Online Appendix for
"How Do Mortgage Refinances Affect Debt, Default, and Spending? Evidence from HARP"

Joshua Abel and Andreas Fuster

A Additional Test for Strategic Behavior by Lenders/Servicers

In Section III.A of the main text, we argued that there was little evidence for lenders or servicers trying to alter securitization speeds in response to the HARP announcement in March 2009. In this section, we present additional tests, based on the idea that the strategic incentives to alter securitization speeds varied depending on LTV: the HARP option is relatively more valuable when a borrower has a high LTV as opposed to a low one, as HARP is superfluous to low-LTV borrowers. This implies that even if there is no evidence of an overall shift of guarantees from after the cutoff date to before, strategic behavior may be masked if, say, servicers ensure that their high-LTV borrowers are guaranteed on time and delay the process for low-LTV borrowers. As a result, we would see borrowers in this window being guaranteed relatively quickly if their origination LTV is high.[1]

To investigate this, we look at all FRM originations in the January-June 2009 window that are guaranteed by a GSE by March 2010. We first study this issue graphically. Figure A-2 shows that the guarantee lags (months between origination and GSE guarantee) over this origination window are distributed essentially identically across the different LTV bins. This makes it hard to imagine that there would have been much strategic behavior.

In Table A-7, we turn to regression analyses to more precisely quantify differences in guarantee lags within loans originated in the same month. In columns (1) and (2) we regress guarantee lag on dummies for different LTV bins, cohort fixed effects, and (in column 2) other loan-level controls such as credit score, interest rate, or loan amount. The coefficients in column (1) are small—the magnitude corresponds to less than 1 day—and non-monotonic. In column (2), coefficients become slightly more significant but remain economically small. For instance, a loan with an LTV above 90% is guaranteed just 0.050 months—1.5 days—earlier than one with an LTV below 60%. Additionally, the low values for R^2 tell us that LTV (as well as the other variables) plays only a very small role in determining how long a loan takes to be guaranteed.

Columns (3) and (4) present a stronger test by interacting LTV with "MaxLag," the maximum number of months a loan can go before being guaranteed while still being eligible for HARP (e.g. for the April cohort, MaxLag $= 1$, while for the January cohort, MaxLag $= 4$). The motivation for these specifications is that if there is strategic guarantee activity occurring, high-LTV loans should be affected more if they are originated close to the cutoff date (MaxLag small). Intuitively, strategic motives are less relevant early in

[1] One could also imagine the incentive going in the other direction, if servicers want to minimize the likelihood that their borrowers refinance (since that may lead to a loss of servicing fees). In this case as well, however, we would expect guarantee speeds to vary systematically with LTV.

the origination window because the vast majority of loans will be guaranteed before the cutoff date anyway, while loans originated close to the cutoff date can have their eligibility status affected by small changes in their guarantee speed. Notice that the interaction coefficients for the four highest LTV bins have alternating signs, a result that would be hard to justify with any story of strategic sorting. Furthermore, the magnitudes of the coefficients are again small.[2]

Finally, if we look at the binary HARP eligibility indicator itself as the outcome rather than the guarantee lag, there is again no evidence of strategic manipulation. This is arguably the more important outcome, as manipulation of the guarantee lag is only relevant to the extent that it has an impact on borrowers' eligibility. Columns (5) and (6) show that there is no meaningful difference in eligibility across the LTV distribution. Again, the coefficients on high-LTV bins have alternating signs, and the magnitudes of the coefficients are mostly well under 1 percentage point.[3] In all then, LTV does not seem to be an important determinant of HARP eligibility.[4]

All of this evidence, together with the arguments in Section III.A and daily guarantee evidence shown in Karamon, McManus and Zhu (2016), provides strong support for the argument that there was not strategic sorting of guarantees caused by the HARP eligibility cutoff date.

B Estimating Dynamic Effects of Refinancing

In Section III.E, we discussed how our main empirical specification imposes the restriction that the treatment effect of refinancing does not depend on how long ago the refinance occurred. In this section, we discuss how one can relax that restriction.

