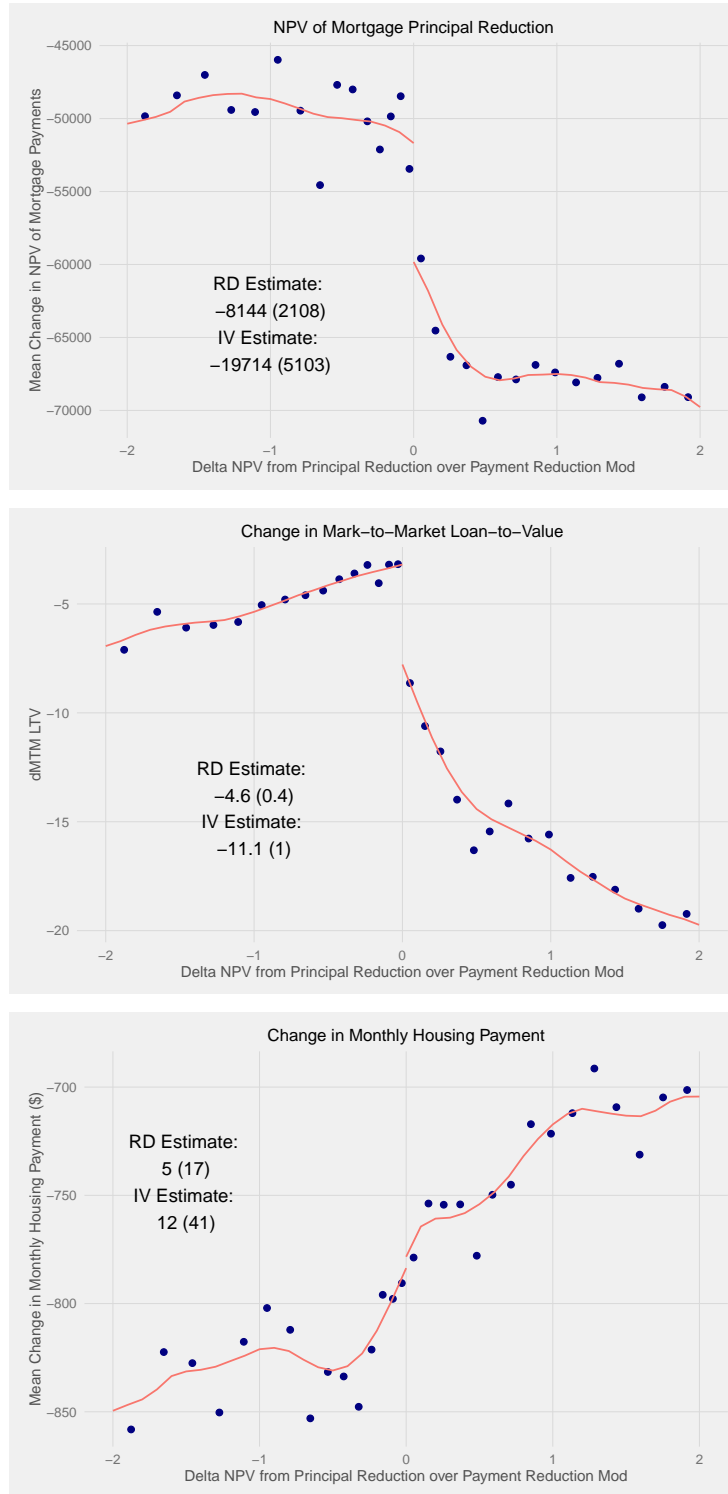
Notes: This figure plots the share of borrowers in the Treasury HAMP dataset successfully matched to their credit bureau records. The horizontal axis shows the normalized predicted gain to investors of providing principal reduction to borrowers from equation (1). The blue dots are conditional means for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

Figure A.4: Mortgage Delinquency Over Time



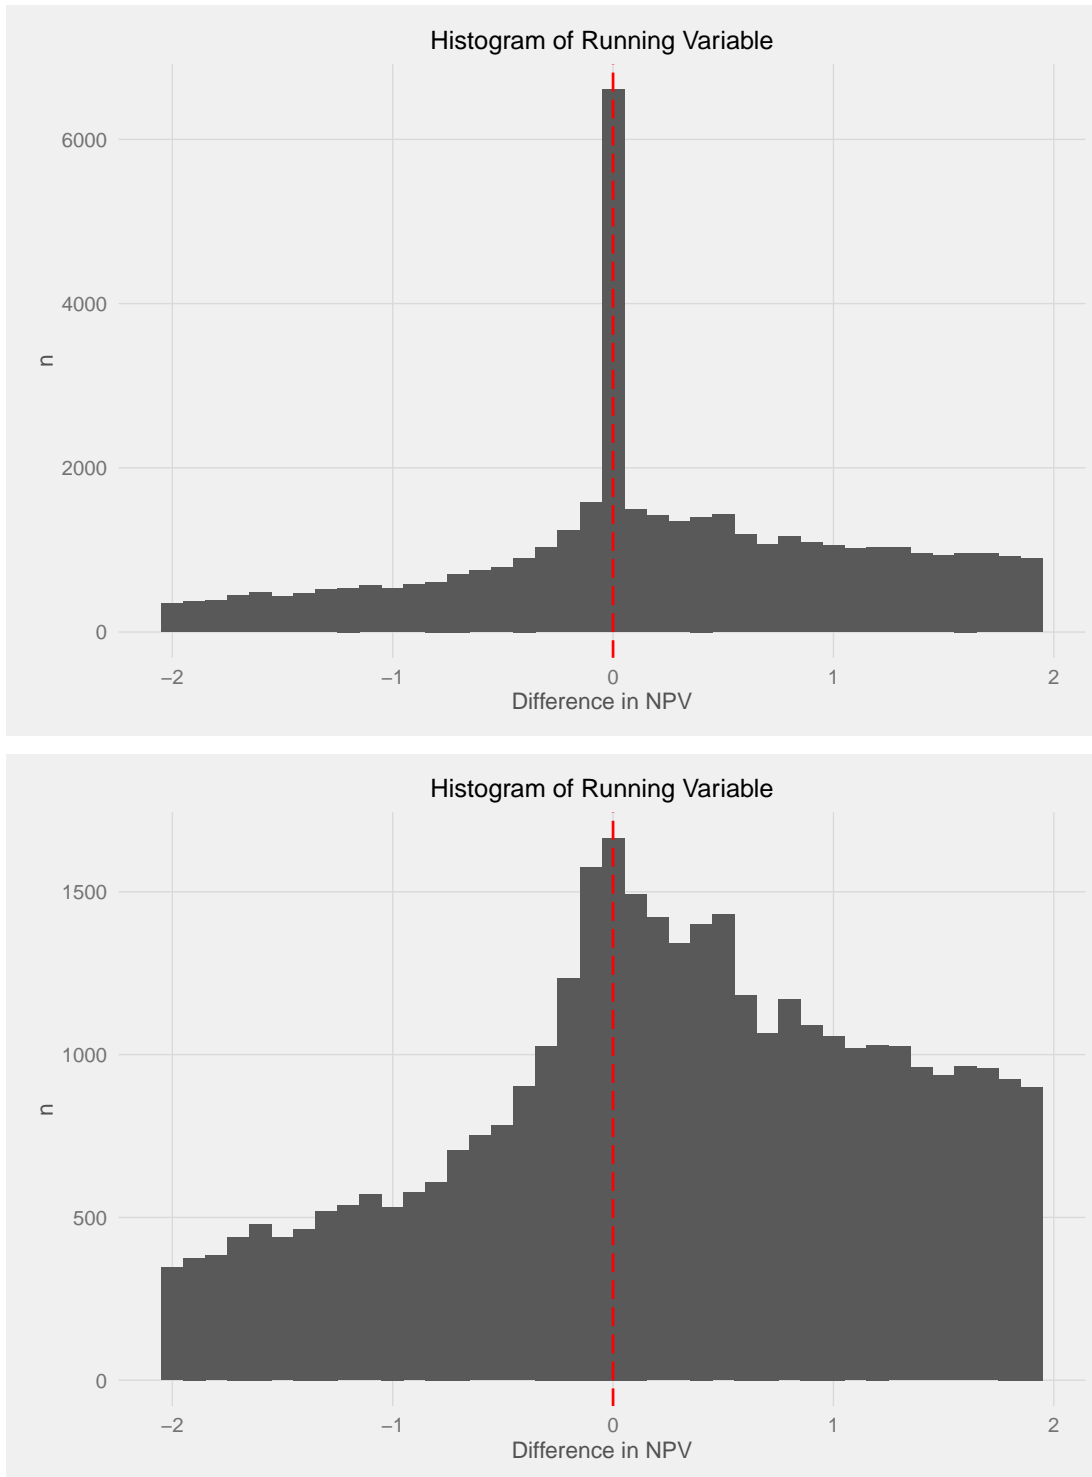
Notes: This figure plots the share of U.S. residential mortgages more than 30 days delinquent as reported by the Federal Reserve Board. The purple shaded region denotes the period where borrowers in our regression discontinuity sample had their first pre-modification delinquencies.

