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

Peter Ganong and Pascal Noel

Contents

A Appendix Figures and Tables

Figure 1: Match Rate around Principal Reduction Discontinuity

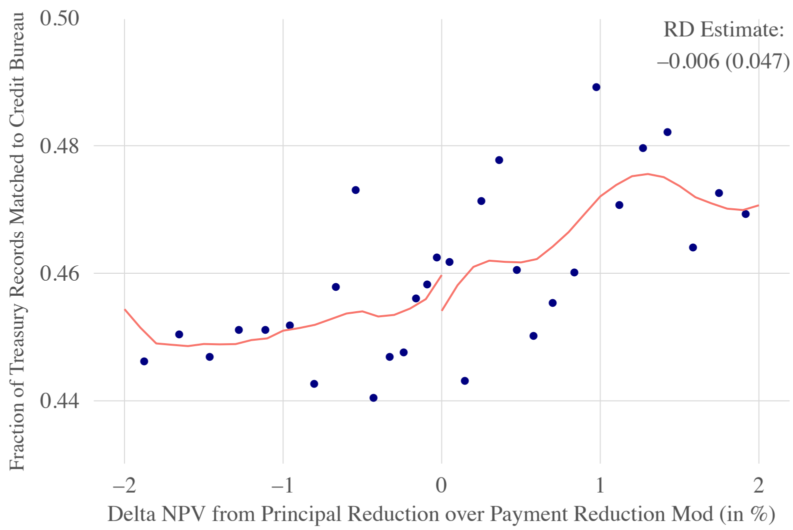

Notes: This figure plots the share of borrowers in the Treasury HAMP dataset successfully matched to their credit bureau records. The horizontal axis shows the normalized predicted gain to lenders of providing principal reduction to borrowers from equation (1). The blue dots are conditional means for 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 2: Modification Terms Summary

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

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

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

Figure 3: Mortgage Delinquency over Time

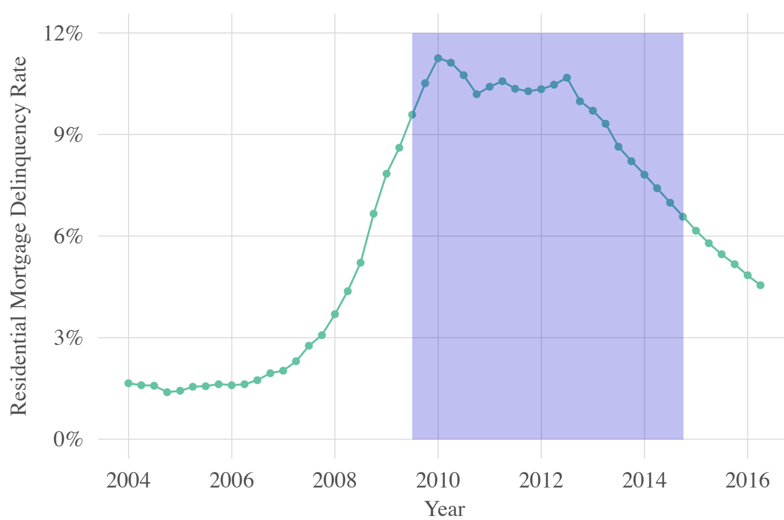

Notes: This figure plots the share of U.S. residential mortgages more than 30 days delinquent as reported by the Federal Reserve Board. The shaded region denotes the period where borrowers in our principal reduction sample had their first pre-modification delinquencies.

Figure 4: Pre-Modification Characteristics around Principal Reduction Discontinuity

(a) Pre-Mod FICO Score

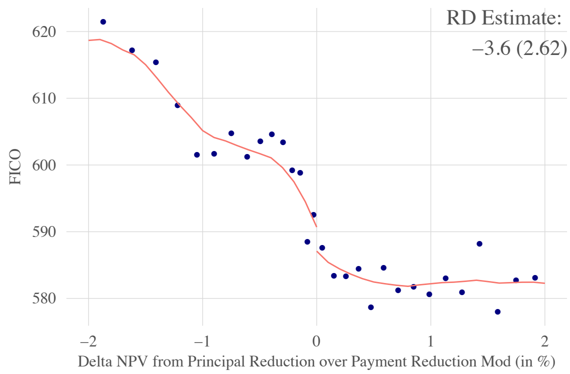

(b) Monthly Income at Mod Date

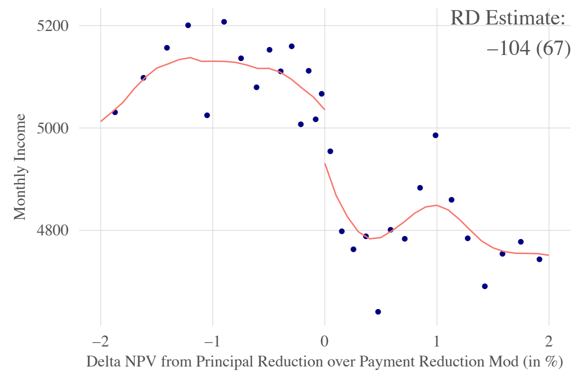

(c) Pre-Mod Payment-to-Income Ratio

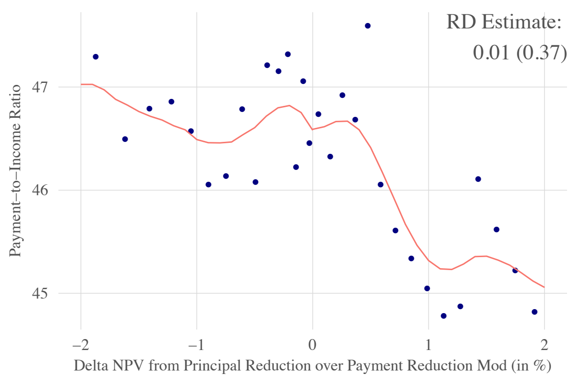

(d) Pre-Mod Mark-to-Market Loan-to-Value

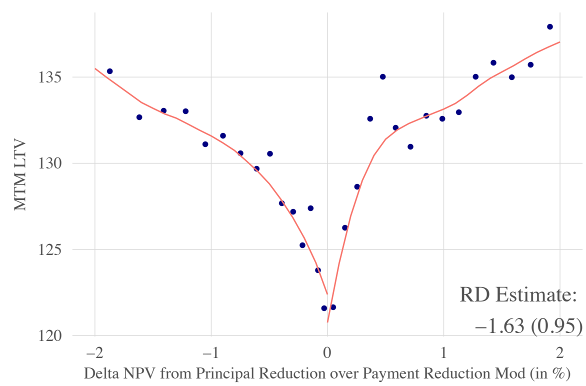

(e) Pre-Mod Months Past Due

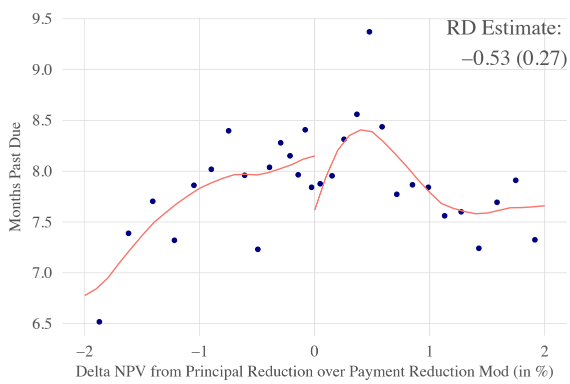

(f) Predicted Default Rate Using Covariates

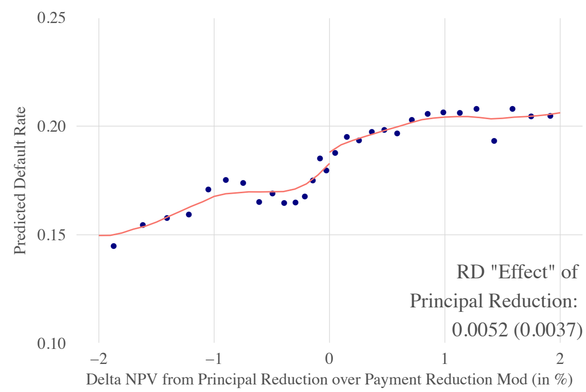

Notes: This figure shows average pre-treatment characteristics around the regression discontinuity cutoff in the matched HAMP credit bureau dataset. The horizontal axis shows the normalized predicted gain to lenders of providing principal reduction to borrowers from equation (1). The vertical axis in the first five panels shows borrower credit score, monthly income, the ratio of monthly mortgage payments to monthly income, the ratio of unpaid principal balance to the market value of the house (mark-to-market loan-to-value ratio), and the number of monthly mortgage payments the borrower is past due at application date. The final panel shows predicted default rates from a linear regression of default on the first five borrower characteristics. The 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). See online Appendix B.1.2 for details.

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

(a) Histogram of Running Variable

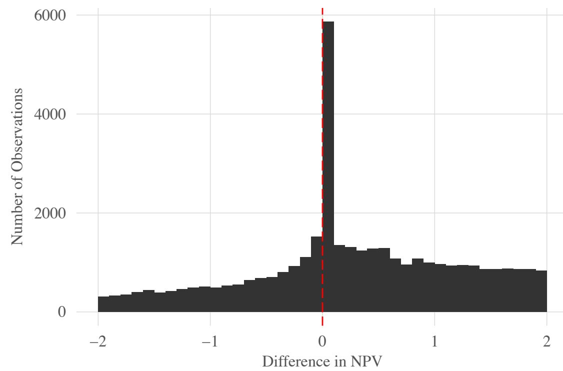

(b) Histogram of Running Variable Excluding Zeros

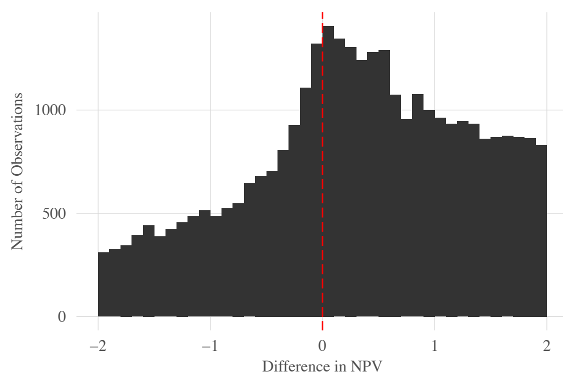

(c) Take-up Rate

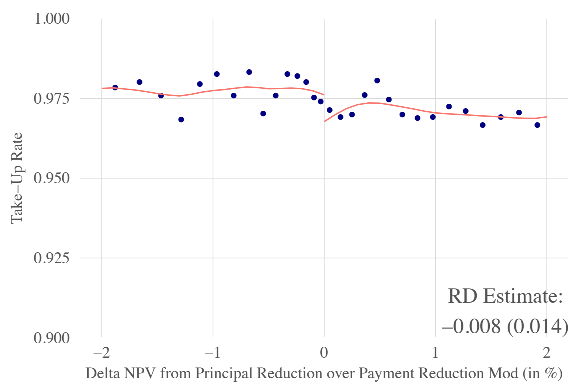

Notes: Panel (a) plots the histogram of the running variable from our regression discontinuity strategy in the matched HAMP credit bureau dataset. The horizontal axis shows the normalized predicted gain to lenders of providing principal reduction to borrowers from equation (1). HAMP program officers in the U.S. Treasury Department explain that the mass at exactly zero is due to data misreporting. Some servicers reported a single number as the calculation for both the payment reduction and principal reduction modifications, meaning that the estimated gains from principal reduction were calculated to be zero. Panel (b) plots the same histogram dropping observations exactly at zero, which is our analysis sample. Online Appendix B.1.2 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. Panel (c) shows the take-up rate conditional on borrowers being offered a modification in the Treasury HAMP dataset.

Figure 6: Treatment Size around Principal Reduction Discontinuity

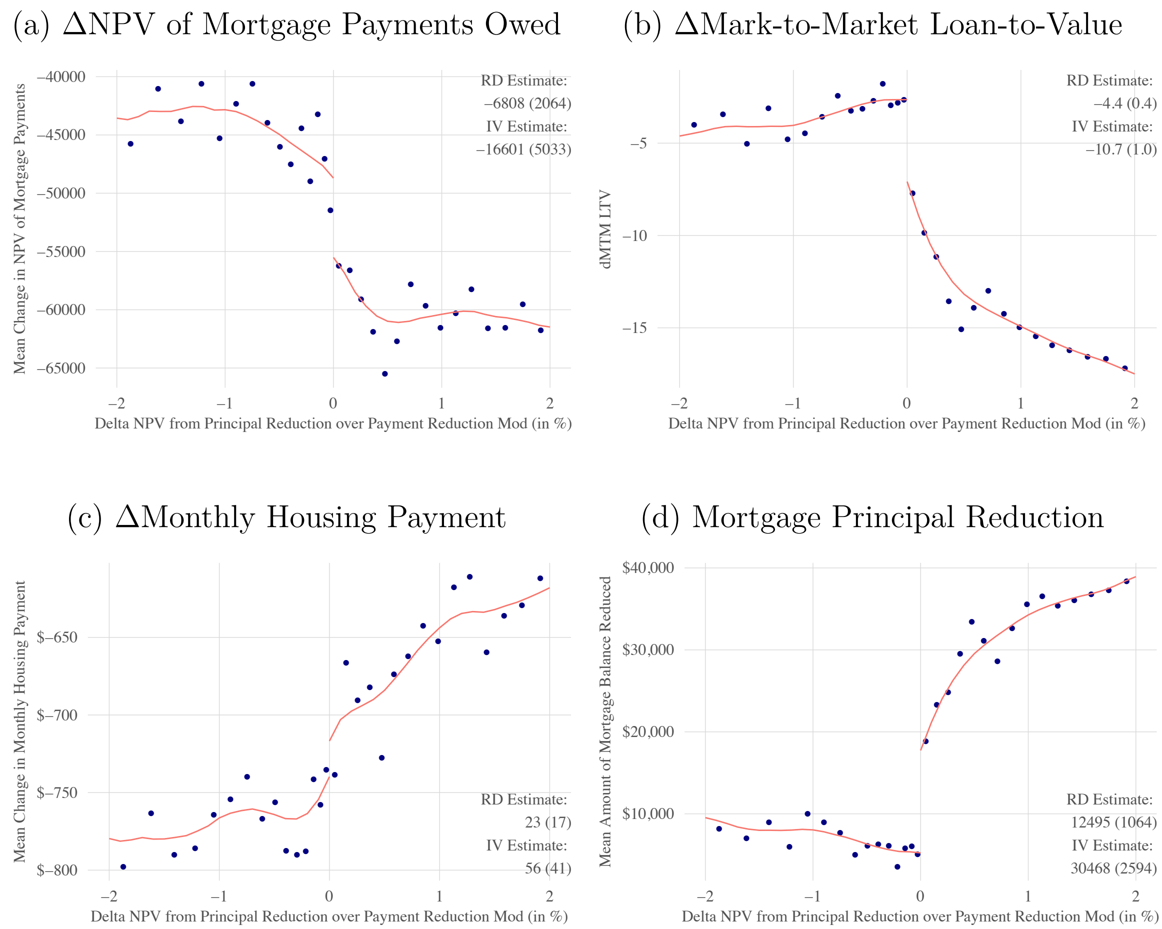

Notes: This figure shows the treatment size at the regression discontinuity cutoff in the matched HAMP credit bureau dataset. The horizontal axis shows the normalized predicted gain to lenders of providing principal reduction to borrowers from equation (1). Panel (a) shows the change in the net present value (NPV) of mortgage payments owed under the modified contract relative to the status quo mortgage contract, discounted at a 4 percent interest rate, panel (b) shows the change in the loan-to-value ratio, panel (c) shows the change in initial monthly housing payments, and panel (d) shows the average amount of principal reduction per borrower. The 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 3.2.