Recall that the treatment variable, Refied_{it}, is an indicator for whether borrower i refinanced in some month $\tau \leq t$. We instrument for this using HARP eligibility interacted with time (quarter) to estimate the causal effect. Intuitively, we find a treatment effect if the eligible and ineligible groups differ substantially in some outcome (e.g. first difference of auto debt) in the months when the eligible group is far more likely to have refinanced (basically, post-2011). However, in principle, the empirical framework allows for estimation of a more flexible response, where the treatment effect can differ depending on how long ago the borrower refinanced. This is done by simply including multiple treatment variables of the form:

$$\text{Refied}_{it}^{j_1,j_2} = \begin{cases} 1 & \text{if refinanced between } j_1 \text{ and } j_2 \text{ months ago} \\ 0 & \text{otherwise.} \end{cases} \tag{A-1}$$

For example, we can define two such variables, $\text{Refied}_{it}^{0,12}$ and $\text{Refied}_{it}^{13,\infty}$. By jointly estimating these, we allow the treatment effect in the first year to differ from its longer-term impact. The same set of instruments is used to estimate the causal effects: HARP eligibility interacted with time. Intuitively, since the

[2]The coefficient 0.007 in the last row of column (4) implies that a high-LTV loan ($\geq 90\%$) will be guaranteed $4 \times 0.007 = 0.028$ months (0.8 days) faster if originated in May rather than January, a tiny effect.

[3]These columns have large R^2s because, as discussed in the text, the date of a guarantee, and therefore eligibility, is determined to a very large extent simply by origination cohort—not LTV.

[4]In unreported regressions, we have further found that eligible high-LTV loans were not guaranteed meaningfully faster (nor slower) than eligible loans with lower LTVs.

refinancing boom occurred in 2012, the estimator will assign a strong effect to $\text{Refied}_{it}^{0,12}$ if the outcome variable is starkly different between the eligible and ineligible groups in 2012/2013, and it will assign a strong effect to $\text{Refied}_{it}^{13,\infty}$ if the difference between the groups is strong later in the sample period.

Table A-8 shows the results of estimating this more flexible specification for our main outcomes. The first takeaway is simply that the standard errors on the difference between the coefficients for the two periods ($\leq$ 12 months, $>$ 12 months) are very large. Intuitively, we are able to identify a statistically significant effect in the main results of the paper because we are using all of the post-refinancing-boom data (2012-2016) to identify a single coefficient—the one on Refied_{it}. The more flexible specification forces each coefficient to be estimated based on only some of the post-2011 variation, and so it is under-powered. This is why we do not focus on these results in the main text, and refrain from attempting to allow for even more flexible models.

Nonetheless, the results in Table A-8 are suggestive of some interesting effects. In particular, while the difference is not statistically significant, it seems that the reduction in mortgage defaults was not strong in the first year after a borrower refinanced but quite a bit stronger thereafter. This may be because, as discussed at length in the text, refinancing borrowers tended to be in relatively good financial health and so were not at heavy risk of default at the time of the refinance. However, the resulting payment reduction may have helped them withstand negative shocks that subsequently befell them.

Furthermore, it seems as though borrowers responded to their improved cash flow by immediately drawing on their retail consumer debt instruments, and this may be the case for auto loans as well, though the difference is not statistically significant. HELOC debt seems to show the opposite pattern, with the strong response being delayed by at least a year.

C Heterogeneity in HARP's Overall Impact on Borrowers

In Sections IV, V and VI, we studied heterogeneity in treatment effects of refinancing on default outcomes and debt accumulation, as well as heterogeneity in the take-up of the refinancing opportunity offered by HARP. In this section, we present evidence from an alternative econometric strategy to separately analyze the complementary question of how HARP eligibility itself differentially affected different borrower types in terms of their refinancing propensity and debt outcomes.