Figure A.5: Treatment Size Around Regression Discontinuity Cutoff



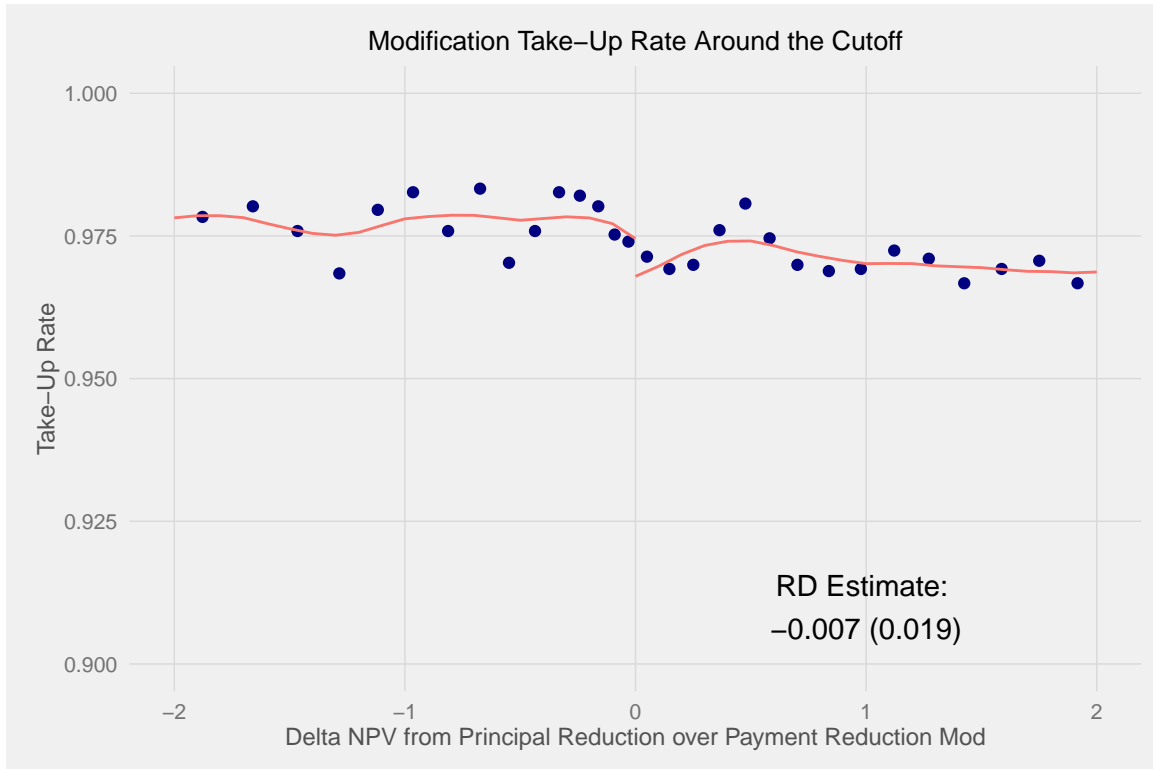
Notes: This figure shows the treatment size at the regression discontinuity cutoff in the matched HAMP credit bureau dataset. The horizontal axis shows the normalized predicted gain to investors of providing principal reduction to borrowers from equation (1). The top panel shows the change in the net present value of mortgage payments owed under the modified contract relative to the status quo mortgage contract, discounted at a 4% interest rate. The middle panel shows the change in the loan-to-value ratio. The bottom panel shows the change in initial monthly housing payments. The blue dots are conditional means for 15 bins on each side of the cutoff. The red line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (2). Construction of the IV estimate $\hat{\tau}$ is described in section 4.2.

Figure A.6: Borrower Density Around Regression Discontinuity Cutoff




Notes: The top panel plots the histogram of the running variable from our regression discontinuity strategy in the matched HAMP credit bureau dataset. The horizontal axis shows the normalized predicted gain to investors of providing principal reduction to borrowers from equation (1). HAMP program officers in the U.S. Treasury Department explain that the mass at exactly zero is due to data misreporting. Some servicers reported a single number as the calculation for both the payment reduction and principal reduction modifications, meaning that the estimated gains from principal reduction were calculated to be zero. The bottom panel plots the same histogram dropping observations exactly at zero, which is our analysis sample. Section 4.3 discusses three additional arguments for why the mass at zero is unlikely to pose a challenge for the validity of the regression discontinuity research design.

Figure A.7: Takeup Rate Around the Regression Discontinuity Cutoff



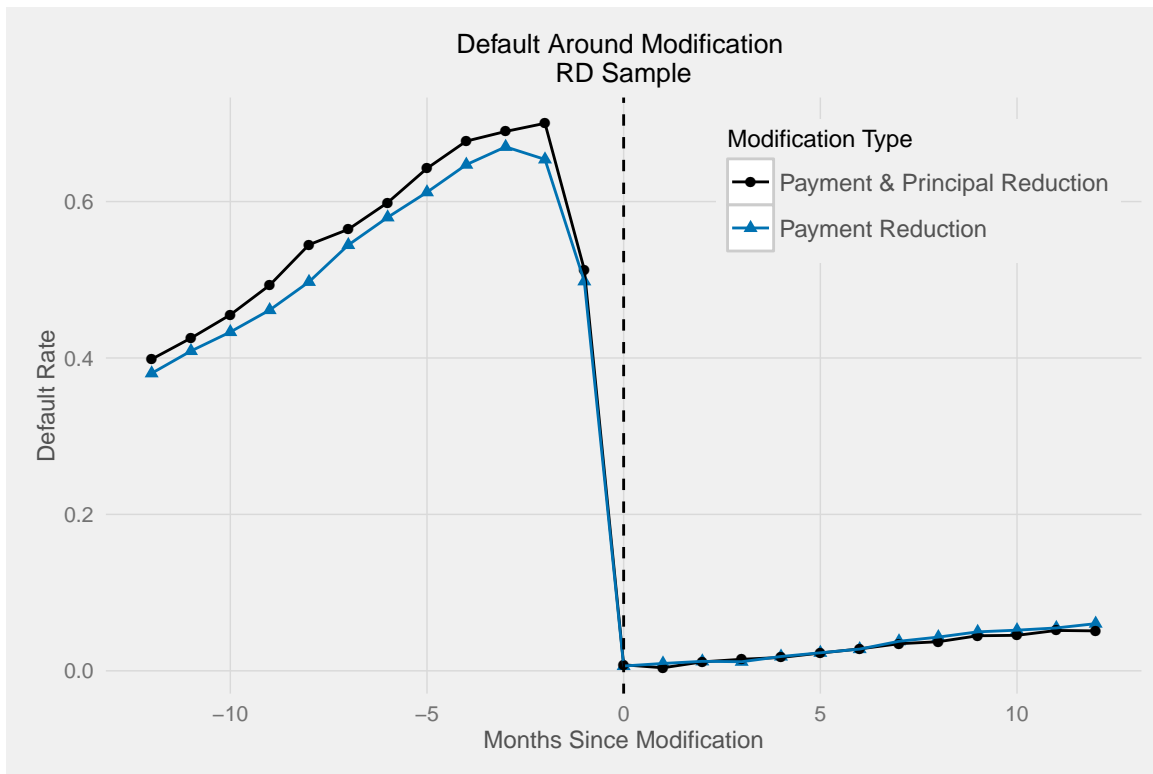
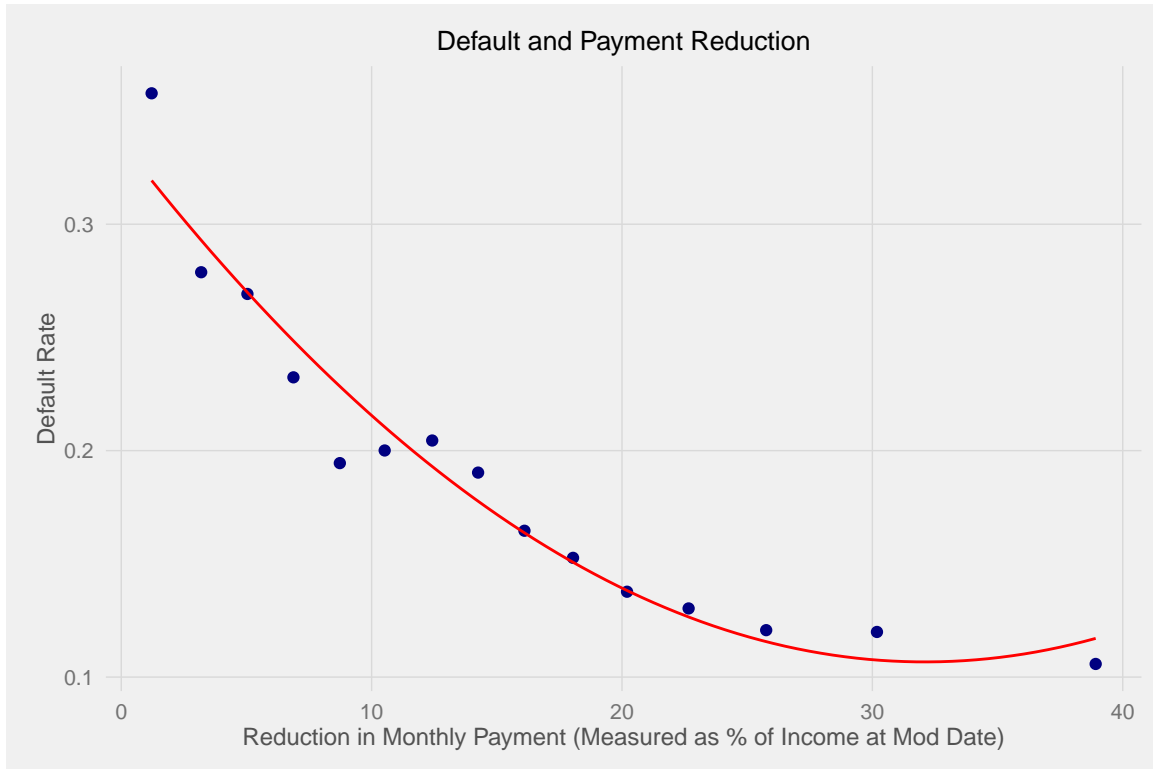
Notes: This figure shows the take-up rate conditional on borrower's being offered a modification in the Treasury HAMP dataset. The horizontal axis shows the normalized predicted gain to investors of providing principal reduction to borrowers from equation (1). The blue dots are conditional means for 15 bins on each side of the cutoff. The red line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (2).