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

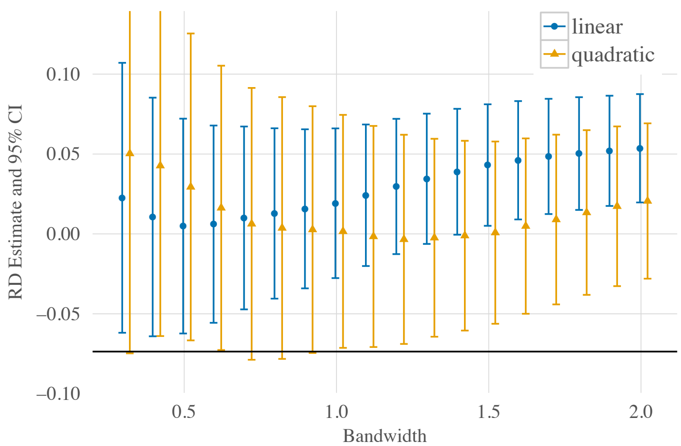

Notes: This figure shows the estimated impact of principal reduction on default under various specifications and bandwidths in the matched HAMP credit bureau dataset. Each line plots the IV estimate and associated 95 percent confidence interval from a local linear or quadratic regression on either side of the cutoff. The optimal bandwidths for the linear specification from Imbens and Kalyanaraman (2012) and Calonico 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 8: Effect of Principal Reduction on Foreclosure Initiation

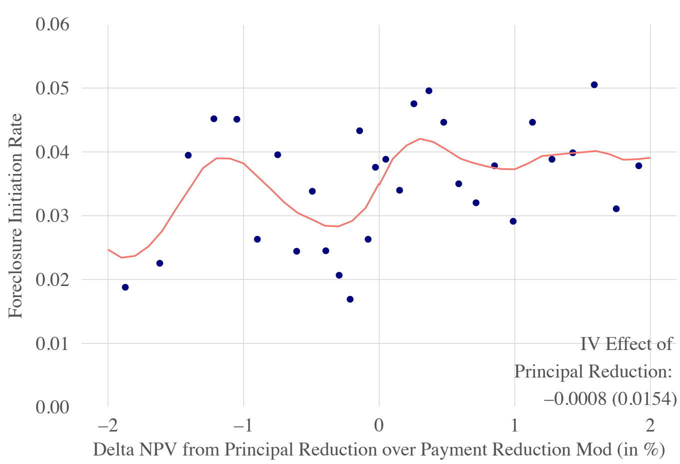

Notes: This figure shows the effect of principal reduction on foreclosure initiation in the matched HAMP credit bureau dataset. The foreclosure initiation rate is plotted on the vertical axis and the normalized predicted gain to lenders of providing principal reduction is on the horizontal axis. The 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 3.2.

Figure 9: Spending around Modifications with and without Principal Reduction

(a) Credit Card Spending around Modification – Normalized

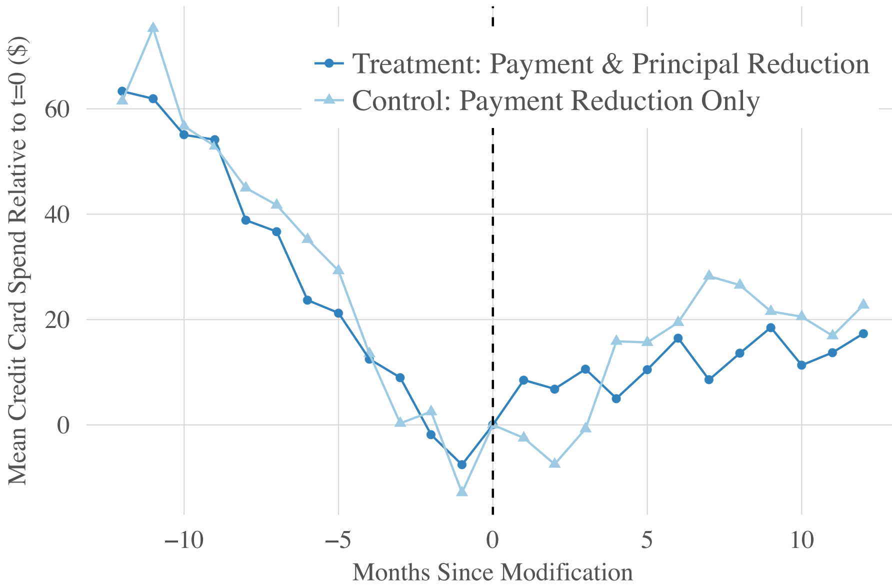

(b) Auto Spending around Modification

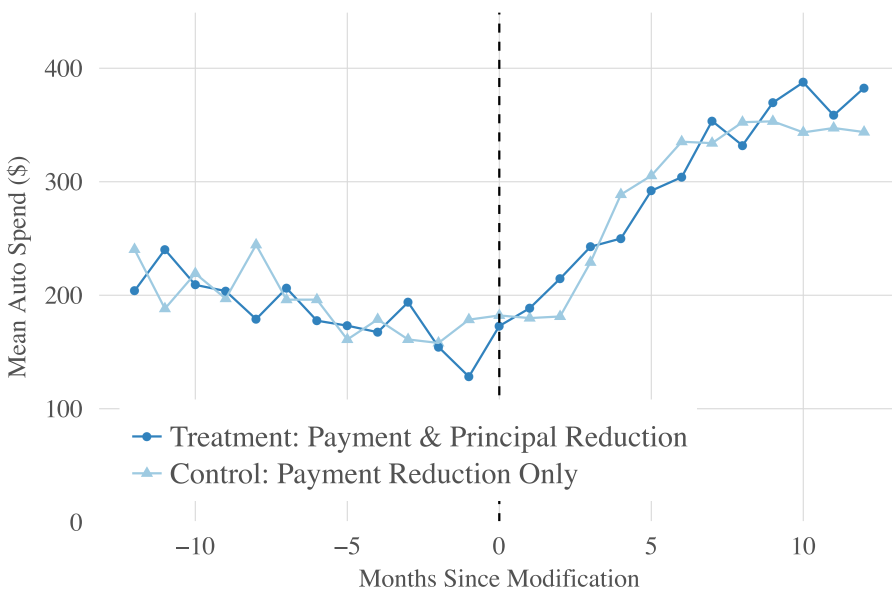

Notes: This figure shows the event study of monthly spending around modification for borrowers receiving each type of modification in the matched HAMP credit bureau dataset. The top panel plots credit card expenditure in dollars as measured from credit bureau records relative to the month of modification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

Figure 10: Spending around Modifications with and without Principal Reduction using JPMCI Data

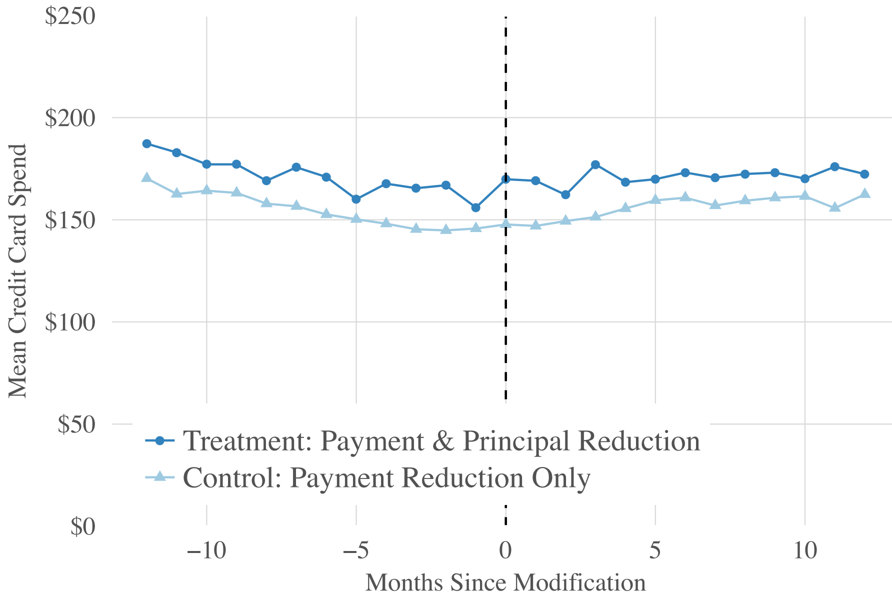

Notes: This figure shows the event study of monthly credit card expenditure around modification for borrowers receiving each type of modification in the JPMCI bank account dataset. For further details see sections 2.2 and 4.2.

Figure 11: Credit Card and Auto Spend around Principal Reduction Discontinuity

(a) Change in Mean Monthly Credit Card Spending

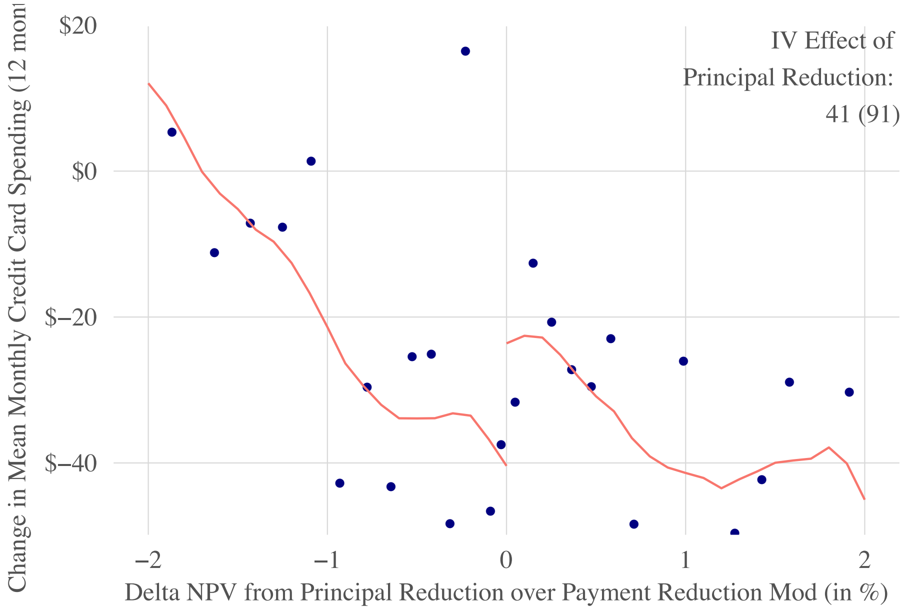

(b) Change in Mean Monthly Auto Spending

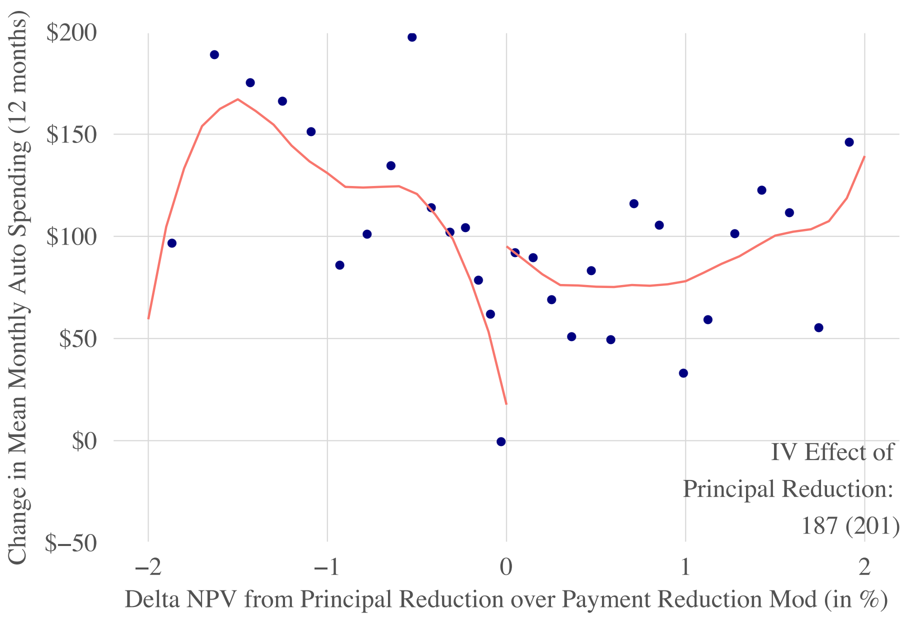

Notes: This figure shows the estimated impact of principal reduction on expenditure using the fuzzy regression discontinuity strategy in the matched HAMP credit bureau dataset. The horizontal axis shows the normalized predicted gain to lenders of providing principal reduction to borrowers from equation (1). The vertical axis on the top panel shows the average change in credit card expenditure between the twelve months before modification and the twelve months after modification. Credit card expenditure is measured from credit bureau records as discussed in section 2.1. The vertical axis in the bottom panel shows the average change in auto spending between the 12 months before modification and the 12 months after modification. Auto spending is measured from new auto loans, as described in Section 2.1. The 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 3.2.