In a textbook treatment effect estimation scenario, the local average treatment effect (LATE) is equal to the ratio of a reduced form coefficient (like the effect of HARP eligibility on mortgage default) to a first stage coefficient (like the effect of HARP eligibility on refinancing take-up). The reduced form coefficient, sometimes referred to as the "intent-to-treat," would show the overall impact of HARP *eligibility*, rather than *take-up*, on outcomes. This essentially shows the overall impact of HARP on a set of borrowers, combining their take-up rate with their treatment effect. This decomposition of a LATE into an intent-to-treat and a first stage seems particularly valuable in our heterogeneity analysis because of our finding that borrowers with stronger LATEs (those with low credit scores, high credit utilization rates, and high CLTVs) also tend to have weaker take-up rates, so it appears unclear which group is affected more by the policy overall.

Such a decomposition is not easily applicable in our non-standard environment, however. As described at length in Section III, our sample is not partitioned into "treatment" and "pre" periods, as borrowers could refinance at any time. Rather, we use an empirical model that takes advantage of the full time profile of refinancing rates in a more continuous fashion. The result is a large set of instruments (HARP eligibility and HARP eligibility interacted with quarter dummies; and in the heterogeneity regressions, those terms are then further interacted with HighCreditScore/CreditUtilization/CLTV); additionally, the heterogeneity regressions have multiple endogenous regressors (Refied and Refied × HighCreditScore/CreditUtilization/CLTV). As a result there is no single coefficient of interest that comes out of a first stage or reduced form regression, and so that powerful econometric intuition of a LATE being comprised of an intent-to-treat and a first stage breaks down.

However, we run auxiliary regressions to pursue that intuition despite this technical issue. Define

$$\text{Post}_t = \begin{cases} 1 & \text{if } t \geq 201206 \\ 0 & \text{otherwise,} \end{cases} \tag{A-2}$$

which is an indicator for whether the observation is before or after June 2012, which marks the beginning of HARP 2.0 and largely coincided with the beginning of the refinancing boom. Further define

$$\text{HighC}_{it} = \begin{cases} 1 & \text{if borrower } i \text{ has characteristic C in period } t \text{ above the sample median} \\ 0 & \text{otherwise,} \end{cases} \tag{A-3}$$

where characteristic C is either credit score, credit utilization rate, or CLTV. We then run the following regression for different outcomes of interest, Y_{it}:

$$\begin{aligned} \mathbb{E}[Y_{it}|t, \text{Eligible}_i, \text{X}_{it}] = {} & \nu_{11}\text{Eligible}_i + \nu_{12}\text{Post}_t + \nu_{13}\text{HighC}_{it} \\ & + \nu_{21}\text{Post}_t \times \text{HighC}_{it} + \nu_{22}\text{Eligible}_i \times \text{HighC}_{it} + \nu_{23}\text{Eligible}_i \times \text{Post}_t \\ & + \nu_{31}\text{Eligible}_i \times \text{Post}_t \times \text{HighC}_{it} \\ & + \text{X}_{it}\mu, \end{aligned} \tag{A-4}$$

where X_{it} is the standard set of controls we use throughout the paper. This sort of specification is often referred to as a "triple differences" regression, as it will produce a positive (negative) value of ν_{31} if the differential effect of being in the HARP 2.0 period on the eligible sub-group is higher (lower) for those for which HighC_{it} is 1. In other words, this shows which group was more responsive to HARP eligibility in the time period where HARP take-up was strong. The advantage of these regressions is that, by collapsing all the time periods into simply pre vs. post, the time interactions become easier to interpret. As an example, if Y is mortgage default and C is credit score, a positive ν_{31} means that the mortgage default propensity fell relatively less for eligible high-credit score borrowers from the pre-201206 period to the post-201206 period, compared to eligible low-credit score borrowers.

On the other hand, the cost of this ease of interpretation is that, by collapsing the timing component into the coarse Post_t variable, these are not first-stage and reduced-form coefficients corresponding directly to the 2SLS output from our main analysis. While they are intuitively evaluating the same things, they are fundamentally new regressions, and by collapsing the time dimension, we lose a lot of useful information, often leading to large standard errors and in some cases coefficients that are stronger or weaker than what one would expect from our core heterogeneity results (the 2SLS results shown in Tables 4 and 7). Despite this caveat, the results provide some useful insights.