Figure A.8: Regression Discontinuity Robustness



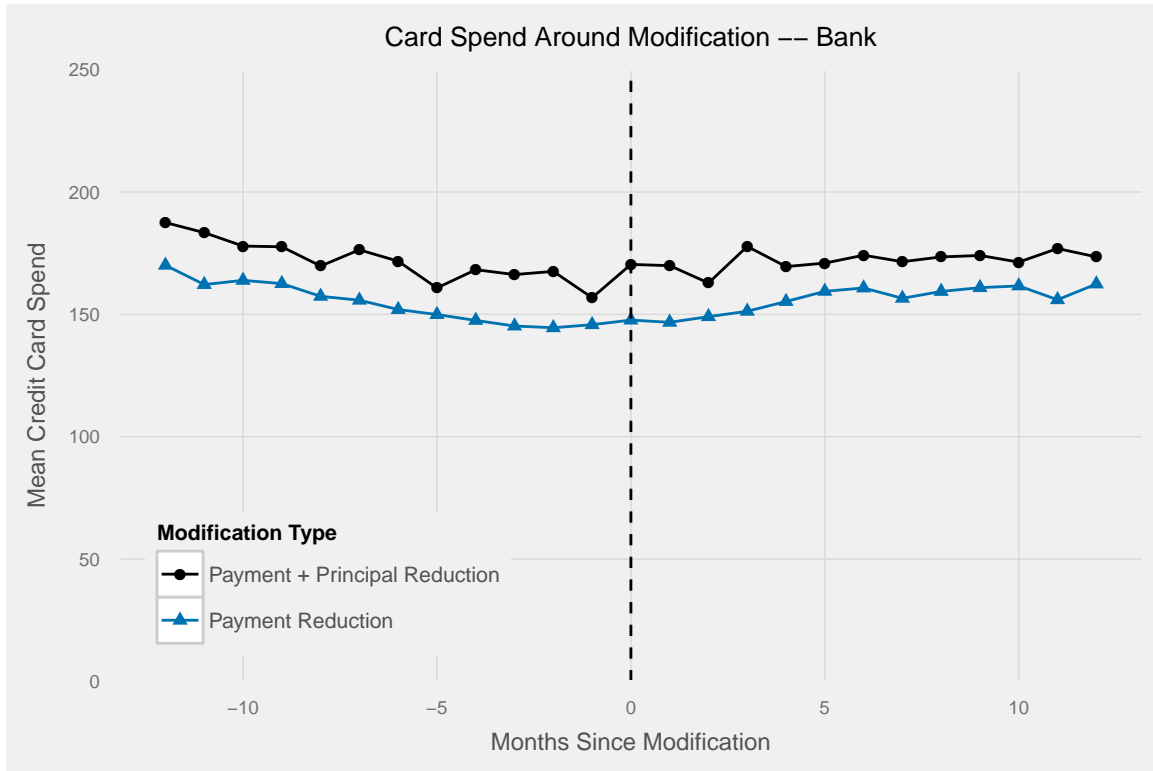
Notes: This figure shows the estimated impact of principal reduction on default under various specifications and bandwidths in the matched HAMP credit bureau dataset. Each line plots the IV estimate and associated 95% confidence interval from a local linear or quadratic regression on either side of the cutoff. The optimal bandwidths for the linear specification from Imbens and Kalyanaraman (2012) and Calonico et al. (2014) are 0.5 and 0.4, respectively. The optimal bandwidths for a quadratic specification from Imbens and Kalyanaraman (2012) and Calonico et al. (2014) are 0.8 and 1.0, respectively. The black horizontal line is the predicted impact of principal reduction on default from Treasury’s redefault model.

Figure A.9: Impact of Liquidity on Default



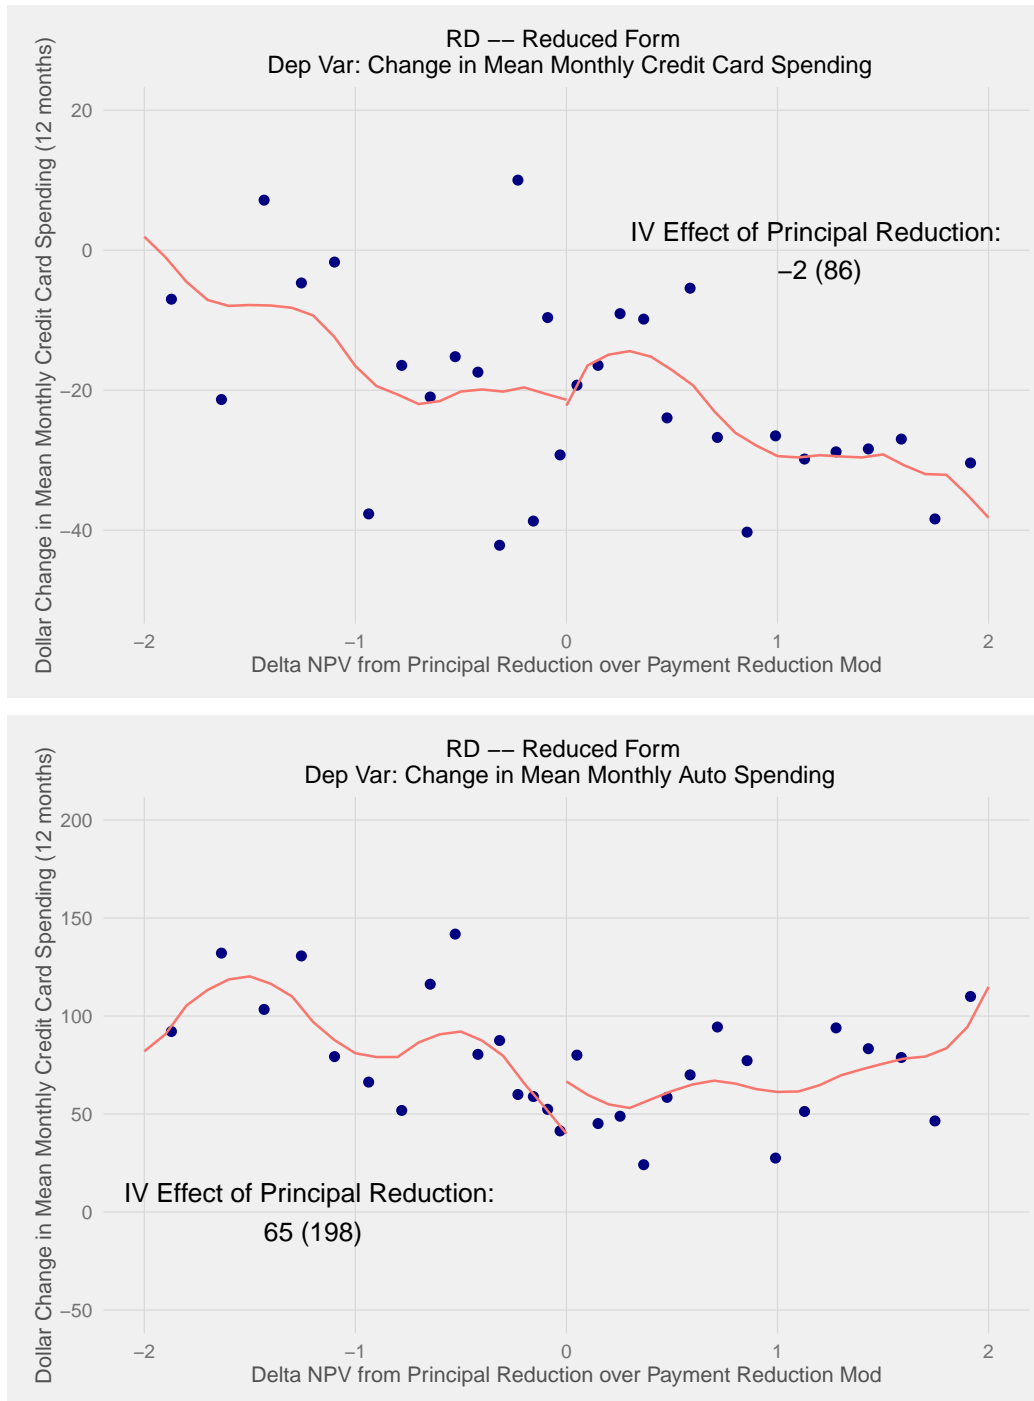
Notes: The top panel plots the cross-sectional relationship between change in a borrower's debt-to-income (DTI) ratio after modification and their probability of default. The bottom panel plots the event study of "default", defined as being 90 or more days behind on a mortgage payment, around date of modification. Modification mechanically brings borrowers "current". Both panels use data from the matched HAMP credit bureau dataset.

Figure A.10: Consumption Event Study Using Bank Data




Notes: This figure shows the event study of monthly credit card expenditure around modification for borrowers receiving each type of modification in the JPMCI bank account dataset. For further details see sections 3.2 and 5.2.

Figure A.11: Regression Discontinuity Design for Credit Card and Auto Expenditure





Notes: This figure shows the estimated impact of principal reduction on expenditure using the fuzzy regression discontinuity strategy in the matched HAMP credit bureau dataset. The horizontal axis shows the normalized predicted gain to investors of providing principal reduction to borrowers from equation (1). The vertical axis on the top panel shows the average change in credit card expenditure between the twelve months before modification and the twelve months after modification. Credit card expenditure is measured from credit bureau records as discussed in section 3.1. The vertical axis in the bottom panel shows the average change in auto spending between the twelve months before modification and the twelve months after modification. Auto spending is measured from new auto loans, as described in Section 3.1. The blue dots are conditional means for 15 bins on each side of the cutoff. The red line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. Construction of the IV estimate $\hat{\tau}$ is described in section 4.2.