Figure 12: Pre-Modification Characteristics around Payment Reduction Discontinuity

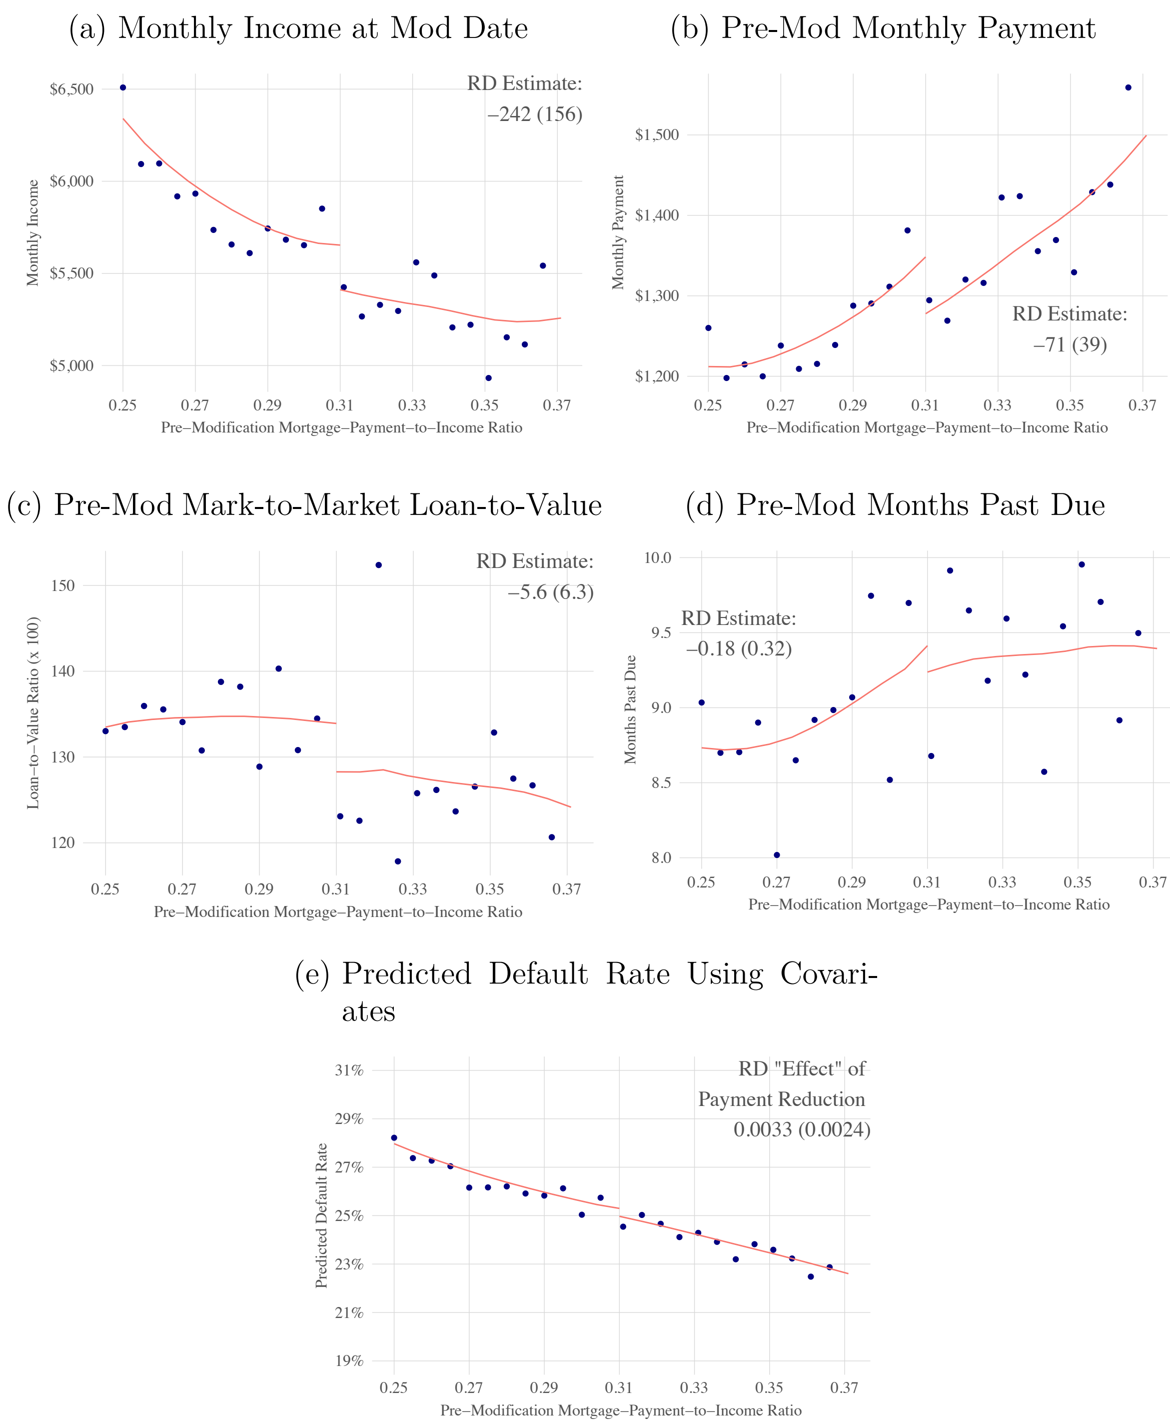

Notes: This figure shows average pre-treatment characteristics around the 31 percent PTI regression discontinuity cutoff in the JPMCI bank dataset for non-GSE-backed loans. The horizontal axis shows pre-modification borrower PTI. The vertical axis in the first four panels shows monthly income, monthly payment, the ratio of unpaid principal balance to the market value of the house (mark-to-market loan-to-value ratio), and the number of months past due at modification date. The final panel shows predicted default rates from a linear regression of default on the first four borrower characteristics. The 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 (4).

Figure 13: Borrower Density around the Payment Reduction Discontinuity

(a) Histogram of Running Variable – Near Discontinuity

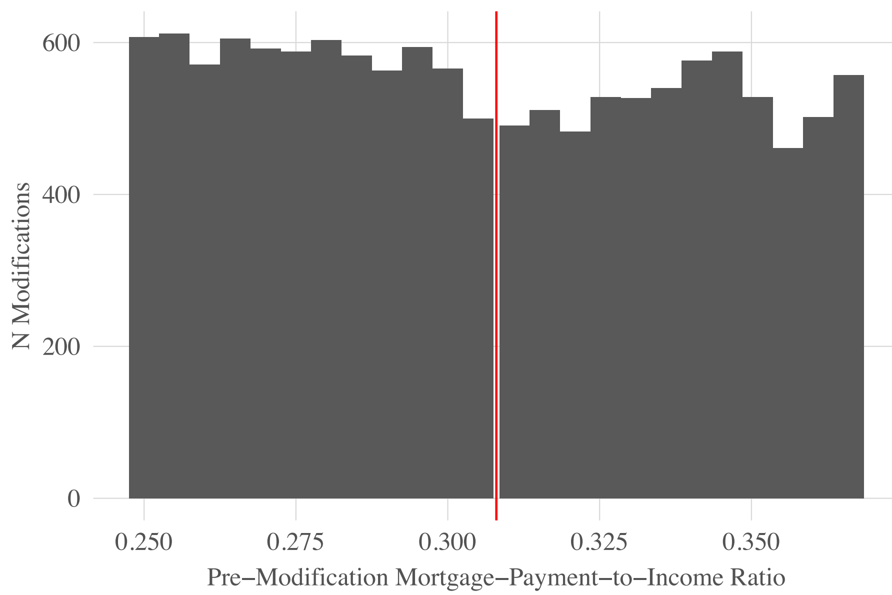

(b) Histogram of Running Variable – Full Support

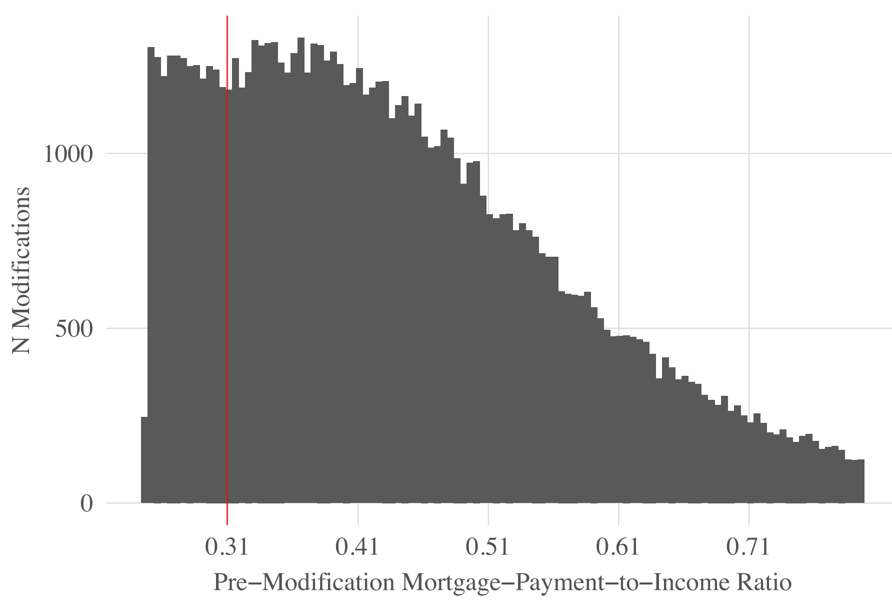

Notes: This figure plots the histogram of the running variable from our 31 percent PTI regression discontinuity strategy in the JPMCI bank dataset. The horizontal axis shows pre-modification borrower PTI. The top panel shows borrowers in the main analysis sample. This sample is restricted to pre-modification PTI ratio between 25 percent and 37 percent (dropping the 258 observations between 31.0 percent and 31.1 percent), pre-modification terms 30 years or less, and fixed rate loans. This is our main analysis sample. The bottom panel shows the density for the full sample.

Figure 14: Treatment Size around Payment Reduction Discontinuity

(a) ΔNet Present Value of Payments Owed

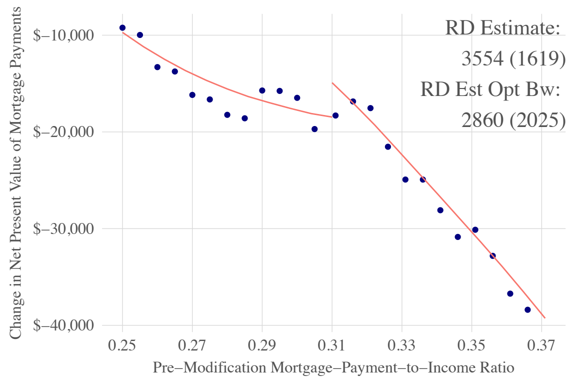

(b) Mortgage Principal Forgiveness

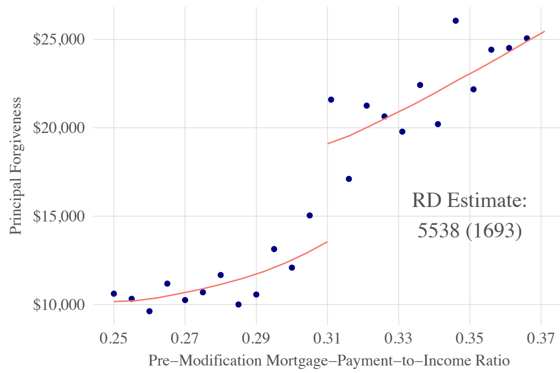

(c) ΔInterest Rate

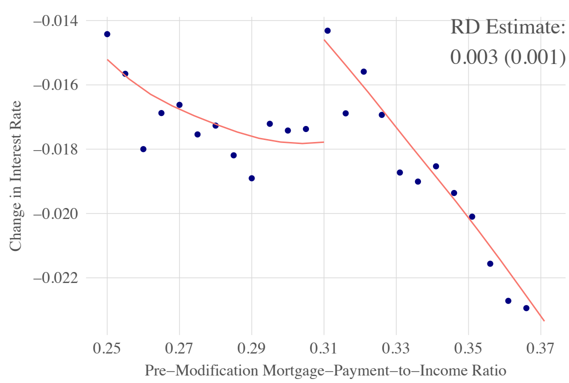

Notes: This figure describes the treatment in terms of long-term obligations around the 31 percent PTI discontinuity in the JPMCI bank dataset for non-GSE-backed loans. The blue dots are conditional means for 12 bins on each side of the cutoff. The red line shows the predicted value from a local linear 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 (4) using the IK-optimal bandwidth for delinquency of 0.057. Panel (a) shows the change in the NPV of payments owed under the mortgage contract before and after modification. The IK-optimal bandwidth for this outcome variable is 0.039 and the label also includes a second RD estimate using this optimal bandwidth of 0.039. Panel (b) shows mortgage principal forgiveness. Panel (c) shows the change in the interest rate.

Figure 15: Effect of Payment Reduction on Default: Robustness to Broader Sample

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

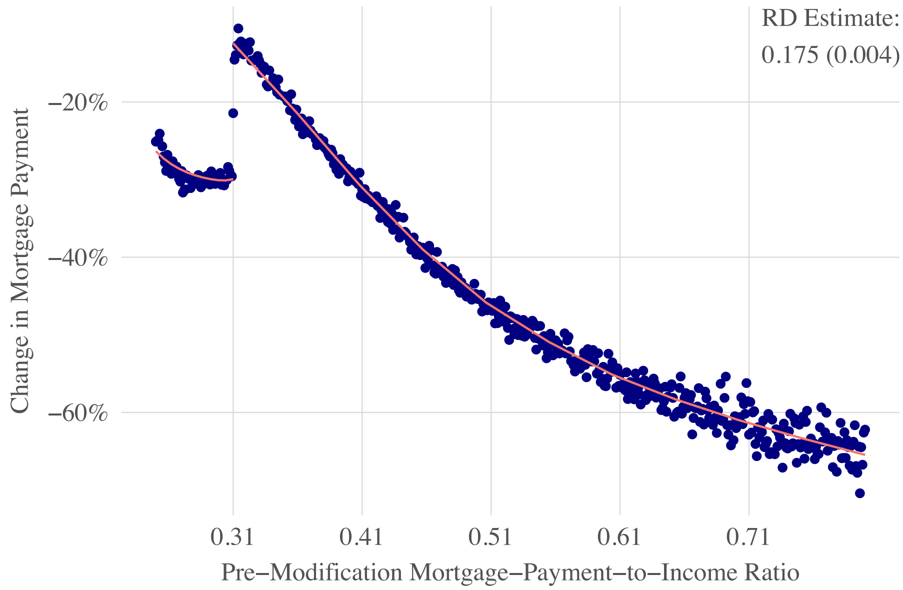

(b) Reduced Form -- Mortgage Default

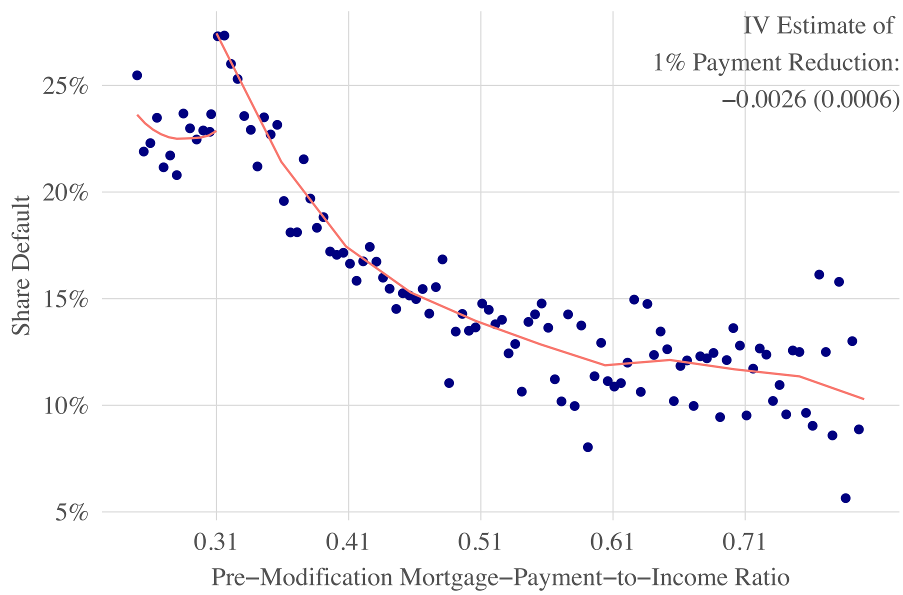

Notes: This figure shows the estimated effect of payment reduction on default using the 31 percent PTI regression discontinuity in the JPMCI bank dataset for a broader sample of non-GSE-backed loans. It includes loans with pre-modification terms greater than 30 years, loans with variable interest rates,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