Table A-9 shows ν_{31} for the different outcomes of interest. The first panel shows the effect on whether a borrower has refinanced, while the remainder of the table looks at our outcomes of interest from Sections IV and V. The first panel shows that financially constrained borrowers are *more* likely to be induced to refinance by HARP eligibility; marginally so based on credit score and credit utilization, substantially so based on CLTV. This result provides an interesting contrast to our finding in the take-up analysis in Section VI that struggling borrowers are *less* likely to refinance. To understand this contrast, consider CLTV: among the eligible group, high-CLTV borrowers refinance less than low-CLTV borrowers (as shown in Table 8); however, when comparing to the ineligible group, the high-CLTV borrowers refinance differentially more than do low-CLTV borrowers. This is because among HARP-ineligible borrowers with high CLTVs, essentially nobody refinances, while ineligible borrowers with relatively low CLTVs are often able to refinance without HARP. In other words, while eligible high-CLTV borrowers are less likely to refinance than their low-CLTV counterparts, they become *far* less likely to refinance when ineligible, so the triple interaction coefficient is actually positive for CLTV. So while it is true that there is in general a tension between who is more likely to refinance and who responds more after having done so, it appears that HARP—by targeting a constrained group (those with high LTVs)—in fact was effective at differentially reaching those with the largest treatment effects.

The bottom panel shows the triple interaction coefficients for other outcomes. By not utilizing much of the information contained in the timing and simply treating pre-201206 as a homogeneous period and post-201206 as another, we lose a lot of the power to precisely estimate effects in the reduced form. Nevertheless, the results are largely consistent with those from the 2SLS analysis. For instance, HARP-eligible borrowers with low credit scores, high credit utilization, and high LTVs disproportionately reduce their default incidence in the post-201206 period. Auto debt increases more for eligible borrowers with low credit scores or high credit utilization, while bank card debt gets paid down by these borrowers.

Bibliography

Karamon, Kadiri, Douglas McManus, and Jun Zhu. 2016. "Refinance and Mortgage Default: A Regression Discontinuity Analysis." *The Journal of Real Estate Finance and Ecnonomics*, 1–19.

State	Share of Sample	Cumulative Share	$\text{HPA}_{2009Q1}^{2010Q1}$
IL	7.2%	7.2%	−6.58%
CA	6.1%	13.2%	0.50%
TX	6.0%	19.2%	0.21%
WA	5.7%	24.9%	−7.04%
PA	4.2%	29.2%	−2.29%
FL	4.0%	33.2%	−8.36%
NJ	3.9%	37.1%	−5.71%
NC	3.8%	40.9%	−2.45%
NY	3.5%	44.4%	−3.41%
GA	3.3%	47.7%	−3.50%
MD	3.1%	50.8%	−7.83%
OR	3.0%	53.8%	−7.61%
AZ	2.9%	56.7%	−8.38%
MN	2.9%	59.6%	−2.97%
UT	2.7%	62.3%	−7.43%
WI	2.6%	64.9%	−2.49%
CO	2.5%	67.4%	1.42%
MO	2.5%	79.9%	−1.40%
VA	2.4%	72.2%	4.15%
MI	2.3%	74.5%	−5.40%
TN	2.0%	76.6%	−2.30%
OH	1.8%	78.4%	1.06%
SC	1.7%	80.1%	−5.32%
ID	1.6%	81.7%	−12.31%
MA	1.6%	83.3%	0.59%

Table A-1: Top 25 states in our sample, in descending order of observations. The final column shows home price appreciation (from CoreLogic's Home Price Indexes) in that state from the beginning of our window of originations (2009Q1) to the time when we determine whether loans have high LTVs and are included in our sample (2010Q1).

	A. Sample				B. Population			
	Mean	Median	10^{th} Pctl	90^{th} Pctl	Mean	Median	10^{th} Pctl	90^{th} Pctl
CLTV	81.6%	81.9%	67.4%	94.5%	70.9%	71.5%	35.8%	102.4%
Interest Rate	4.96%	4.88%	4.63%	5.38%	5.45%	5.50%	4.38%	6.5%
Mortgage Principal	$231k	$220k	$119k	$357k	$190k	$160k	$66k	$351k
Credit Score	768	780	708	808	751	769	672	806
Credit Utilization	23.7%	12.5%	0.4%	68.0%	31.4%	19.0%	0.6%	83.6%
30-year? (share)	91.5%				77.5%			
Junior lien? (share)	6.7%				23.8%			
Purchase mortgage? (share)	26.3%				33.6%			

Table A-2: Moments of the distributions of key observables in the CRISM dataset for borrowers who refinanced between March 2010 and February 2016. Panel A looks at our sample, while Panel B looks at a 1% random sample of all FRM borrowers in CRISM. Statistics are calculated as of 1 month prior to the refinance.