Figure A.12: Credit Score Around Modification in Bank Data



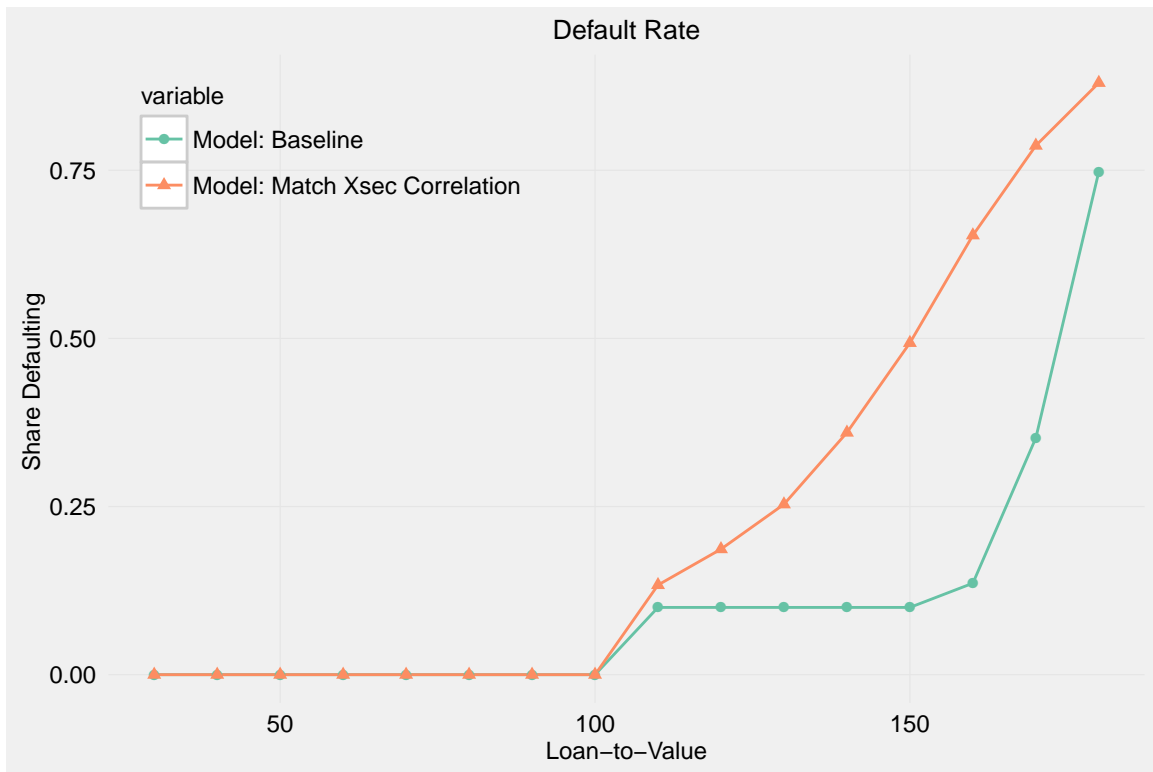
Notes: This figure shows the event study of credit scores for borrowers receiving each type of modification in the JPMCI bank dataset. For further details see sections 3.2 and 5.2.

Figure A.13: Mortgage Credit Availability



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

Figure A.14: Default With Heterogeneous Utility Cost of Default



Notes: This figure plots default rates by LTV ratio in the model under alternative parameterizations. The LTV is moved according to the same policy exercise described in the notes to Figure 11. The baseline parameterization corresponds to that in Table 6. The “Match Xsec Correlation” assumes a distribution of default costs across the population instead of a constant default cost.

B Effect of Payment Reduction on Default

In this section we analyze the effect of payment reduction on borrower default. Using a regression discontinuity empirical strategy, we find that immediate payment reduction with no change in long-term obligations significantly reduces default. Our findings indicate that additional payment reduction is effective irrespective of a borrower’s payment-to-income ratio, which is contrary to the view of affordability that has defined government housing policy since 1937. Furthermore, our findings suggest that the arbitrary income-ratio target embodied in the government program crowded out private bilateral modifications that could have been more generous, and led to fewer defaults, for up to 40% of HAMP borrowers. Finally, we show that by extending mortgage maturities, modifications can be structured so as to reduce defaults in a way that is likely to leave borrowers, investors, and taxpayers better off.

B.1 Variation in Payment Reduction Between HAMP and Chase Modifications

We analyze the effect of immediate payment reduction by comparing borrowers with HAMP modifications to those with proprietary modifications offered by Chase. Although servicers were required to offer all eligible borrowers HAMP modifications, HAMP had strict eligibility rules as described in section 2.1. Borrowers ineligible for HAMP were often eligible for alternative proprietary modification programs offered by their servicer. We focus in particular on proprietary modifications offered by Chase. This section discusses two important features of the Chase modification program which allow us to carefully study the effect of modifications that reduce immediate payments without reducing long-term obligations, while the next section describes the specific HAMP eligibility rule we rely on in our empirical strategy.

The design of HAMP and the Chase modification program reflect different views about the most effective way to reduce defaults. HAMP was designed with an explicit 31% *payment-to-income (PTI) ratio target*. This income ratio target evolved from the National Housing Act of 1937, which established a PTI limit in the federal government’s public housing program (Schwartz and Wilson 2008). Over the subsequent 80 years, the target spread throughout federal housing policies, first in the rental market, then in the owner-occupied market, and finally in the design of mortgage modifications during the Great Recession. 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 default than borrowers with lower PTI ratios, and that borrowers with PTI ratios below 31% already have “affordable” mortgages and hence do not need assistance to avoid default.

The first important feature of the Chase modification program is that, in contrast to the rigid income ratio target in HAMP, Chase joined a wide range of servicers and investors in the housing market (including Fannie Mae and Freddie Mac) in pursuing a strategy which added a *payment-reduction target*.⁵¹ 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

⁵¹For more detail about the evolution of different modification strategies, see Cowley (2017).

an “affordable” share of income according to the government’s metric. Reflecting this view, the Chase program mimics HAMP in modifying loans such that new payments hit the 31% PTI target, but adds the stipulation that if this results in a payment reduction of less than 30%, that the modification instead targets a 30% payment reduction.⁵² Under this rule, Chase modifications result in immediate payment reductions that are at least as great as in HAMP, and greater for borrowers with low PTI ratios who are either ineligible for HAMP or receive less than 30% payment reductions in HAMP.

The second important feature of the Chase modification program is that maturity extension was used 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 2.1, HAMP reduced payments by first providing interest rate reductions, then maturity extensions, and finally principal forbearance as necessary to achieve the 31% PTI ratio target. In contrast, Chase modifications targeted larger payment reductions by first extending maturity and amortization terms (which we call maturity extension for simplicity).⁵³ In this way deeper immediate payment reductions are offset by continued payments in the long-term. The structure of Chase modifications can be seen visually in Appendix Figure B.3, which shows the average path of payments owed before and after modification.⁵⁴ Low payments are sustained throughout the early mortgage term (since few borrowers have interest rates re-setting from temporary reductions), and continue for up to 40 years (since many borrowers have extended maturities). We compare cash-flows for various modification groups in more detail in section B.3.

B.2 31% Payment-to-Income Regression Discontinuity Empirical Strategy

Our strategy to identify the effect of immediate payment reduction on default uses variation generated by HAMP’s 31% PTI eligibility cutoff and the different payment reduction targets used in HAMP and the Chase modification program. Since HAMP was designed to reduce monthly payments to 31% of a borrower’s income, borrowers whose PTI ratio was already below 31% were ineligible for HAMP. These borrowers were only evaluated for Chase’s proprietary modification program.

The difference in modification program rules generates substantial variation in the amount of

⁵²The Chase program specifically targets a 25% reduction in principal, interest, taxes, and insurance (PITI) payments, which translates to a 30% principal and interest reduction for most borrowers. We follow the convention in the prior literature and characterize the reductions relative to prior principal and interest payments.

⁵³Our language follows prior work by ? who discuss maturity choices in the auto loan market in a context where amortization and maturity terms are equivalent. In both HAMP and the Chase modification program servicers extended mortgage maturities and re-amortized the mortgage to the new term, unless either step was prohibited by investor agreements (in which case only the permitted step, if any, was pursued). In theoretical work analyzing consumer debt restructuring policies ? use “maturity extension” to mean something different. They consider a modification with interest-only payments for a number of periods, with a corresponding increase in loan maturity and no change in the amortization schedule during the periods when principal payments are made.

⁵⁴While the HAMP program ended in December 2016, the Chase modification program described in this section is ongoing and its rules are subject to change. Our description of the program reflects guidelines in place during our analysis period.

immediate payment reduction received by borrowers on either side of HAMP’s 31% PTI eligibility cutoff. Above the cutoff borrowers receive a mix of HAMP modifications with small payment reductions (since these borrowers are already close to HAMP’s 31% PTI target) and Chase modifications with large payment reductions (since Chase had a minimum 30% payment reduction target). Below the cutoff, all borrowers receive Chase modifications with large payment reductions. The result is that the average payment reduction jumps sharply at the cutoff.