Figure 16: Effect of Payment Reduction on Default Using Alternative Bandwidths

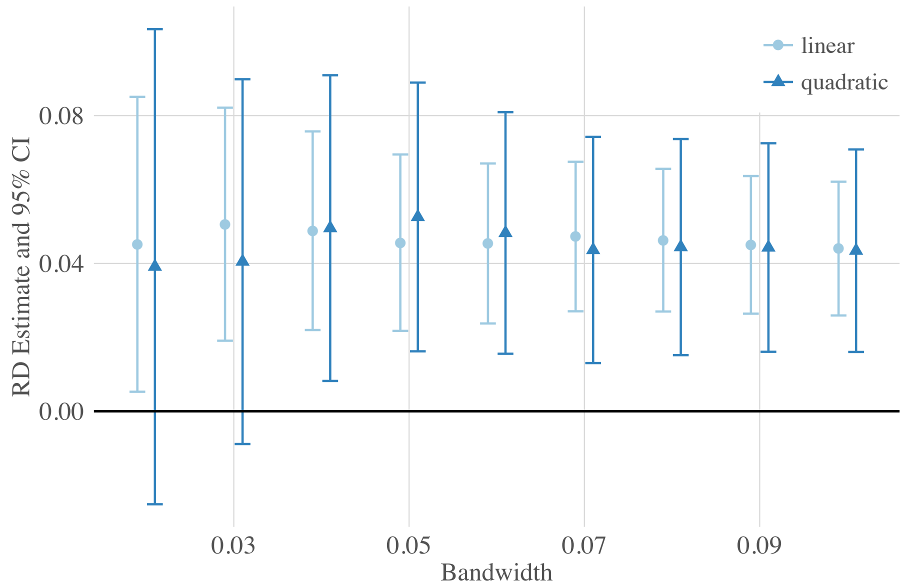

Notes: This figure plots the estimated reduced form jump in default and the associated 95 percent confidence interval at the 31 percent PTI regression discontinuity cutoff calculated using alternative bandwidths in the JPMCI bank dataset for non-GSE-backed loans. The optimal bandwidths for the linear specification from Imbens and Kalyanaraman (2012) and Calonico et al. (2014) are 0.057 and 0.023, respectively. The optimal bandwidths for a quadratic specification from Imbens and Kalyanaraman (2012) and Calonico et al. (2014) are 0.055 and 0.031, respectively.

Figure 17: Effect of Payment Reduction on Default for GSE-Backed Loans

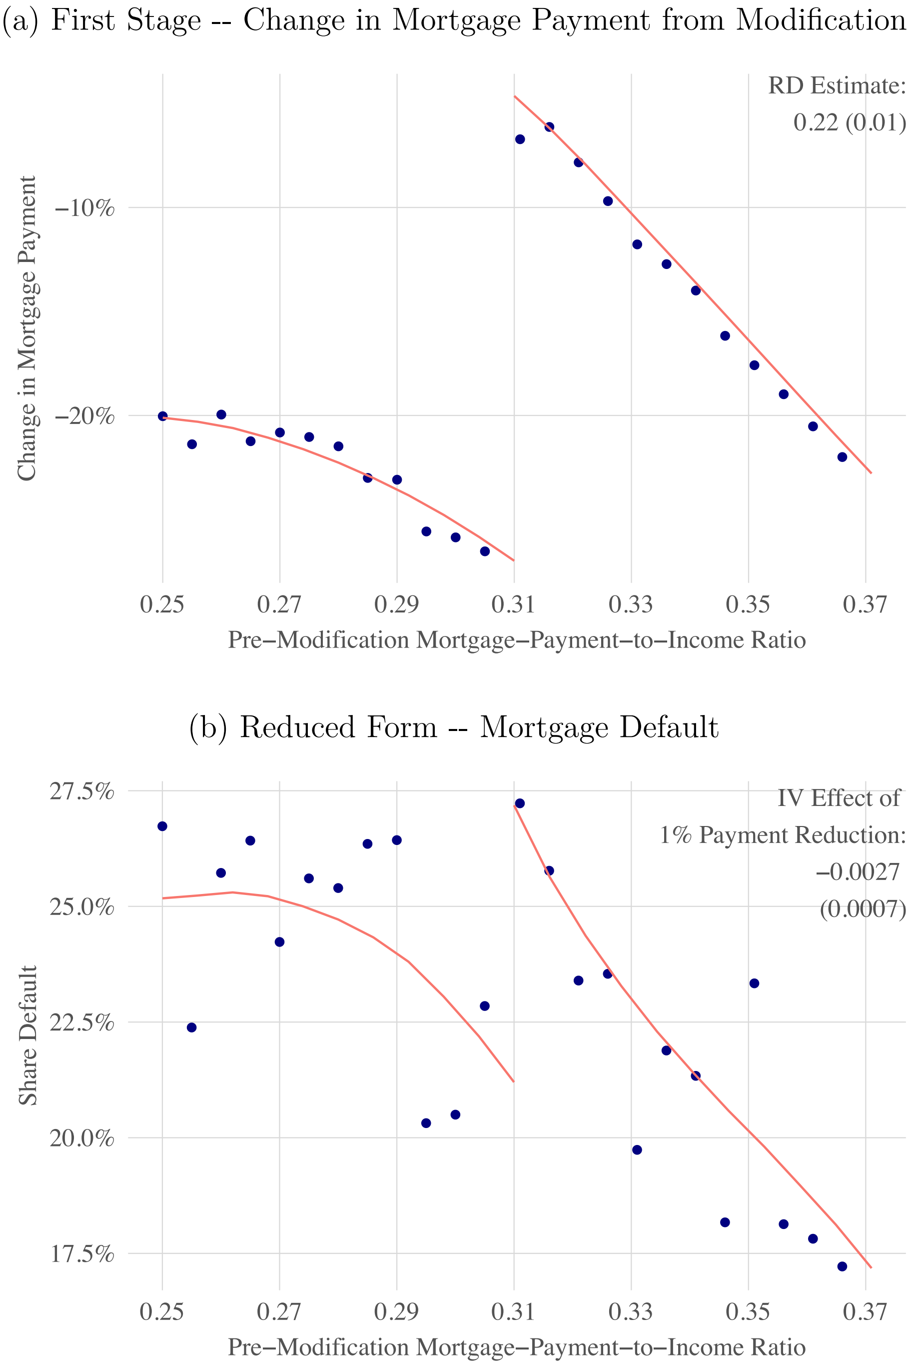

Notes: This figure evaluates the impact of payment reduction on default using a regression discontinuity at the 31 percent payment-to-income (PTI) in the JPMCI bank dataset for GSE-backed loans. The horizontal axis shows borrower PTI. The blue dots are conditional means for 12 equally spaced bins on 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 4. The top panel plots mean payment reduction and the bottom panel plots the default rate on the vertical axis, which is defined as being 90 days delinquent at any point within two years of the modification date. Construction of the IV estimate $\hat{\tau}$ is described in section 5.2.

Figure 18: Treatment Size around Payment Reduction Discontinuity for GSE-Backed Loans

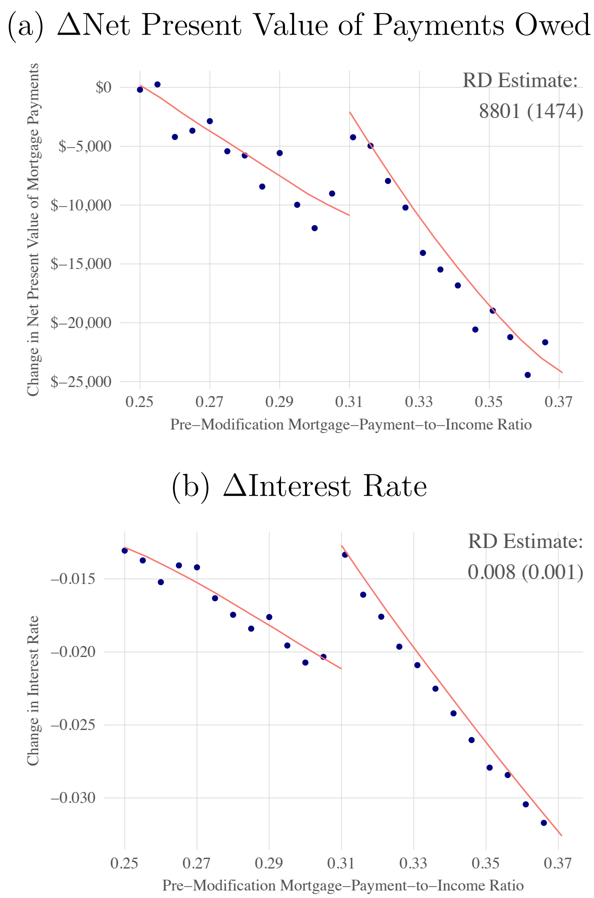

Notes: This figure describes the treatment in terms of long-term obligations around the 31 percent PTI discontinuity in the JPMCI bank dataset for GSE-backed loans. The blue dots are conditional means for 12 bins on each side of the cutoff. The red line shows the predicted value from a local linear regression 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 (4). Panel (a) shows the change in the NPV of payments owed under the mortgage contract for all loans. Panel (b) shows the change in the interest rate. We omit principal forgiveness because the GSEs did not offer mortgage principal forgiveness for either private modifications or HAMP modifications.

Figure 19: Pre-Modification Characteristics around Payment Reduction Discontinuity for GSE-Backed Loans

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

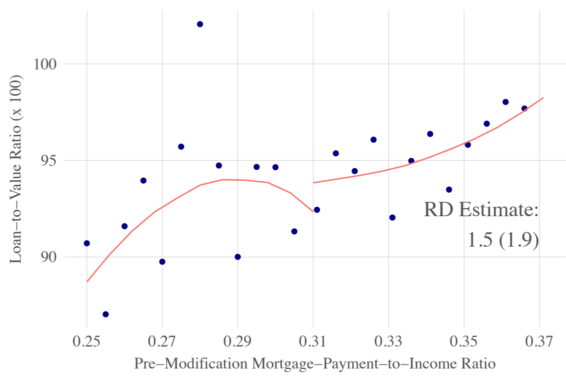

(b) Pre-Mod Months Past Due

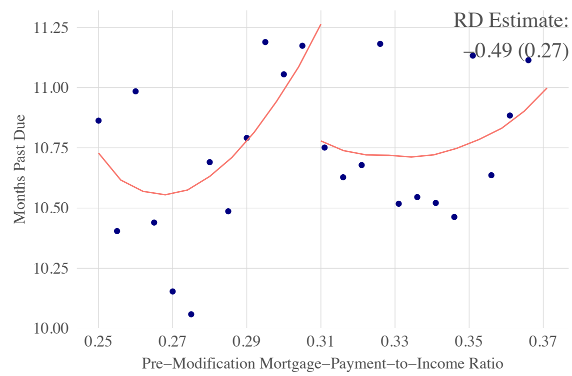

(c) Monthly Income at Mod Date

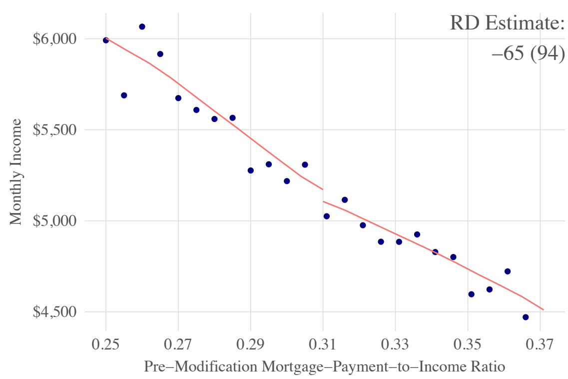

(d) Pre-Mod Monthly Payment

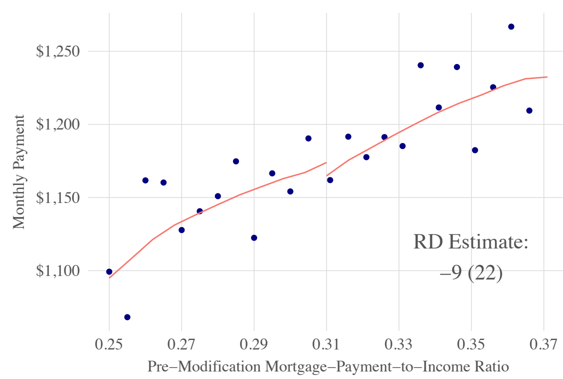

(e) Predicted Default Rate Using Covariates

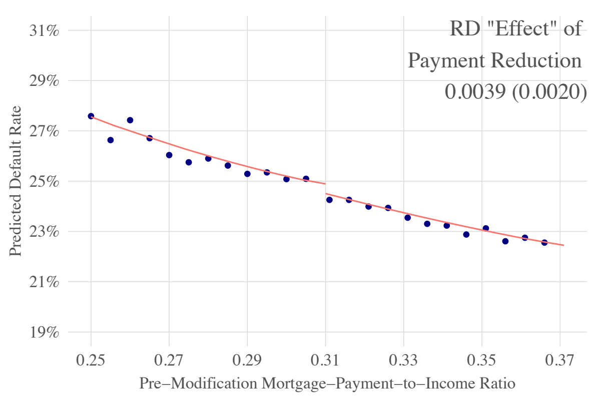

Notes: This figure shows average pre-treatment characteristics around the 31% PTI regression discontinuity cutoff in the JPMCI bank dataset for GSE-backed loans. The horizontal axis shows pre-modification borrower PTI. The vertical axis in the first four panels shows the ratio of unpaid principal balance to the market value of the house (mark-to-market loan-to-value ratio), the number of months past due at modification date, monthly income, and monthly payment. The final panel shows predicted default rates from a linear regression of default on the first four borrower characteristics. The 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 (4).

Figure 20: Borrower Density around the Payment Reduction Discontinuity for GSE-Backed Loans

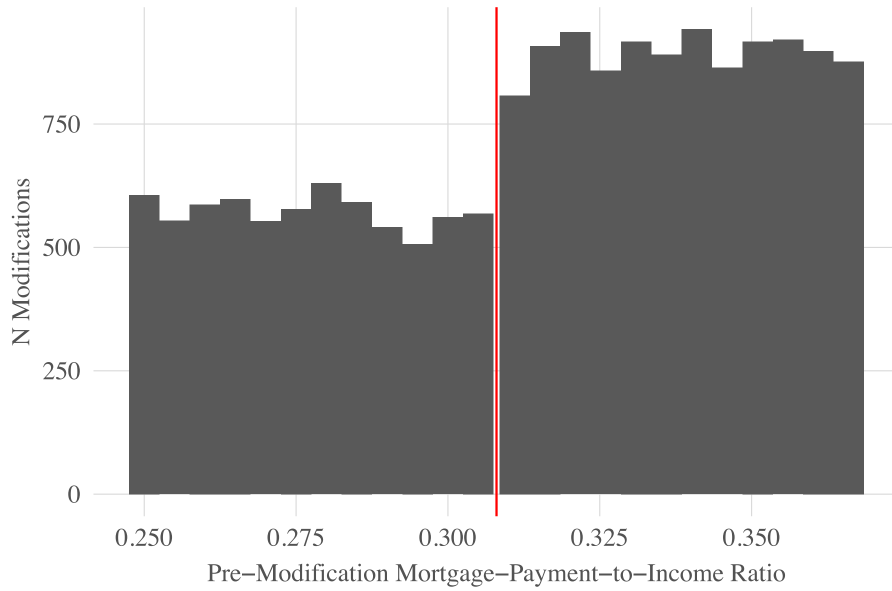

Notes: This figure plots the histogram of the running variable from our 31 percent PTI regression discontinuity strategy in the JPMCI bank dataset for GSE-backed loans. The horizontal axis shows pre-modification borrower PTI.

Figure 21: Projected 40-Year Mortgage Interest Rates

(a) Actual and Projected Loan Interest Rates

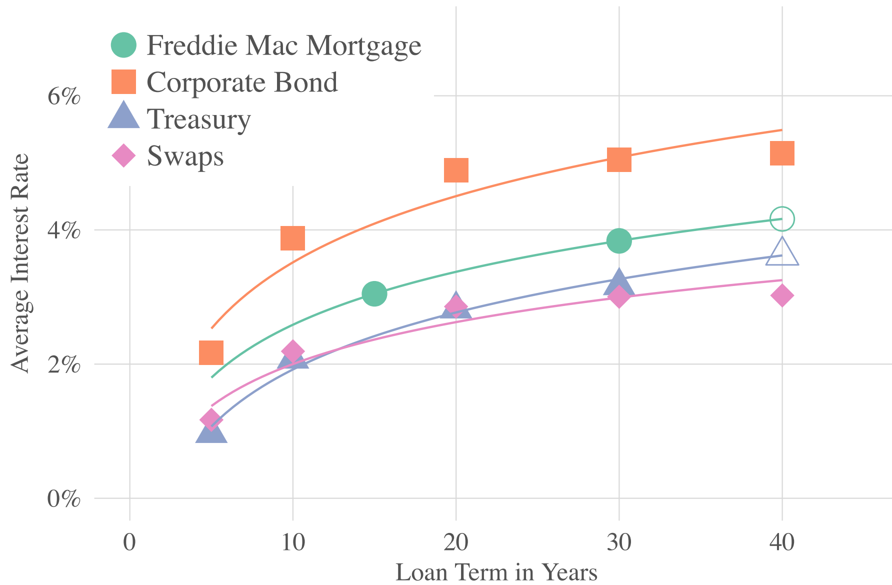

(b) Actual and Modeled Spread Between 30- and 40-year Rates

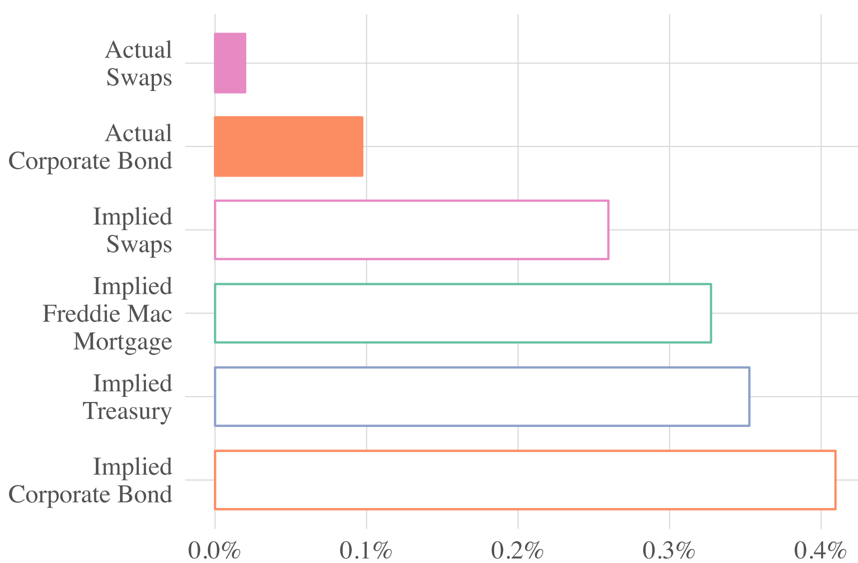

Notes: Panel (a) shows interest rates for various loan terms. Solid dots are data, lines are the best fit of $y = log(x)$ to the solid dots, and hollow dots are projections of 40-year interest rates. Green dots show mortgage rates from the Freddie Mac Conforming Loan Survey, red squares show corporate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

Figure 22: Mortgage Credit Availability

(a) Mortgage Originations by Credit Score

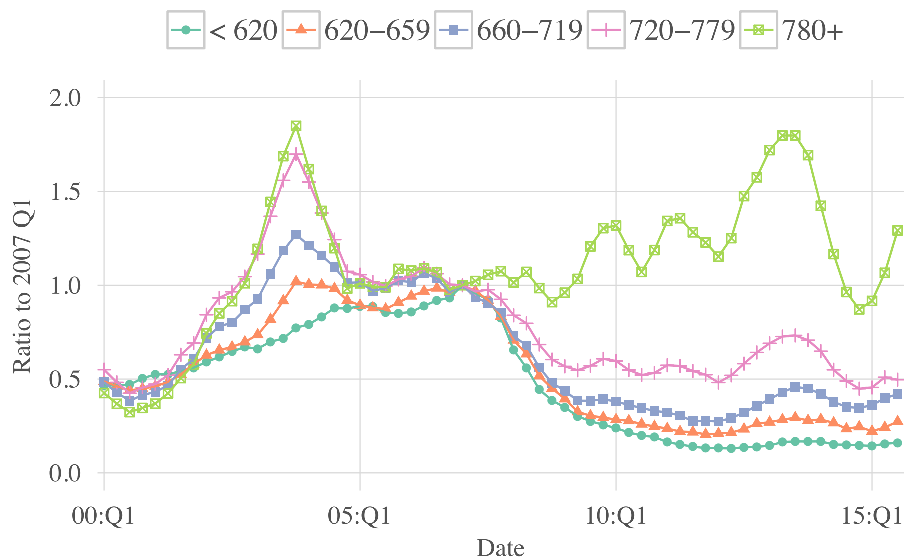

(b) Combined Loan-to-Value for New Home Equity Lines of Credit

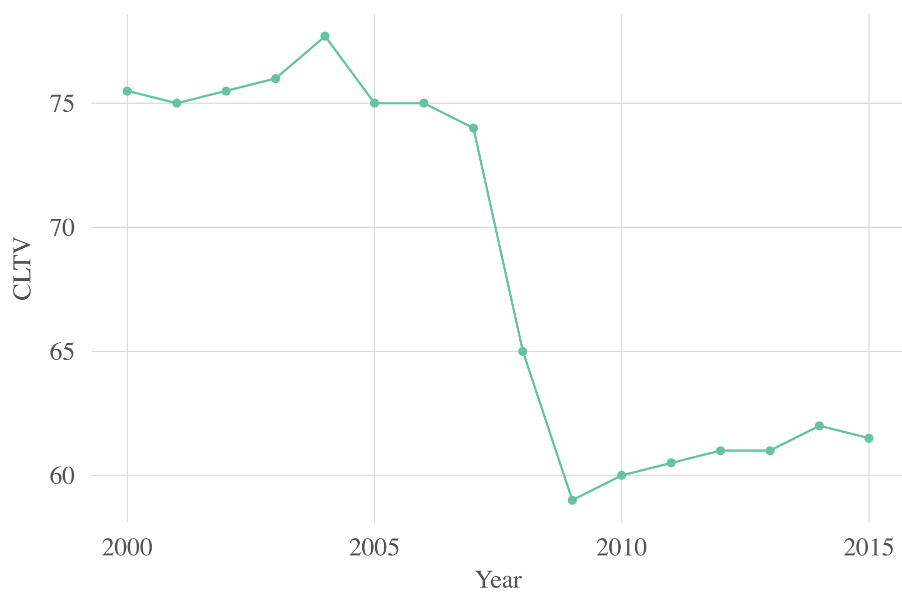

Notes: The top panel plots mortgage origination by borrower credit score from the New York Fed Consumer Credit Panel (Federal Reserve 2015). This includes first mortgages, second mortgages, and home equity installment loans. The bottom panel plots the average combined loan-to-value (CLTV) ratio for new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

Figure 23: Amount of Payment Reduction and Default

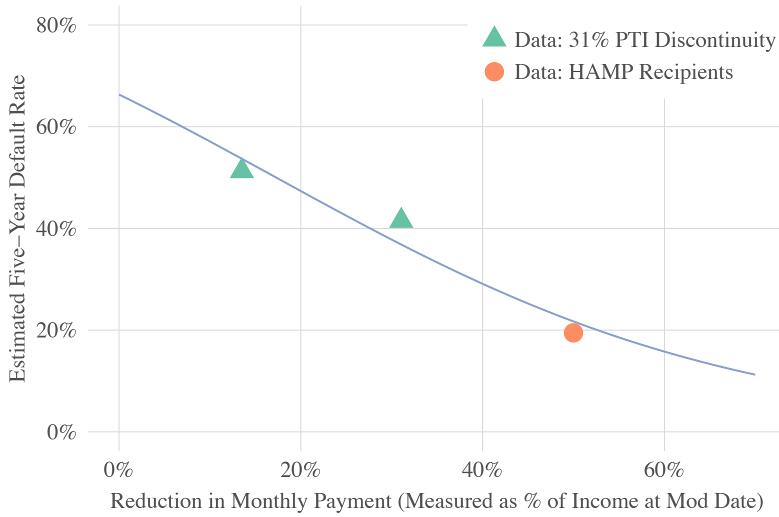

Notes: This figure shows estimated five-year default rates for various amounts of payment reduction. The green triangles are from the two sides of the discontinuity in Figure 5b and the orange circle is borrowers with PTI of 55 percent from online Appendix Figure 15. We take the two-year default rates and multiply them by 1.62, which is the ratio of five-year default rates to two-year default rates among HAMP modifications performed in 2010. The line is a best fit of a logit model to the data.

Figure 24: Effect of 10% Payment Reduction on NPV: Robustness

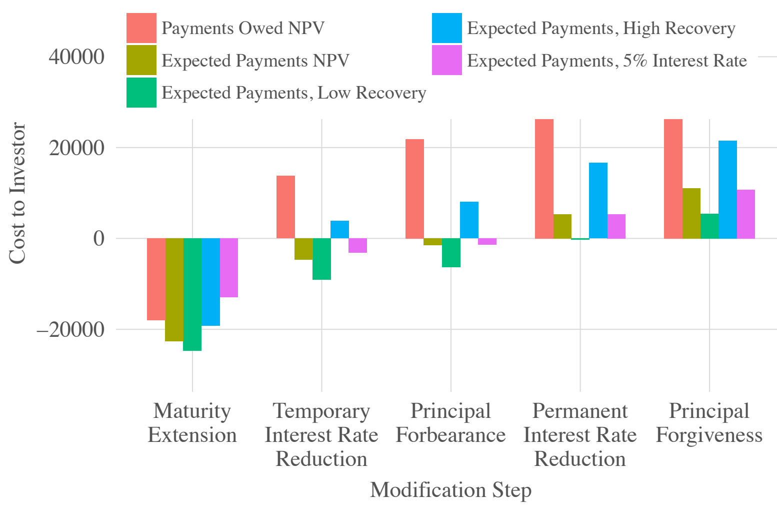

Notes: This figure shows the impact of a 10 percent payment reduction on the NPV of the loan to the investor under various assumptions. The red and yellow bars reproduce Figure 6a. The yellow bars assume a 39 percent self-cure rate on post-modification defaults and a 56 percent loss if the loan is liquidated. The green bars assume a self-cure rate of 18 percent and a liquidation loss of 61 percent. The blue bars assume a 61 percent self-cure rate and a 48 percent liquidation loss. See online Appendix C for the data sources for each of these assumption. The purple bars use the same assumptions as baseline, except a 5 percent initial interest rate.

Figure 25: NPV Cost of Payment Reduction for Various Sequences of Modification Steps

(a) Add Social Value of Payment Reduction

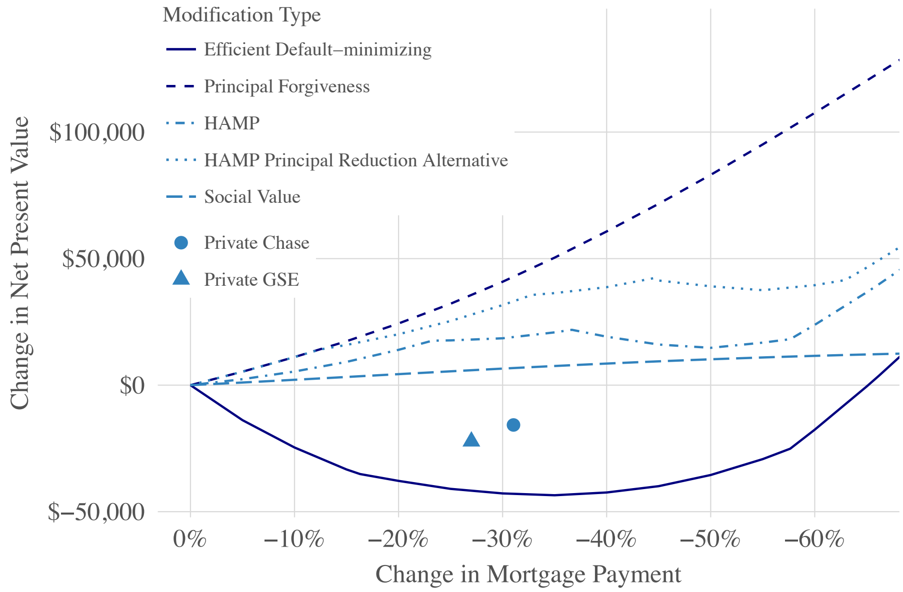

(b) NPV of Payments Owed

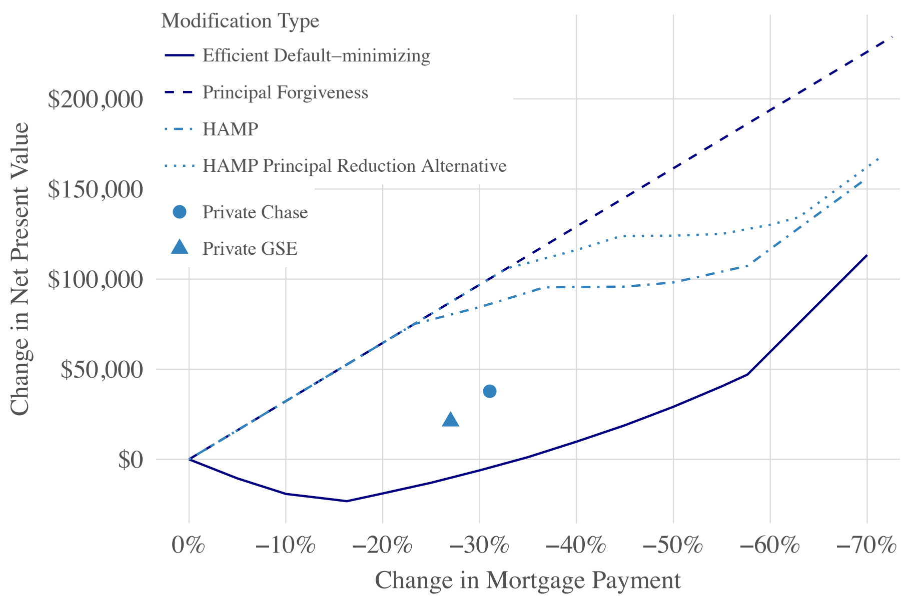

Notes: The top panel takes Figure 6b and adds a line reflecting the social value of payment reduction, assuming a $51,000 social cost per foreclosure as estimated in U.S. Department of Housing and Urban Development (2010). The bottom panel recomputes Figure 6b using the NPV of payments owed.