Employing a fuzzy regression discontinuity strategy, the treatment effect of receiving payment reduction is determined by the jump in default divided by the jump in the amount of payment reduction received at the cutoff. Similar to equation 2, let Y be the outcome variable of interest, and let ΔPay be the percent change in mortgage payment from the modification. The estimand for the effect of a 1% payment reduction is given by

$$\tau = \frac{\lim_{PTI \downarrow 31\%} E[Y|PTI] - \lim_{PTI \uparrow 31\%} E[Y|PTI]}{\lim_{PTI \downarrow 31\%} E[\Delta Pay|PTI] - \lim_{PTI \uparrow 31\%} E[\Delta Pay|PTI]}. \quad (7)$$

We follow the procedures described in section 4.2 to estimate $\hat{\tau}$ using borrowers in the JPMCI bank dataset. The full sample includes all non-GSE borrowers who received a modification through either the HAMP or Chase programs. We further limit the sample to modifications performed in the fourth quarter of 2011 or later, and who are observed at least two years after modification. The Chase proprietary program we study was instituted after HAMP was in place, and by the fourth quarter of 2011 it was sufficiently established such that any borrower ineligible for HAMP due to the PTI cutoff could plausibly receive a Chase modification.⁵⁵ The sample includes borrowers with PTI ratios between 25 and 80.⁵⁶ For symmetry around the 31% PTI cutoff, most figures in the following sections plot data for borrowers with PTI ratios between 25% and 37%. The optimal bandwidth according to the Imbens and Kalyanaraman (2012) procedure is 4.5 percentage points of PTI, though we report estimates for multiple bandwidths.

B.3 Internal Validity of Design

As with the investor NPV regression discontinuity strategy, the validity of this research design rests on two assumptions. First, there is a discontinuous jump in the amount of immediate payment reduction received at the cutoff. Second, all other variables that could affect outcomes are continuous at the cutoff.

The top panel of Figure B.1 validates the first assumption. It shows that borrowers below the cutoff receive payment reductions that are 17.4 percentage points more generous than those received by borrowers above the cutoff.⁵⁷ The figure also shows that payment reductions are approximately

⁵⁵Chase had a variety of other modification programs with different designs that preceded HAMP.

⁵⁶Borrower density thins above PTI ratios of 80, and borrowers with PTI ratios below 25 are evaluated according to different program rules.

⁵⁷This figure, and all subsequent figures in the section, remove the 275 observations with PTI ratios between 31% and 31.1%. Borrowers this close to the cutoff receive an amount of payment reduction which is about halfway between that received by borrowers clearly above and clearly below the cutoff. This can be seen in the top panel of Appendix Figure B.4. The bottom panel of this figure shows that the estimated effect of payment reduction on default is nearly

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. Above the cutoff, about half of borrowers receive deep payment reductions similar to households below the cutoff, while the other half face a 31% PTI ratio target and therefore receive virtually no payment reduction. Although borrowers above the cutoff who receive HAMP modifications with a 31% PTI ratio target would have received larger payment reductions in the Chase program, HAMP rules prohibited servicers from offering private sector alternatives to any HAMP-eligible borrower, under the assumption that private modifications wouldn't be as effective as those subsidized by the government.⁵⁸ Thus the only borrowers above the cutoff receiving Chase modifications are those failing to meet one of the other eligibility criteria described in section 2.1.⁵⁹ Our estimate is a local average treatment effect for the borrowers moved from HAMP to Chase modifications due to the 31% PTI cutoff.

Figure B.2 shows that the treatment we study is an immediate payment reduction with no reduction in long-term payments. Because much of the additional payment reduction received by borrowers below the cutoff is achieved via maturity extension, deeper immediate payment reductions are offset by continued payments further out in the mortgage term, such that the total NPV of the cash-flows owed under the mortgage contract is similar for borrowers just above and just below the cutoff. The top panel shows that payments for borrowers just below the cutoff are lower for the first 22 years, before rising relative to those of borrowers just above the cutoff. The bottom panel summarizes the financial impacts of the modifications for these two groups,⁶⁰ and reports that the NPV of the mortgage payments owed under the new mortgage contract is nearly identical for borrowers just above and just below the cutoff.⁶¹ Figure B.2 also shows that borrowers below the cutoff receive less generous principal reductions than those above the cutoff. This is because these borrowers are ineligible for the subsidized principal reduction in HAMP. Hence borrowers below the cutoff receive a relative reduction in immediate mortgage payments with a relative increase in mortgage principal, with no change in the NPV of payments owed under the full contract. Because the NPV of long-term payments is unchanged, and we found no independent impact of principal reductions on default in section 4.4, we will attribute the causal impact of this treatment on default to the immediate payment reduction portion of the treatment. In the next section we provide a

identical when these borrowers are included.

⁵⁸Although the government prevented servicers from offering private alternatives to HAMP-eligible borrowers, servicers could offer deeper payment reductions within HAMP. However they could only do so if they followed the priority of modification steps allowed within HAMP (i.e. servicers couldn't reduce payments of HAMP borrowers below 31% PTI by offering more maturity extension until they had reduced the interest rate all the way to 2%, which is more costly to the investor).

⁵⁹We do not observe the full set of eligibility criteria in our data and include all borrowers who received either modification type.

⁶⁰Appendix Figure B.5 reports the RD estimates for the jump in principal reduction and the NPV of payments owed at the cutoff.

⁶¹This NPV calculation only reflects the amount owed under the mortgage contract, and does not reflect any behavioral response such as different default probabilities or prepayment behavior. It uses a 4.1% discount rate, equal to the average market mortgage interest rate over the program period, which is the rate used to discount cash flows in the government-mandated model described in Section 2.2. A higher discount rate would lead to larger savings for borrowers below the cutoff relative to those above the cutoff, since savings were front-loaded for these borrowers.

bound on our estimate of the effect of payment reduction under the alternative assumption that the principal increase actually had an offsetting effect.

Appendix Figure B.6 provides evidence supporting the second assumption. The first four panels show the LTV ratio, months past due, borrower monthly income, and pre-modification monthly payment around the cutoff.⁶² In all cases these loan and borrower characteristics trend smoothly. In the bottom panel we use these observable borrower characteristics to predict default, and show that predicted default is also smooth at the cutoff. Finally, Appendix Figure B.7 shows that borrower density is also smooth around the cutoff.

B.4 Results: Effect of Payment Reduction on Default

We find that immediate payment reduction significantly reduces default rates. The bottom panel of Figure B.1 shows the reduced form of the fuzzy regression discontinuity specification, plotting the default rate on the y-axis. The figure provides visual evidence that payment reduction reduces default in two ways. First, the default rate falls sharply by 4.7 percentage points at the 31% PTI cutoff. Second, the slope of the estimated default rates in the bottom panel mirrors the slope of the payment reductions in the top panel; default rates are approximately constant on the left-hand side of the cutoff, consistent with a constant amount of payment reduction, and 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 Appendix Figure B.4, which plots the first stage and reduced form for borrowers with PTI ratios all the way out to 50%).

Our point estimate of $\hat{\tau}$ from equation 7 is that an extra 1% payment reduction reduces default rates in the two years after modification by 0.27 percentage points, or by 0.95% of the mean above the cutoff. Appendix Figure 4 tests the sensitivity of our results to the bandwidth chosen for the local linear regression. Our central estimates are constructed using the optimal bandwidth from the Imbens and Kalyanaraman (2012) procedure, which is 4.5 points of PTI. We test alternative bandwidths between 1 and 10 and find that the point estimate is stable. Furthermore, our point estimate is similar (rising to 1.2%) if we use the upper bound of our confidence interval from section 4.4 to adjust from any potential effect of the relative increase in mortgage principal at the cutoff.⁶³

Our results suggest that immediate liquidity provision is the key driver of the default reductions found in prior work analyzing the effect of sustained payment reductions. To our knowledge, this is

⁶²Unlike in the matched HAMP credit bureau dataset used for the investor NPV strategy, borrower credit score is not available in the JPMCI bank dataset.

⁶³ If we take the upper bound of our 95% confidence interval for the impact of principal reduction on default from section 4.4, and scale it by the amount of relative principal increase received by borrowers just below the 31% PTI cutoff, we find that a principal increase of this magnitude would have led to at most a 1.2 percentage point increase in default rates. If the payment reductions had to offset this effect, this would mean that the reduced form jump in default at the cutoff would have been 5.9 percentage points without the principal increase rather than 4.7 percentage points, or alternatively that each 1% reduction in payment reduced default rates by 0.34 percentage points (1.2%), similar to our baseline estimate of 0.27 percentage points (0.95%).

the first estimate of the effect of an immediate payment reduction with no reduction in long-term payments. Our results are similar to two types of estimates in the prior literature which analyzes sustained payment reductions. First, Agarwal et al. (2011) and Haughwout et al. (2016) analyze modifications provided to delinquent borrowers before the implementation of the HAMP program. They find that a 1% immediate payment reduction that also reduces long-term debt obligations is associated with a 1.1% to 1.3% reduction in default rates.⁶⁴ Second, other authors have analyzed the effect of sustained payment reductions for non-distressed borrowers. Although differences in borrower characteristics and baseline default rates makes it difficult to directly compare magnitudes between distressed and non-distressed borrowers, this literature has found that a 1% payment reduction is associated with a 1.1% to 2.75% reduction in default rates.⁶⁵ The similarity between our estimates of the effect of immediate payment reduction and those of the prior literature analyzing the effect of sustained payment reductions suggests that immediate liquidity provision is a main driver of the default reductions documented in previous work.

B.5 Implications

Our finding that the type of payment reduction we study can substantially reduce defaults has three broader implications. This section discusses each implication in turn.