Figure 26: Consumption Functions with Cash-on-Hand and Collateral Grants at Various Dates

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

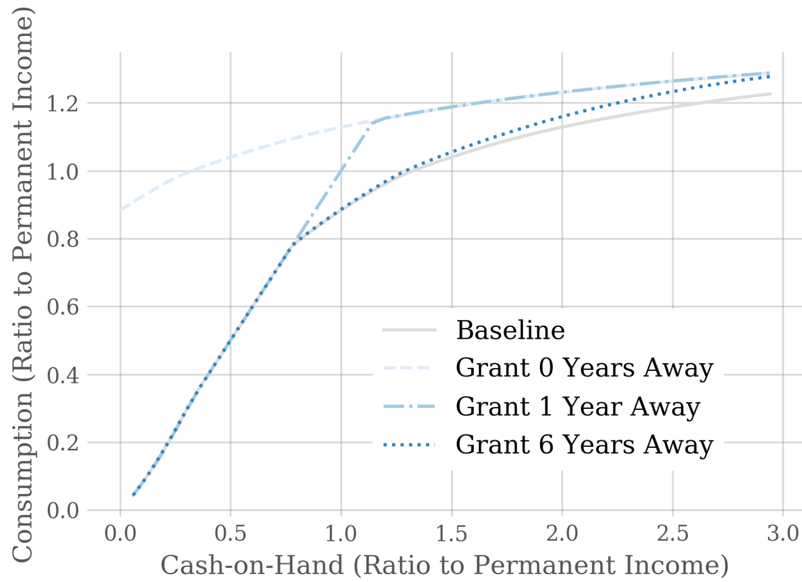

(b) Consumption Function out of Future Collateral

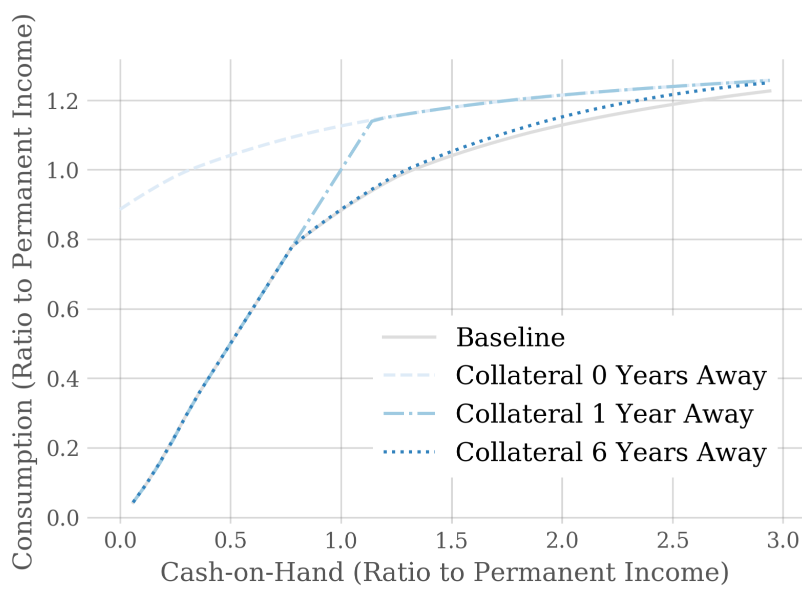

Notes: The top panel plots the consumption function out of cash-on-hand under various alternative scenarios from the model described in online Appendix D. Both the horizontal and vertical axes are measured relative to permanent income. The baseline case considers a household with no home equity (and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

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

	Pre-Match		Post-Match		Normalized
	Mean	SD	Mean	SD	Difference
Income	54,564	24,605	51,225	23,574	-0.14
Home Value	196,992	123,918	178,229	114,756	-0.15
Loan to Value Ratio	150	35	151	35	0.02
Monthly Mortgage Payment	1,716	875	1,552	789	-0.19
Monthly Payment to Income Ratio	0.48	0.12	0.47	0.12	-0.11
Mortgage Interest Rate	0.063	0.020	0.063	0.020	-0.00
Mortgage Term Remaining (Years)	25.9	4.6	25.8	4.7	-0.01
ARM (d)	0.49	0.50	0.46	0.50	-0.05
Months Past Due	11.4	12.9	9.8	11.6	-0.12
Credit Score	584	74	581	75	-0.03
Male (d)	0.57	0.50	0.56	0.50	-0.02
Age	48.8	10.8	48.6	10.9	-0.01
Monthly Payment Reduction ($)	737	544	641	483	-0.18
Monthly Payment Reduction (%)	42	20	41	20	-0.07
Principal Forgiveness Amount	53,072	70,385	46,097	62,187	-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	0.201	0.401	0.00
N	222,695		106,122		

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

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

	Payment Reduction		Payment and Principal Reduction	
	Mean	SD	Mean	SD
Principal Forgiveness Amount	13,215	36,426	80,415	64,678
NPV Payment Reduction	62,990	60,033	97,380	74,116
Monthly Payment Reduction ($)	677	478	673	505
Monthly Payment Reduction (%)	38.5	18.3	41.8	21.4
Loan to Value Ratio	150	33	153	37
Post Modification LTV	148	34	122	29
Monthly Payment to Income Ratio	0.47	0.11	0.47	0.12
Post Modification DTI	0.31	0.03	0.30	0.05
Income	55,641	23,738	52,951	23,665
Credit Score	598	83	578	72
Home Value	205,442	118,803	174,932	111,145
Monthly Mortgage Payment	1,727	803	1,594	780
Mortgage Interest Rate	0.061	0.018	0.064	0.019
Mortgage Term Remaining (Years)	26.2	4.5	26.1	4.4
Male (d)	0.58	0.49	0.55	0.50
Age	48.3	11.2	48.8	10.8
N	35,606		33,890	

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

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

This table reports difference-in-differences estimates of the effect of principal reduction on credit card expenditure in the JPMCI bank 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 ($n_{borrower} = 22,924$). See the text for additional detail on the specification, outcome measures, and sample.

Credit Card Expenditure ($/month)

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (Principal Reduction x Post)	−1.156	−0.993	−3.083	−5.718	−3.565	−3.299
	(3.762)	(3.865)	(3.866)	(4.258)	(4.534)	(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 4: Pareto Improvement at Payment Reduction Discontinuity: Robustness

Scenario	dNPV($)	dNPV (%)	Breakeven Discount Rate (%)
Preferred Estimate	5350	3.31	6.02
Robustness to Default Assumptions			
Low Default Reduction	1944	1.20	5.28
High Default Reduction	8756	5.42	6.65
Optimistic Recovery	-1040	-0.64	4.56
Pessimistic Recovery	8750	5.42	7.09
Robustness to Discounting Assumptions			
Flatter Yield Curve (Actual Swaps Spread)	9263	5.73	6.02
Steeper Yield Curve (Implied Treasury Spread)	5037	3.12	6.02
Discount at Treasury Rates	8741	5.41	6.02
Discount at Swap Rates	14694	9.10	6.02
Robustness to Prepayment Assumptions			
Low Prepayment	5184	3.21	5.88
High Prepayment	5994	3.71	9.60
Crosswalk to Payments Owed NPV			
Payments Owed	-2860	-1.77	3.59
Payments Owed, with Default	8301	5.14	5.63
Payments Owed, with Default and Yield Curve	5279	3.27	5.63

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

B Empirical Appendix

B.1 Effect of Principal Reduction on Default: Representativeness and Robustness

B.1.1 Representativeness of HAMP Participants Relative to Typical Delinquent Underwater Borrowers

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

The median borrower in our sample has a higher LTV than delinquent borrowers in the PSID (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, is important for interpreting our empirical results.

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

B.1.2 Robustness

Balance Plots – Pre-determined covariates trend smoothly through the cutoff, as shown in online Appendix Figure 4. The first five panels show the distribution of pre-modification borrower credit score, monthly income, months past due, monthly mortgage payments to monthly income (payment-to-income, or PTI) 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 PTI) the sign points to slightly worse-off borrowers to the right of the cutoff, while for two variables (LTV and months past due) the sign points to better-off borrowers to the right of the cutoff. The lack of any systematic correlation supports the validity of the design. The only covariate with a statistically significant jump at the 95% level is months past due at application date, and even here the jump is not economically significant. Pre-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.

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

Density – Another relevant issue in regression discontinuity settings is the possibility that the running variable could be manipulated (McCrary 2008). The usual test is to plot a histogram of the running variable to examine whether there is an unusual increase in mass to the right of the cutoff. We show such a plot in online Appendix Figure 5a. 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 "payment & 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.[55]

We attribute the bunching of borrowers at zero to data mis-reporting and drop observations exactly at zero. Online Appendix Figure 5b shows the distribution for the resulting sample, which is our analysis sample. There is no noticeable change in density around the cutoff.

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

Alternative Bandwidths – Online Appendix Figure 7 tests the sensitivity of our results to the bandwidth chosen for the local linear regression. Our central estimates are constructed using the optimal bandwidth from the Imbens and Kalyanaraman (2012) procedure, which is 0.5. The optimal bandwidth recommended by the Calonico et al. (2014) procedure is 0.4. The point estimate is stable out to a bandwidth of 0.6 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. 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.[56]

[55]Two 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, if servicers were 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.

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

B.2 31 Percent Payment-to-Income Discontinuity: Robustness and Other Outcomes

B.2.1 Robustness

Balance Plots – We show the trend in pre-determined covariates through the cutoff in online Appendix Figure 12. The first four panels show the LTV ratio, months past due, borrower monthly income, and pre-modification monthly payment around the cutoff. These balance plots differ in two ways from the balance plots for the discontinuity for principal reduction. First, unlike in the matched HAMP credit bureau dataset used for the investor NPV strategy, borrower credit score is not available in the JPMCI bank dataset. Second, we cannot show balance on PTI because it is the running variable. Instead, we show balanced on pre-modification monthly payment.

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

Density – Online Appendix Figure 13 shows that borrower density is also smooth around the cutoff.

Alternative Bandwidths – Our point estimate of $\hat{\tau}$ from equation 4 is that an extra 1 percent payment reduction reduces default rates in the two years after modification by 0.34 percentage points. Online Appendix Figure 16 tests the sensitivity of our results to the bandwidth chosen for the local linear regression. Our central estimates are constructed using the optimal bandwidth from the Imbens and Kalyanaraman (2012) procedure, which is 0.057 points of PTI. We test alternative bandwidths between 0.01 and 0.1 and find that the point estimate is stable.

Adjusting for Upper Bound of Potential Principal Forgiveness Impact – If we take the upper bound of our 95 percent confidence interval for the impact of principal reduction on default from section 3.3, and scale it by the amount of relative principal increase received by borrowers just below the 31 percent PTI cutoff, we find that a principal increase of this magnitude would have led to at most a 1.2 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 7.2 percentage points without the principal increase rather than 6.0 percentage points, or alternatively that each 1 percent reduction in payment reduced default rates by 0.41 percentage points (1.3 percent), similar to our baseline estimate of 0.34 percentage points (1.1 percent).

B.2.2 Impact of Payment Reduction on Consumption

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

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

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

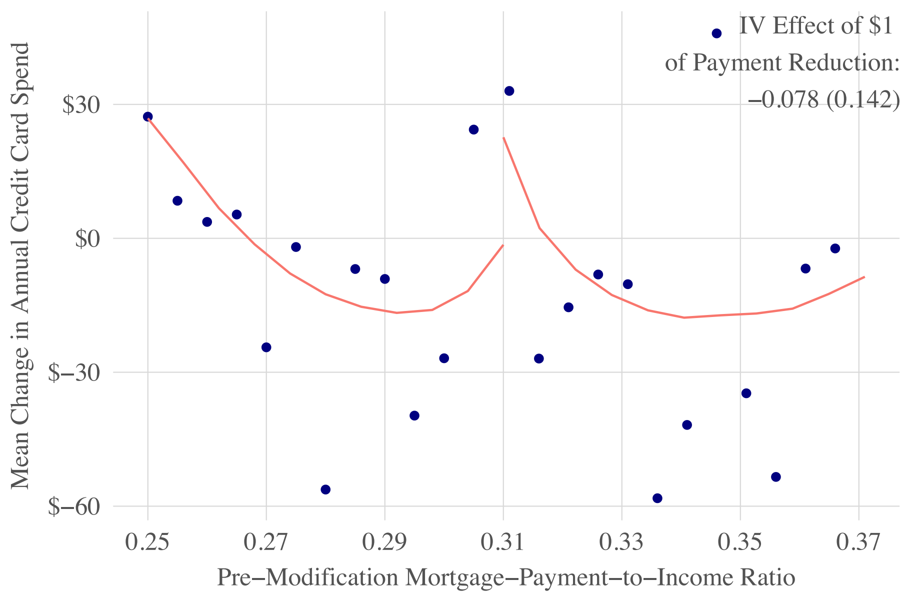

Notes: This figure shows the reduced form of the estimated impact of payment reduction on credit card expenditure using the 31 percent PTI regression discontinuity strategy in the JPMCI bank dataset. The blue dots are conditional means for 12 bins on each side of the cutoff. The red line shows the predicted 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 5.2. This strategy is unable to detect economically meaningful changes in expenditure.

C Net Present Value Calculations

In this section we provide more detail on our NPV calculations. Section C.1 discusses the basic setup, which is applicable to our analysis in Sections 3.1 and 6.2.2. Section C.2 provides more detail for calculating the NPV to lenders at the HAMP eligibility cutoff, which we use in Section 6.2.2.

C.1 Net Present Value of Expected Payments

We use two equations to estimate the NPV of the loan. Equation 5 estimates the value of a mortgage that "cures," meaning that the borrower repays on time or early:

$$NPV_{Cure}(\delta) = \sum_{i=1}^{T} \frac{1}{(1+\delta)^i} \left[(UPB_{i-1} - Prin_i)(s_{i-1} - s_i) + (Prin_i + I_i)s_{i-1}\right] \tag{5}$$

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

We use the Treasury NPV model to estimate annual prepayment rates s_i. This is the same model used by servicers to calculate the expected cash flows to lenders under various HAMP modification types, which we use for identification in Section 3.1, and is documented in U.S. Department of Treasury (2015). The model uses a logit equation for predicting prepayment rates (Section V of U.S. Department of Treasury 2015) and we use the coefficients for owner-occupied homes reported in Appendix C of U.S. Department of Treasury (2015) for borrowers that are 90+ days delinquent at modification date.[57] Adding prepayment shrinks the gains from treatment in two ways. First, if prepayment rates are the same in the treatment and control group, prepayment reduces the value of the investment when the average interest rate on the mortgages is greater than the investor's discount rate. We argue below that this is likely to be the case. Second, the Treasury NPV model assumes that borrowers who owe more debt relative to the value of their home are less likely to prepay. However, the quantitative magnitude of these forces is small and overall, incorporating prepayment risk has little effect on default rates.