First, our findings reject the homogeneous-preferences view of housing affordability that has been embodied in government housing policy since 1937. A one-size-fits-all PTI ratio target could be justified by a model where agents have Cobb-Douglas preferences over housing and non-housing consumption and all have identical tastes for housing.⁶⁶ Such a model predicts that all households will choose to spend the same share of their income on housing. Payments above this share would lead to defaults while payments below this share would be sustainable. With a 31% PTI target, this view implies that default rates should be converging to zero as average PTI approaches 31%. In contrast, we show that default rates are well above zero for households below this cutoff.

Instead, our results are more consistent with a model with heterogeneous tastes for housing where liquidity shocks are the main drivers of default. A model with heterogeneous preferences for housing would suggest that different households have different preferred income shares of housing.

⁶⁴In addition, Agarwal et al. (2016a) show that HAMP modifications with both short and long-term payment reduction led to significant reductions in default rates, though they do not estimate an elasticity with respect to a given quantity of payment reduction.

⁶⁵See Ehrlich and Perry (2015), Tracy and Wright (2012), Fuster and Willen (2017), and DiMaggio et al. (2017). The latter three of these papers analyze the effect of payment reductions caused by downward adjustments of interest rates for borrowers with adjustable rate mortgages. As Fuster and Willen (2017) discuss, to the extent that borrowers are aware of their mortgage terms and follow the movement of underlying index rates, then these estimates may already be capturing only the liquidity effects of lower payments since borrowers would already have been anticipating and responding to the long-term payment reductions. Similarly, in contemporaneous work Scharlemann and Shore (2017) analyze the effect of sustained payment increases for non-distressed borrowers in a setting where these payment increases are pre-determined five years ahead of time (resulting from a step up in interest rates as part of a previous mortgage modification). They find that a 1% increase in payments is associated with a 0.8% reduction in default rates.

⁶⁶An alternative justification specific to the owner-occupied market is the value of PTI limits from a macroprudential policy perspective. Greenwald (2017) finds that a cap on PTI ratios is more effective than a cap on LTV ratios at limiting boom-bust cycles.

Households could support a wide range of potential PTI ratios, but temporary income or expense shocks would induce them to default even at potentially low PTI ratios. In these cases, a payment reduction would plausibly allow them to maintain payments. Our empirical results are consistent both with liquidity-shock driven defaults (section 4.6) and the effectiveness of payment reductions even for low-PTI borrowers. We find that payment reductions led to a substantial reduction in default rates, even for borrowers already at the 31% PTI ratio target who supposedly had “affordable” mortgages according to the government’s rule. Furthermore, for borrowers above the 31% cutoff, those who started with higher PTI ratios and thus received more payment reduction had lower default rates than those starting with lower PTI ratios, even though HAMP reduced the PTI ratio of both groups to the same 31% target.

Second, our data showing that the government’s rigid income ratio target generated less generous modifications than private sector alternatives, coupled with the government’s requirement that borrowers be offered HAMP before private alternatives, raise concerns about substantial crowd-out which may have had the unintended effect of increasing defaults for up to 40% of HAMP borrowers. The 31% PTI target in HAMP means that borrowers with pre-modification PTI ratios below 41.5% received payment reductions of less than 30%. For these borrowers, an alternative private bilateral modification program such as the one we studied would have resulted in more generous immediate payment reduction and, according to our empirical results, lower default rates without requiring any government intervention or subsidy. About 40% of all borrowers with HAMP modifications (625,000 borrowers) were in this region, and the government spent approximately \$7 billion subsidizing potentially less effective modifications for these borrowers.⁶⁷ Furthermore, the government explicitly prohibited servicers from offering alternative private modifications to any HAMP-eligible borrower under the assumption that the private modifications wouldn’t be as generous or effective as those subsidized by the government. Our data on proprietary modifications and the HAMP-first requirement indicate that HAMP crowded out private modifications that could have been more effective for a large fraction of borrowers in the government program.⁶⁸

Third, our results imply that since the amount of immediate payment reduction appears more important than the way that payment reduction is achieved, maturity extension can be used to reduce short-term defaults in a way that is likely to leave borrowers, investors, and taxpayers better off. To understand why this is true, it is useful to revisit the particular structure of mortgage modifications around the cutoff we study. Because Chase modifications achieved larger payment

⁶⁷This figure includes subsidies paid to investors, servicers, and borrowers. See Treasury Department (2014c) for a detailed breakdown of these subsidies.

⁶⁸While our emphasis is on crowd-out on the intensive margin of the quality of modifications, Agarwal et al. (2016a) provide evidence on the extensive margin in terms of the quantity of modifications. They show that each HAMP modification is associated with about 0.84 more “net” modifications, suggesting that the program did not substantially reduce the total number of private modifications that would have been offered in the absence of the program. They argue that by inducing more borrowers to apply for modifications, HAMP increased the private modification rate of HAMP-ineligible borrowers, offsetting the reduced private modification rate of HAMP-eligible borrowers. Our evidence suggests that at least some of the borrowers pushed from private to HAMP modifications would have received larger payment reductions, and had lower short-term default rates, if they had remained in private modifications designed to emphasize payment reduction rather than an income-ratio target.

reductions primarily by extending mortgage maturities, relative to borrowers above the 31% PTI treatment discontinuity, those below the cutoff traded more generous immediate payment reduction for continued payments in the long-term. A maturity extension with no other changes to the mortgage contract is equivalent to the investor “lending” the borrower their monthly payment reductions at the mortgage interest rate, with repayments on this “new” loan beginning at the end of the original mortgage term and continuing through the end of the new loan term.

This particular form of loan restructuring is likely to leave borrowers, investors, and taxpayers better off. For borrowers, a maturity extension which moves their funds from the future to the present might be particularly valuable in periods of acute economic distress. In the short term, lower default rates indicate that more borrowers will remain in this contract, which by revealed preference indicates that they found it more attractive than the contract with less generous payment reductions. Borrowers may also be better off in the long-term because payments are still lower on average 22 years after modification, by which point borrowers will likely have regained positive home equity.⁶⁹ It is thus likely that foreclosures are permanently lower under this modification type. Investors will also be better off as long as the reduction in default rates offsets any increase in prepayment risk and any potential reduction in the net present value of the cash-flows owed on the mortgage contract.⁷⁰ Finally, the government spent substantial resources subsidizing HAMP modifications above the 31% PTI cutoff with small payment reductions and high default rates, whereas borrowers below the cutoff received proprietary modifications requiring no government assistance which had large payment reductions and low default rates.

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

Our empirical results about ex-post debt restructuring help inform the theoretical debate about optimal ex-ante mortgage design. A number of recent papers have analyzed alternative mortgage contracts with built-in features designed to assist households overcome periods of financial distress. Eberly and Krishnamurthy (2014) propose a fixed rate mortgage (FRM) with a one-time option to convert to an adjustable-rate mortgage (ARM). Guren et al. (2017) develop an equilibrium model of the housing market and compare the Eberly-Krishnamurthy convertible mortgage contract to a FRM with an option to costlessly refinance to another FRM. They find that the option to convert to an ARM outperforms the option to refinance to a FRM specifically because relative to the

⁶⁹Once borrowers regain their footing, they can always make the higher mortgage payments they would have owed under the less generous modification and avoid the larger payments at the end of the mortgage term. On the other hand, behavioral biases could make borrowers less likely to depart from the status-quo and make higher payments than required under their modified contracts even when it might be optimal for them to do so.

⁷⁰With no change in the default rate or prepayment risk, a maturity extension is NPV negative (positive) to the lender if the lender’s current discount rate is higher (lower) than the interest rate on the loan being extended.

refinance option, the conversion to an ARM front-loads the payment reduction to the borrower, and does so at similar cost to the investor. Our results provide empirical evidence that contracts front-loading payment reductions to households in financial distress will be more effective at preventing defaults than an alternative contract with equal cost to the investor that spreads payment reductions throughout the mortgage term. Similarly, Clara and Cocco (2016) compare an ARM with a refinance option to an ARM with the option to temporarily allow for interest-only payments, with a corresponding increase in its maturity. They find that this maturity extension option outperforms a refinance option because it provides similar liquidity-provision benefits to borrowers at much lower cost to investors. Our results show that immediate payment reduction can indeed be effective at reducing defaults even if structured with offsetting payments in the future so as to minimize costs to investors.


B.6 Effect of Payment Reduction on Consumption

Our payment reduction regression discontinuity empirical strategy is under-powered for studying consumption impacts. In Appendix Figure B.9 we plot the reduced form of the 31% PTI strategy with the change in mean credit card spending from the year before modification to the year after modification as the outcome variable. The standard error is so large that, using the same procedure for calculating an MPC as described in section 4.5, we cannot rule out an MPC above 1 or below -1.

Unlike with principal reduction, we are unable to increase the precision of our payment reduction estimates by using a difference-in-differences design. The difference in principal reduction received by borrowers with and without principal reduction remains large when we expand the sample to a wider bandwidth. In contrast, the difference in payment reduction between HAMP and Chase modifications falls when looking at a wider sample (as can be seen by looking at the edges of the top panel of Figure B.1). This is because the PTI target in HAMP generates larger payment reduction for higher PTI borrowers. Hence comparing borrowers who received HAMP and Chase modifications at a wider bandwidth results in a shrinking treatment size. We therefore conclude that our data and available research designs are unsuited for credibly estimating the effect of payment reduction on consumption.