Our second key equation incorporates default risk into our NPV estimate. The NPV of the mortgage is the weighted average of the value of the mortgage if the borrower pays (described above) and the value of the mortgage if borrower defaults:

$$\begin{aligned} NPV = {} & P(Cure) * NPV_{cure} + \\ & P(D) * [P(liquidate|D) * NPV_{liquidate} + (1 - P(liquidate|D))NPV_{cure}] \end{aligned} \tag{6}$$

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

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

C.2 Investor Valuation at HAMP Eligibility Discontinuity

Our choice of the discount rate δ for future cash flows depends on the maturity of the mortgage. Recall that assignment to treatment involves an extension in the term of the mortgage and 80 percent of loans in the treatment group last 40 years after modification. Ideally, we would use the interest rate on 40-year mortgages to discount these cash flows, but unfortunately we are unaware of any publicly available data source with prices for 40-year mortgages. Instead, we estimate the price of a 40-year mortgage by using a simple functional form to extrapolate from the price of 15-year and 30-year fixed mortgages sold by Freddie Mac. The JPMCI payment reduction sample includes modifications from October 2011 through January 2014. The average 15-year rate is 3.06 percent during this period and the average 30-year rate is 3.84 percent.[58] We fit an equation $r = \alpha + \beta log(term)$ to these data and estimate a hypothetical 40-year mortgage rate of 4.16 percent.[59] In the robustness analysis below, we explore the sensitivity of our estimates to alternative assumptions about the discount rate and the yield curve.

We estimate the effect of treatment on default rates using our causal estimates from the regression discontinuity design and HAMP performance data. We estimate a 90-day default rate in the two years after modification of 25.6 percent for the treatment group and 31.6 percent for the control group. Among HAMP modifications done in 2010, the default rate is 28.1 percent two years after modification and 45.6 percent five years after modification (U.S. Department of the Treasury 2017), for a ratio of 1.62. We project default rates five years after modification in our data by multiplying our estimated default rates by 1.62. This calculation assumes that payment reduction is equally effective in years three, four, and five. We project the default rate will be 38.6 percent in the treatment group and 46.2 percent in the control group. We explore alternative assumptions for the impact of treatment on the default rate in the robustness analysis below.

To estimate the probability that a default results in a loss for the investor, we use HAMP performance data. Among HAMP modifications that are disqualified due to default, 26 percent end up in foreclosure, 14 percent end in a short sale, 18 percent self-cure, 33 percent get a proprietary modification, and 10 percent have delayed action, such as a borrower going through bankruptcy (U.S. Department of the Treasury 2017). Of loans whose status is fully resolved, 45 percent are foreclosed on, 24 percent end in a short sale and 31 percent self-cure. We assume that loans which get a proprietary modification or delayed action ultimately have the same distribution of final outcomes. We explore alternative assumptions in the robustness analysis below.

Unfortunately, HAMP does not collect data on losses after disqualification so we draw on GSE performance data to estimate losses. The GSEs report losses on loans that are "liquidated" via either foreclosure or short sale. Goodman and Zhu (2015) document that GSE losses are quite similar on foreclosures and short sales. We use performance data from loans liquidated in 2011 because that was the year in which the GSEs experienced the largest number of liquidations. In that year, the Fannie Mae reported losses at liquidation equal to 41 percent of the unpaid balance on the loan (Fannie Mae 2018). However, this

[58] This is quite similar to the average 30-year rate of 4.11 percent during the time period when mods were performed for our principal reduction sample.

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

includes reimbursements from third parties (mortgage insurers and mortgage originators) to the GSEs equal to 15 percent of the unpaid balance of the loan. Altogether, investors lost 56 percent of the unpaid balance of the loan at liquidation. We explore alternative assumptions in the robustness analysis below.

Our estimates imply a gain to the investor from assigning a borrower to treatment. Recall that treatment is essentially a loan to the borrower in the form of lower mortgage payments for 22 years which is offset by additional mortgage payments extending beyond the pre-modification term of the loan. The change in the NPV arising from this maturity extension treatment is $5,350, as shown in online Appendix Table 4. This is equal to a 3.4 percent increase in the NPV of the loan.

As an alternative to the NPV calculation, we also report the discount rate an investor would need to break even on providing treatment to a group of mortgages. While the prior calculation assumed that the lender discounted future mortgage cash flows at our best estimate of the market interest rate, an alternative approach allows us to be agnostic as to the lender's discount rate. The NPV of a mortgage that cures is a function of the discount rate δ, as shown in equation 5 and the expected NPV of all mortgages in equation 6 relies on this, so we can rewrite NPV in equation 6 as a function $NPV(\delta)$. Then, we can solve for the discount rate that satisfies the lender's indifference condition such that the change in NPV from offering the treatment modification is the same as the change from offering the control modification:

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

In our baseline specification, we estimate that a lender that discounts the future annually by 6.02 percent will be indifferent between offering this modification. This implies that a lender with an annual discount rate less than 6.02 percent will be better off offering the treatment modification.

We explore the robustness of our NPV and discount rate estimates to alternative assumptions on default rates, recovery rates on losses, discounting, and prepayment in online Appendix Table 4. Across almost all scenarios, we find that the NPV of the loan to the investor increases from assigning a loan to treatment instead of control. First, we explore the impact of alternative assumptions about the impact of treatment on mortgage default. Using the lower and upper bounds of our 95 percent confidence interval, we estimate the change in NPV ranges from $1,944 to $8,756.

Second, we show that impact of treatment on NPV is sensitive to our assumptions on the recovery rate on defaulted loans, but is always positive or statistically indistinguishable from zero. Our specification with the most optimistic recovery rates assumes that *every* proprietary modification and *every* action pending self cures, meaning that there is a 61 percent self cure rate, and uses the highest possible recovery rate on GSE loans during the crisis, which was a 48 percent loss in 2009. Our specification with the most pessimistic recovery rates assumes that all proprietary modifications and action pending ends in liquidation, meaning that there is an 18 percent self-cure rate, and the lowest possible recovery rate on GSE loans, which was a 61 percent loss in 2014. Treatment in the optimistic scenario causes an NPV loss to the investor of -$1,222, while in the pessimistic scenario it causes an increase of $8,750. Note that -$1,222 is indistinguishable from zero given our standard errors and therefore the criteria for a Pareto improvement (which is that at least one party is better off and no party is worse off) is still satisfied in this scenario.

Third, we show the impact of using alternative methodologies for estimating the discount rate. Intuitively, the treatment modification defers cash flows from the present to the future and investors require a higher rate of return for deferring these cash flows. Recall that the average interest rate for a 30-year fixed rate mortgage during our sample period is 3.84 percent and in our baseline specification we estimated an additional 32 basis points for a 40-year mortgage. At one extreme, an alternative methodology which relies on a comparison of 30-year and 40-year loans is the swap rate where the yield curve is flatter and the average spread in our sample period is only 2 basis points. At the other extreme, projecting hypothetical spreads using interest rates on debt issued by the U.S. Treasury implies a steeper yield curve with an additional 34 basis points for 40-year mortgage. Both of these projections are shown in online Appendix Figure 21. This flatter yield curve implies a change in NPV of $9,263 and the steeper yield curve implies a change in NPV of $5,037. The figure also shows that, if anything, the log functional form overestimates the term premium at higher maturities. Forty-year maturities are actually observed for swaps and corporate bonds. For these, the “implied” spread between 30 and 40 year maturities using the log functional form assumption are much larger than the actual spreads.

It may be preferable to use the risk-free rate to discount cashflows in our model. The argument for using the risk-free rate here is that lenders offering mortgages charge a premium over the risk-free rate in order to compensate the lender for prepayment risk and default risk. However, our expected payments NPV calculation already takes into account default and prepayment risk. The average rate on 30-year Treasury notes during this time period is 3.17 percent and we project that that the rate on a 40-year note would be 3.52 percent. The average rate on fixed-for-floating swaps is 3.00 percent for 30 years and 3.02 percent for 40 years. Under these assumptions, we calculate changes in NPV of $8,741 and $14,694 respectively. The value is greater to the investor under this scenario because a maturity extension delays cashflows and switching to a lower discount rate makes cashflows far in the future more valuable.

Fourth, we show that prepayment rates have little effect on the change in NPV from treatment. At one extreme, we assume an annual prepayment rate of 0.9 percent (the observed prepayment rate after HAMP modification) for the life of the loan. At the other, we assume an annual prepayment rate 6.8 percent (the observed prepayment rate on all Fannie Mae loans 1999-2017). The change in NPV varies from $5,184 under the low prepayment scenario to $5,994 under the high prepayment scenario.

Finally, we crosswalk our expected payments NPV estimate to the payments owed NPV estimates reported elsewhere in the text. Recall that the investor’s return from treatment is a $5,350 gain in terms of expected payments NPV, but a loss of $2,860 when using the payments owed NPV estimate reported in online Appendix Figure 14. (To be precise, the figure shows that the investor loses $2,860 more from treatment.) This assumes that the loan is repaid on schedule (no default or prepayment) and the investor discounts cashflows at 4.11 percent annually.

The gain from treatment in the expected payments NPV is larger than the payments owed NPV primarily because of the reduction in defaults associated with treatment. If we take the payments owed NPV calculation and allow for default as described in our baseline specification, we find a gain of $8,301. This gain is partially offset by the longer time horizon for cash flows under treatment; incorporating the yield curve calculation above (28 additional basis points for 40-year mortgages over 30-year mortgages) reduces the gain to $5,279. Finally, incorporating prepayment risk has little affect on our estimates, and we recover our benchmark estimate of $5,350.

D Partial Equilibrium Life-cycle Model with Housing

D.1 Setup

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

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

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

where c_t is non-housing consumption, d_i is an indicator variable equal to 1 if the household defaults, and ψ is a utility cost of defaulting. This additive default cost follows the structure in Campbell and Cocco (2015), Hembre (2018), Kaplan et al. (2017), 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.[60] 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.[61] Renters pay the user cost of housing for the equivalent home. Thus, housing payments are given by

$$h_{itj} = \begin{cases} M_{i1} \frac{r(1+r)^{30}}{(1+r)^{30}-1} + \tau_p P_{it} + \tau_m P_{i1}, & j = owner \\ \left(r - g + \tau_p\right) P_{it} + \tau_m P_{i1}, & j = renter \end{cases}. \tag{7}$$

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

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

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

where $(1-\phi)$ is the fraction of a house's value that can be used as collateral.[62] Renters are not able to borrow.

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

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

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

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 $\mathbb{E}\left[\theta_{it}\right] = 1$. 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, is[63]

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

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

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

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

We model default as a one-shot decision. Households begin the first period with a given mortgage, home price, and asset level. They then observe their first-period income shock, and decide whether to default or hold the house until retirement. This provides a simple way to analyze the short-term default decisions which we study empirically in Section 3.3. In Section D.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.

D.2 Parameterization

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

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

$$\theta_{it} = \begin{cases} b & \text{with probability } p \\ \frac{\delta_{it}(1-b \cdot p)}{1-p} & \text{with probability } (1-p) \end{cases} \tag{8}$$

borrowing limit is given by

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

in order to prevent forced deleveraging of liquid assets.

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

Table 5: Baseline Model Parameter Values

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

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 $\sigma_\delta = 0.14$.[64] 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 D.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 percent. This matches the average 30-year mortgage rate from the Freddie Mac Conforming Loan Survey for the period 2010-2014 (4.1 percent) minus the average expected inflation on 30-year Treasury bonds over the same period (2.1 percent). We assume a collateral constraint ϕ of 0.2, such that homeowners can only borrow up to 80 percent of the value of their home. This matches the caps for cash-out refinancing from Fannie Mae and Freddie Mac, 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 percent, which is the average from FHFA's national index between 1991 and 2010, as well as the expected annual price growth from home price futures in 2011, though we test the sensitivity of our results to alternative house price growth rate paths. We follow Himmelberg et al. (2005) and set the property tax rate to 1.5 percent and the maintenance cost to 2.5 percent. These parameters generate a first-period user cost of housing of 5.1 percent, similar to the empirical estimates in Diaz and Luengo-Prado (2008) and Poterba and Sinai (2008), who find 5.3 percent and 6 percent, respectively.

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

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

double-trigger behavior.[65]

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

D.3 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.

D.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.[66] 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

[65] Our choice of a high γ ensures that agents default when they are hit with a bad income shock, but do not default under regular economic circumstances. The model exhibits this behavior because when γ is high, the value function for the agent paying her mortgage is much more concave than the value function for the agent who is defaulting, generating a region where default is sensitive to income. In contrast, when γ is low in our model, LTV is the primary determinant of default decisions, which is inconsistent with our empirical findings. We discuss this choice in more detail in Section D.3.

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

119, the median pre-modification LTV for borrowers in our regression discontinuity analysis (Table 1a). We then vary the LTV, holding mortgage payments for households that have not defaulted fixed for five years, after which payments fall according to the annuity formula in equation 7 applied to the new mortgage balance.

Online Appendix Figure 28a shows that for a given current payment level and LTV ratio, there is a cash-on-hand level below which households will find it optimal to default. The more underwater the household, the smaller the income shock necessary to push them to default. For borrowers in our baseline scenario, the income cutoff for defaulting is both low and relatively insensitive to debt levels. In particular, below LTVs of about 150, low-asset borrowers will only default if their income is less than three-quarters of its permanent level, a shock of about two standard deviations. This means that default is most likely to occur for borrowers who are hit with "unemployment," the large liquidity shock in our income process.

We find that default rates are insensitive to principal reduction for the typical borrower. Online Appendix Figure 28b plots the default rate in the first period after modification for borrowers with various amounts of principal reduction. In our baseline case, additional principal reduction is ineffective below an LTV of about 160. For such moderately underwater borrowers, the gain from defaulting is not worth the cost unless they are hit by a liquidity shock. However, far underwater borrowers have much higher default rates because they default even in the absence of liquidity shocks.