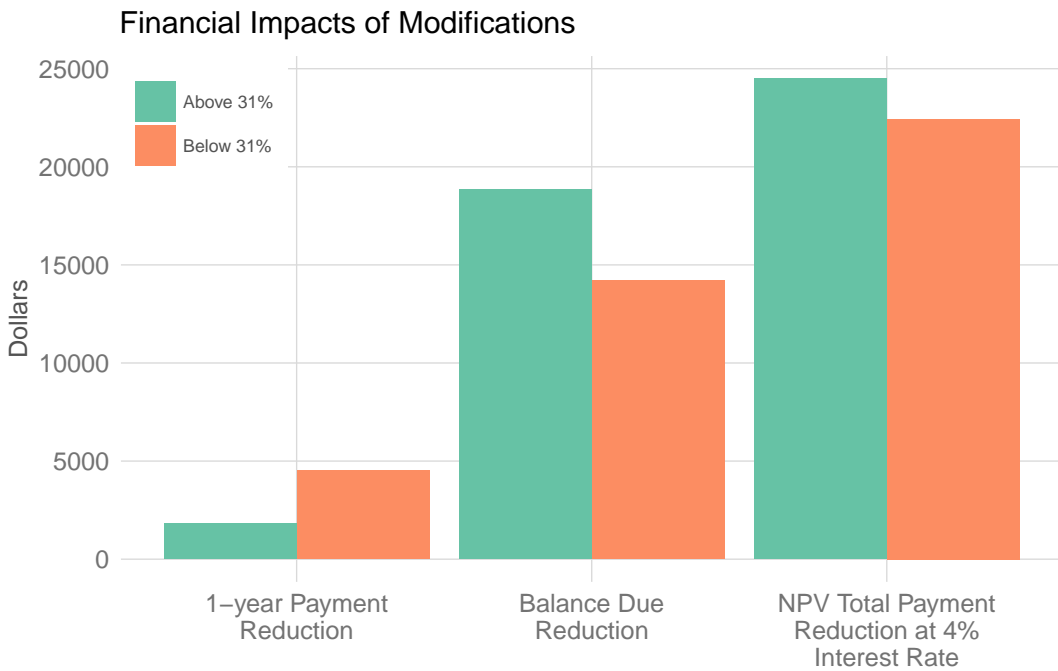

B.7 Payment Reduction Figures

Figure B.1: Effect of Payment Reduction on Default Using 31% PTI Cutoff



Notes: This figure shows the estimated effect of payment reduction on default using the 31% PTI regression discontinuity in the JPMCI bank dataset. The top panel plots the first stage, with payment reduction on the vertical axis and borrower PTI on the horizontal axis. The blue dots are conditional means for 12 equally spaced bins on each side of the cutoff. The red line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (7). The bottom panel plots the reduced form, with the default rate on the vertical axis and borrower PTI on the horizontal axis. Default 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 B.2.


Figure B.2: Financial Impacts of Modifications Around 31% PTI Cutoff



Notes: The top panel plots the difference in average annual payments relative to the pre-modification mortgage for borrowers with PTI ratios from 30%-31% and from 31.1%-32.1% in the JPMCI bank dataset. The bottom panel summarizes the financial impacts of modifications along various dimensions for borrowers above and below the cutoff..

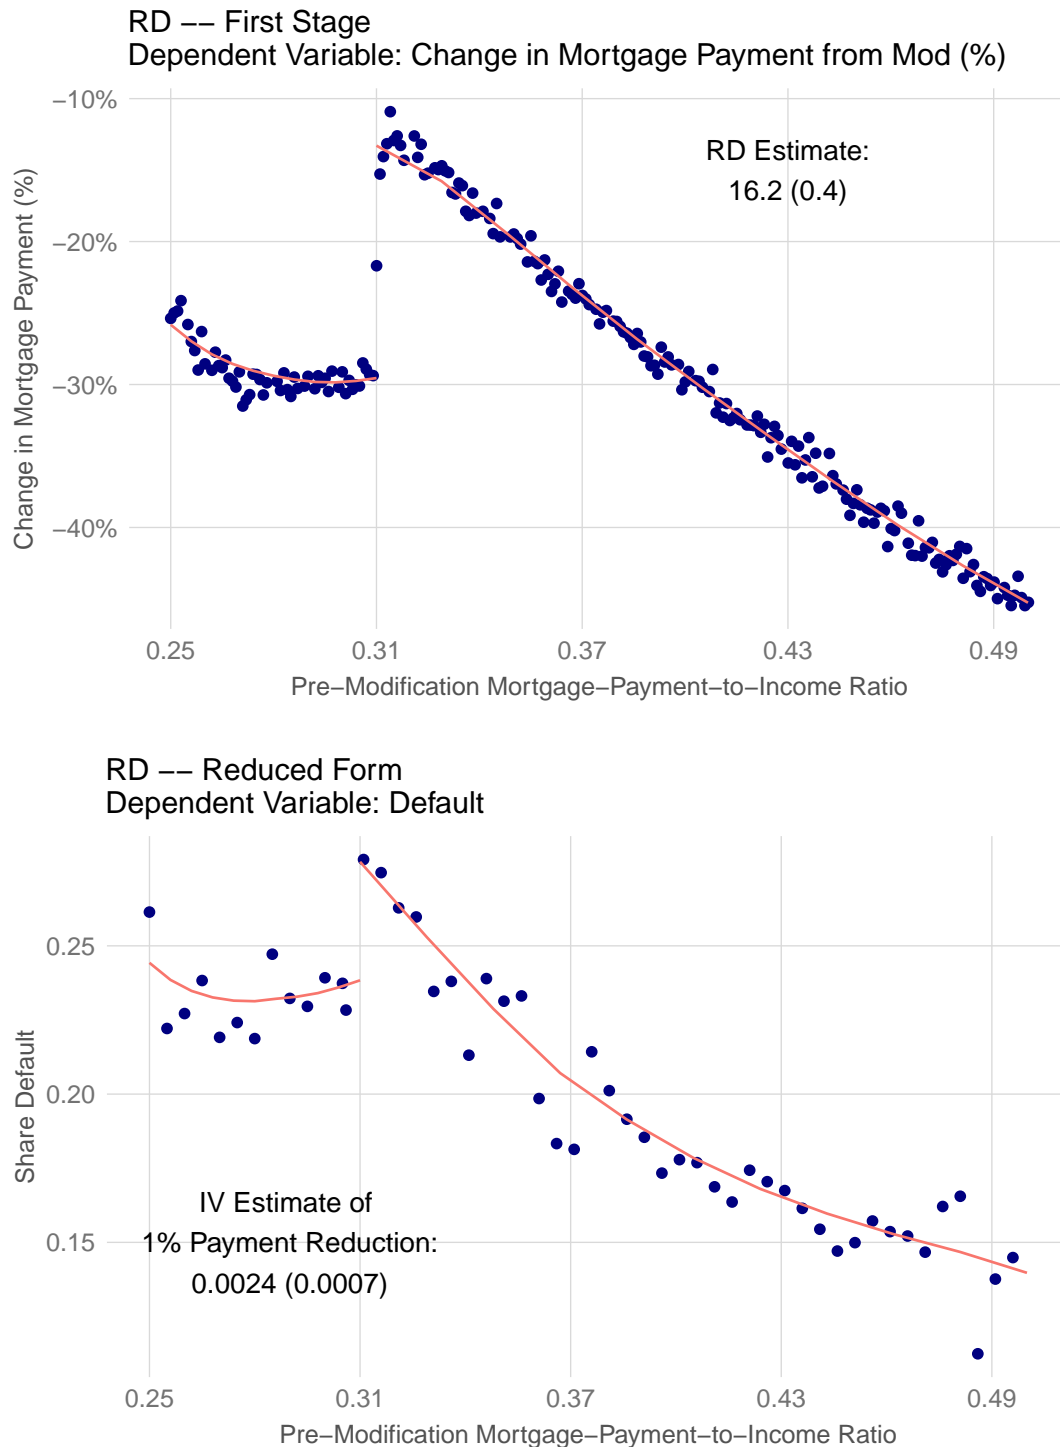
B.8 Payment Reduction Appendix Figures

Figure B.3: Structure of Mortgage Modifications





Notes: This figure plots the average annual payments owed under the status quo and under the modified mortgage contracts for borrowers receiving Chase mortgage modifications in the JPMCI bank dataset. This includes all borrowers in our sample with PTI ratios between 25% and 38% .

Figure B.4: Robustness of Payment Reduction Result to Inclusion of Mixed Treatment Group




Notes: This figure shows the estimated effect of payment reduction on default using the 31% PTI regression discontinuity in the JPMCI bank dataset, including observations with PTI between 31% and 31.1% which are dropped in the main analysis. The top panel plots the first stage, with payment reduction on the vertical axis and borrower PTI on the horizontal axis. The blue dots are conditional means for equally spaced bins on each side of the cutoff. Bins are four times narrower than in the baseline figures in order to visually capture the loans right at the cutoff. The red line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (7). The bottom panel plots the reduced form, with the default rate on the vertical axis and borrower PTI on the horizontal axis. Default 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 B.2.

Figure B.5: Modification Characteristics Around the 31% PTI Cutoff



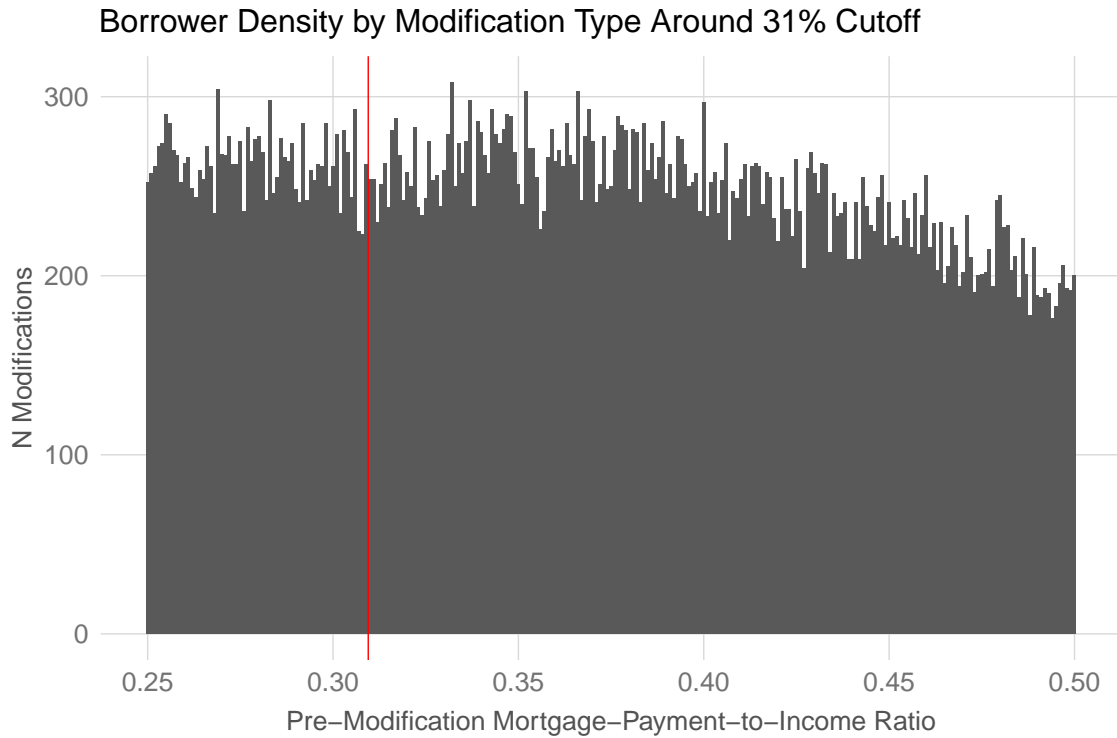
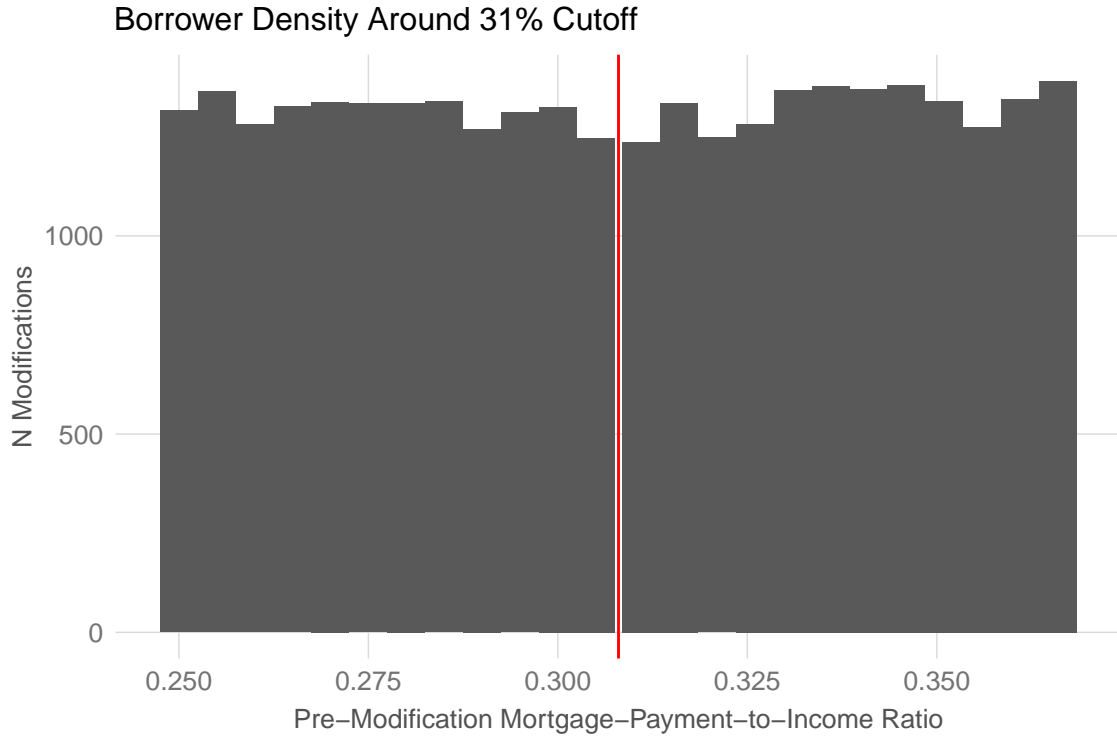
Notes: This figure shows the amount of principal reduction received around the 31% PTI regression discontinuity cutoff in the JPMCI bank dataset. The blue dots are conditional means for 12 bins on each side of the cutoff. The red line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (7).

Figure B.6: Initial Borrower and Mortgage Characteristics Around the 31% PTI Cutoff




Notes: This figure shows average pre-treatment characteristics around the 31% PTI regression discontinuity cutoff in the JPMCI bank dataset. The horizontal axis shows pre-modification borrower PTI. The vertical axis in the first four panels shows the ratio of unpaid principal balance to the market value of the house (mark-to-market loan-to-value ratio), the number of months past due at modification date, monthly income, and monthly payment. The final panel shows predicted default rates from a linear regression of default on the first four borrower characteristics. The blue dots are conditional means for 12 bins on each side of the cutoff. The red line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. The RD estimate is the jump in predicted values at the cutoff, corresponding to an estimate of the numerator in equation (7).

Figure B.7: Borrower Density Around the 31% PTI Cutoff



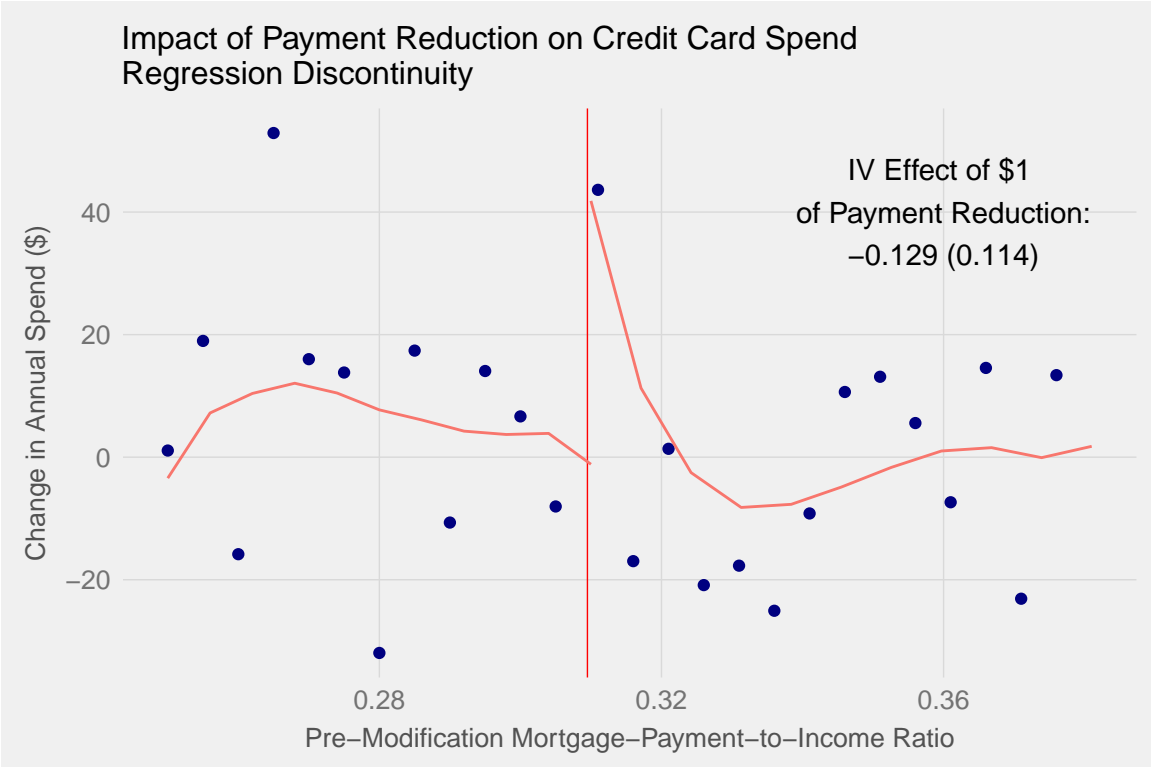
Notes: This figure plots the histogram of the running variable from our 31% PTI regression discontinuity strategy in the JPMCI bank dataset. The horizontal axis shows pre-modification borrower PTI. The top panel shows borrowers with PTI ratios between 25% and 37%, dropping the 250 observations between 31.0% and 31.1%. This is our main analysis sample. The bottom panel shows the density for the full sample.

Figure B.8: Effect of Payment Reduction on Default Using Alternative Bandwidths Around the 31% PTI Cutoff



Notes: This figure plots the estimated reduced form jump in default and the associated 95% confidence interval at the 31% PTI regression discontinuity cutoff calculated using alternative bandwidths in the JPMCI bank dataset. The Imbens and Kalyanaraman (2012) optimal bandwidth is 0.45.

Figure B.9: Effect of Payment Reduction on Credit Card Expenditure Using the 31% PTI Cutoff



Notes: This figure shows the reduced form of the estimated impact of payment reduction on credit card expenditure using the 31% PTI regression discontinuity strategy in the JPMCI bank dataset. The blue dots are conditional means for 12 bins on each side of the cutoff. The red line shows the predicted value from a local linear regression estimated separately on either side of the cutoff. Construction of the IV estimate $\hat{\tau}$ is described in section B.2. This strategy is unable to detect economically meaningful changes in expenditure.