Figure 28: Default and Mortgage Debt Levels

(a) Default if Income + Assets < Threshold

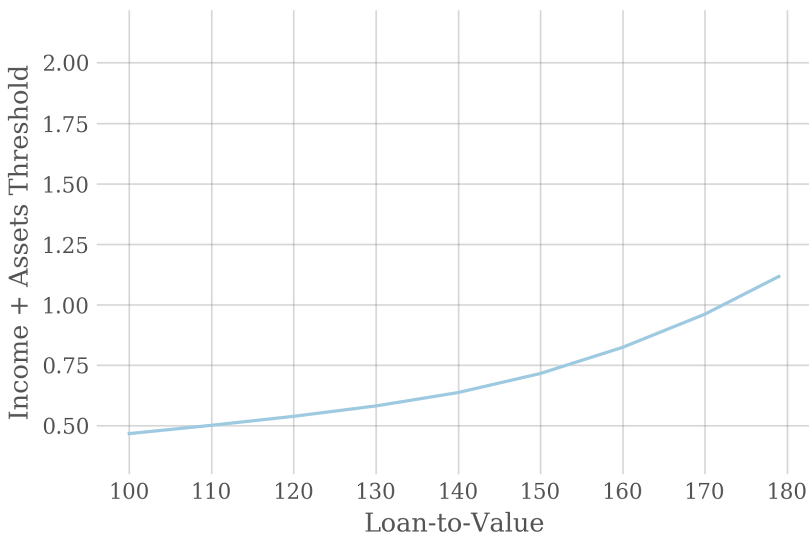

(b) Default Rate

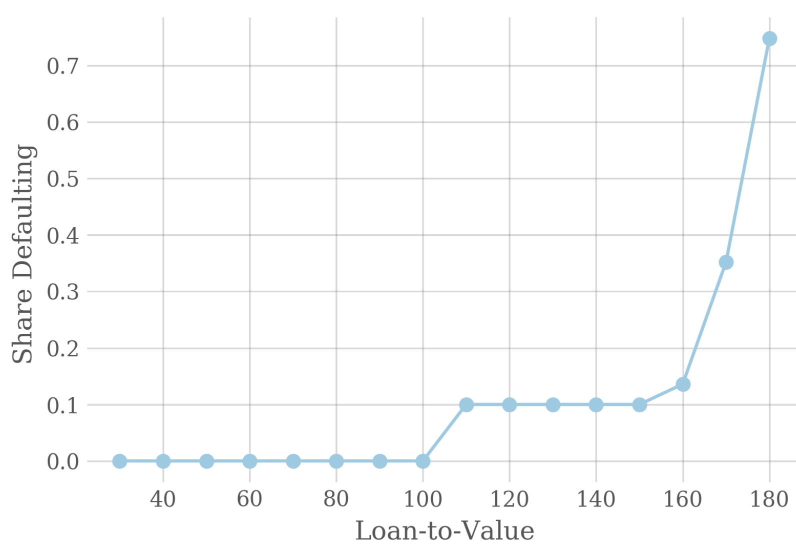

Notes: The top panel plots the cutoff thresholds for borrower default decisions. The vertical axis is relative to permanent income. The line shows the baseline assumptions as described in equation (7). For borrowers with a given LTV ratio, the line shows the cash-on-hand (income plus assets) threshold below which borrowers decide to exercise their default option. The bottom panel plots default rates by LTV ratio. Default rates are calculated by taking the default thresholds shown in the top panel and integrating over the distribution of income shocks described in equation (8).

D.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. 2018). 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 2018).

In our model, agents are optimizing. At moderate levels of underwaterness, it is only optimal for agents with large liquidity shocks (i.e., unemployment in our model) to default. Default is insensitive to negative equity in this region because the costs of default are high and the gains for an employed agent are low. However, beyond about 160 LTV, their optimizing behavior generates a steep causal relationship between LTV and default. These borrowers are defaulting for what is sometimes referred to as "strategic" reasons, 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. (2017), who show that the median homeowner without an income shock does not default until their LTV is greater than 174.[67]

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 (Guvenen et al. 2014).[68] 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.[69]

[67]Similarly, 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.

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

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

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 online Appendix Figure 29, and show that this does generate a smooth relationship between LTV and default. However, our empirical results, which find that default is insensitive to LTVs for moderate amounts of underwaterness, reject this parameterization of our model.

Figure 29: Default with Heterogeneous Utility Cost of Default

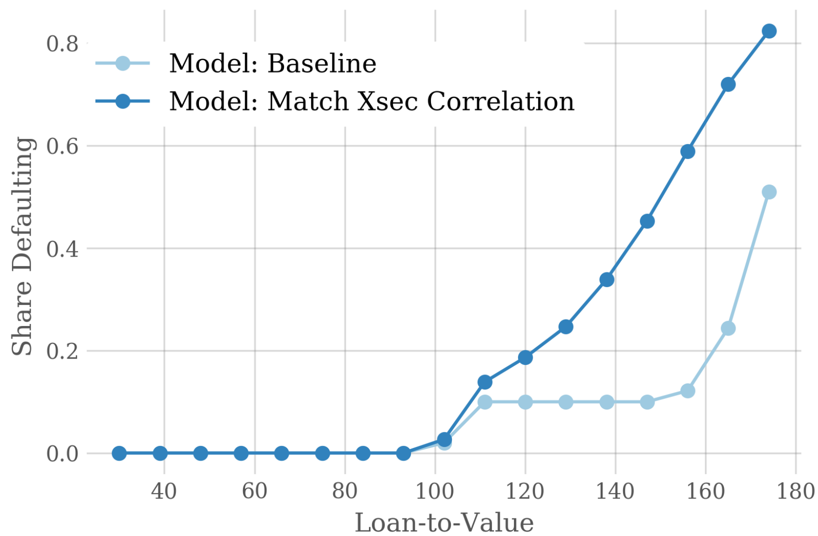

Notes: This figure plots default rates by LTV ratio in the model under alternative parameterizations. The LTV is moved according to the same policy exercise described in the notes to Figure 30a. The baseline parameterization corresponds to that in Table 5. The "Match Xsec Correlation" assumes a distribution of default costs across the population instead of a constant default cost.

D.4 Consumption

In this section, we provide more detail on our consumption-related results.

D.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 permanent income, which is the median non-housing wealth for all

default below an LTV cutoff and *all* agents default above the LTV cutoff.

homeowners in the 2007 Survey of Consumer Finances (SCF).[70] 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).[71]

Online Appendix Table 6 reports the average MPC out of an additional dollar of housing equity for the average borrower as well as for high-leverage (but still above-water) borrowers. We find MPCs of 8 and 15 cents, respectively. These are similar to the average MPC for homeowners of 9 cents reported in Mian 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.

Table 6: Housing Wealth MPC in Model and External Benchmarks

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

Notes: This table shows the marginal propensity to consume out of changes in housing wealth in the model relative to the estimates in the external benchmark from Mian 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).

D.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 $\left(w_{iT_y} = P_{iT_y} - M_{iT_y}\right)$, and the vectors of housing payments and collateral constraints for the rest of life $\left(\vec{h}_i, \vec{\underline{a}_i}\right)$. We can then decompose the effect of a change in mortgage debt level at date t in the following way:

$$\begin{aligned}
\frac{dc_{it}\left(m_{it}, w_{iT_y}, \vec{h}_i, \vec{\underline{a}_i}\right)}{dM_{it}} &= \frac{\partial c_{it}}{\partial w_{iT_y}} \cdot \frac{\partial w_{iT_y}}{\partial M_{it}} + \sum_{s=t}^{T} \frac{\partial c_{it}}{\partial h_{is}} \cdot \frac{\partial h_{is}}{\partial M_{it}} + \sum_{s=t}^{T} \frac{\partial c_{it}}{\partial \underline{a}_{is}} \cdot \frac{\partial \underline{a}_{is}}{\partial M_{it}} \\
&= \underbrace{MPC_{t,w_{iT_y}} \frac{\partial w_{iT_y}}{\partial M_{it}} + \sum_{s=t}^{T} MPC_{t,h_{is}} \cdot \frac{\partial h_{is}}{\partial M_{it}}}_{\text{Future cash-on-hand effect}} + \underbrace{\sum_{s=t}^{T} MPC_{t,\underline{a}_{is}} \cdot \frac{\partial \underline{a}_{is}}{\partial M_{it}}}_{\text{Collateral effect}} . \quad (9)
\end{aligned}$$

Equation 9 shows that a reduction in debt levels affects today's consumption through two channels. The first is a future cash-on-hand effect. Reducing mortgage debt increases

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

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

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

D.4.3 Difficulty of Accessing Housing Wealth During Recovery

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

Second, the time series of mortgage credit origination shows that credit constraints had tightened during the recovery. Online Appendix Figure 22a shows mortgage originations by borrower credit score from 2000 to 2015. This covers all mortgages, including second mortgages and home equity lines of credit (HELOCs). It shows that originations dipped sharply after 2007, and for low-credit score borrowers, originations have never recovered. Borrowers receiving HAMP principal reductions had mean FICO scores of 579, with 85 percent below 660, the cutoff for the red line in the figure. This evidence suggests that even with positive equity, the low credit-score borrowers in our sample may have been unlikely to obtain additional housing-related credit. This is further reinforced by the Online Appendix Figure 22b, which shows the time series of average CLTV ratios for borrowers able to obtain HELOCs in a given year. The average CLTV ratio fell 20 points between 2004 and 2009, indicating a tightening of underwriting constraints. Mian and Sufi (2014) argue that tightening credit conditions could explain why the house price recovery from 2011 onward did not contribute significantly to economic activity, since in this case the borrowing channel is restricted. Our results support this hypothesis for underwater borrowers. Furthermore, Agarwal et al. (2016) show that credit expansions during the recovery were more likely to benefit higher-FICO borrowers, precisely those least likely to respond by increasing borrowing.

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

Online Appendix Figure 30 shows the evolution of borrowing limits and mortgage payments around principal reduction for the average borrower according to our model and using the assumptions described above. We consider an average household with first period income $y_t = 0.85$ units of permanent income, based on Bernstein (2017) who finds that borrowers receiving mortgage modifications during the recent crisis had temporarily low incomes. We set initial LTV equal to 150, the median pre-modification LTV for borrowers receiving principal reduction in our difference-in-differences analysis.[72] For our treatment group, we then reduce their mortgage balance by $70,000, bringing them to an LTV of 106.[73]

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

[73]The median LTV post-modification in our data is actually 114, because borrowers' unpaid mortgage payments are capitalized into the new mortgage balance. We abstract from this in our model, though it

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

---

would only serve to further reduce the effect of principal reduction.

Figure 30: Effect of Modeled Principal Reduction on Borrowing Limits and Mortgage Payments

(a) Borrowing Limits

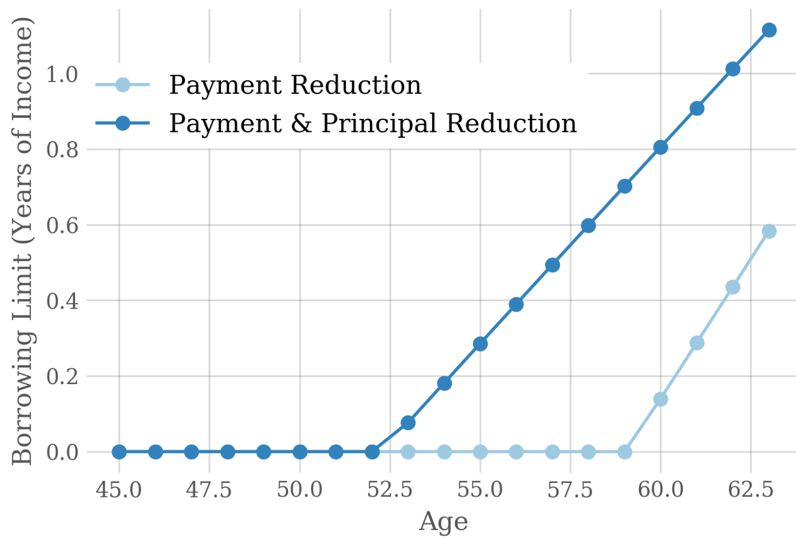

(b) Mortgage Payments

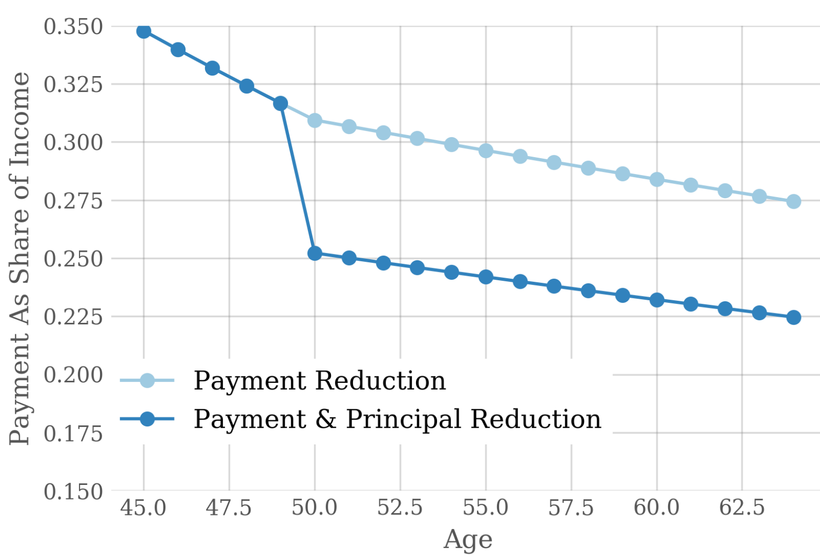

Notes: This figure shows the effect of the modeled principal reduction policy on borrowing limits and mortgage payments. We assume homeowners receive modifications at age 45. We set initial LTV equal to 150. For our treatment group, we then reduce their mortgage balance by $70,000, bringing them to an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

D.4.4 Consumption Response to Principal Reduction Under Alternative Parameterizations

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

Table 7: MPC out of Principal Reduction in the Model

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

Notes: This table compares the MPC out of principal reduction in the model under alternative parameterizations to the MPC calculated in our data (discussed in Section 4.2). The "Baseline Model" corresponds to the parameterization shown in Table 5 and the modeling of principal reduction policy discussed in Section D.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%.

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.[74]

References

Agarwal, S., Chomsisengphet, S., Mahoney, N., and Stroebel, J. (2016). Do Banks Pass Through Credit Expansions? Asymmetric Information and the Marginal Profitability of Consumer Lending During the Great Recession. *SSRN Electronic Journal*.

Bernstein, A. (2017). Negative Equity, Household Debt Overhang, and Labor Supply. SSRN Scholarly Paper ID 2700781, Social Science Research Network, Rochester, NY.

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

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. (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 Universtiy.

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.

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

Corelogic (2016). Home Equity Lending Landscape. White Paper.

Department of Housing And Urban Development, U. (2016). Housing Scorecard. HUD Report, September 2016, Availabel 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.

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

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

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.

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

Goodman, L. S. and Zhu, J. (2015). Loss Severity on Residential Mortgages: Evidence from Freddie Mac's Newest Data. SSRN Scholarly Paper ID 2784290, Social Science Research Network, Rochester, NY.

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. (2018). How do Forecolosures Exacerbate Housing Downturns? Working Paper, Boston University and Stanford University.

Guvenen, F., Ozkan, S., and Song, J. (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy*, 122(3):621–660.

Hembre, E. (2018). HAMP, Home Attachment, and Mortgage Default. Working Paper, University of Wisconsin.

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

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

Kaplan, G., Mitman, K., and Violante, G. L. (2017). Consumption and House Prices in the Great Recession. Working Paper.

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.

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

Mae, F. (2018). Statistical Summary Tables.

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

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. (2014). House Price Gains and U.S. Household Spending from 2002 to 2006. Working Paper 20152, National Bureau of Economic Research.

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.

Schelkle, T. (2016). Mortgage Default During the U.S. Mortgage Crisis. Working Paper, University of Cologne.

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

United States Department of Housing and Urban Development (2010). Economic Impact Analysis of the FHA Refinance Program for Borrowers in Negative Equity Positions. Technical report, US Department of the Treasury, Washington, DC.

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

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