	GSE Purchase Month								
Origination Mo.	Jan 09	Feb 09	Mar 09	Apr 09	May 09	Jun 09	Jul 09	Aug 09 +	
Jan 09	867	16,107	15,173	528	175	86	30	1,389	34,355
	(3%)	(47%)	(44%)	(2%)	(1%)	(0%)	(0%)	(4%)	
Feb 09	0	1,053	27,663	3,369	676	243	51	1,771	34,826
	(0%)	(3%)	(79%)	(10%)	(2%)	(1%)	(0%)	(5%)	
Mar 09	0	0	6,152	15,413	7,905	1,731	136	1,234	32,571
	(0%)	(0%)	(19%)	(47%)	(24%)	(5%)	(0%)	(4%)	
Apr 09	0	0	0	1,399	17,600	11,589	725	1,378	32,691
	(0%)	(0%)	(0%)	(4%)	(54%)	(35%)	(2%)	(4%)	
May 09	0	0	0	0	1,050	23,403	2,071	1,052	27,576
	(0%)	(0%)	(0%)	(0%)	(4%)	(85%)	(8%)	(4%)	
Jun 09	0	0	0	0	0	7,514	30,535	19,339	57,388
	(0%)	(0%)	(0%)	(0%)	(0%)	(13%)	(53%)	(34%)	
	867	17,160	48,988	20,709	27,406	44,566	33,548	26,163	219,407

Table A-3: Month of GSE purchase, by cohort, in the CRISM sample. Parentheses show percentage of origination cohort c purchased in month t. The vertical line between May and June indicates the eligibility cutoff.

	Interest Rate (%)			Monthly Payment ($)			
	Before Refi	After Refi	Diff	Before Refi	After Refi	Diff	(in %)
Credit Score							
Above Median	4.94	3.90	-1.04	1,538	1,371	-168	(-10.9%)
Below Median	5.03	4.00	-1.03	1,539	1,358	-179	(-11.6%)
Credit Utilization							
Above Median	5.01	3.99	-1.01	1,539	1,364	-174	(-11.3%)
Below Median	4.95	3.90	-1.05	1,538	1,366	-172	(-11.2%)
CLTV							
Above Median	5.01	4.01	-1.01	1,547	1,371	-178	(-11.5%)
Below Median	4.94	3.88	-1.06	1,528	1,359	-168	(-11.0%)

Table A-4: Average values of payments and interest rates for refinancing borrowers, split by credit score, credit utilization, and CLTV. Credit score, credit utilization, and CLTV are measure 3 months prior to the refinance, while the Before and After values are measured 1 month before and after the refinance, respectively. Last column expresses the payment decrease in % of the monthly payment before the refinance.

	Net Change				Positive	Negative
Auto	9	10	10	10	30	-20
(Std. Err.)	(2)	(2)	(2)	(2)	(2)	(2)
HELOC	3	2	2	1	19	-18
	(2)	(2)	(2)	(2)	(2)	(2)
Bank Card	-28	-26	-26	-27	20	-46
	(1)	(1)	(1)	(2)	(2)	(2)
Student	-1	-1	-1	-1	0	-1
	(1)	(1)	(1)	(1)	(1)	(1)
Retail Consumer Debt	-1	0	0	0	13	-13
	(1)	(1)	(1)	(1)	(1)	(1)
Initial Bal.	✓	✓	✓	✓	✓	✓
Quarter FEs	✓	✓	✓	✓	✓	✓
ZIP-code FEs	✓	✓	✓	✓	✓	✓
Observables	✓	✓	✓	✓	✓	✓
Q-by-ZIP FEs		✓	✓	✓	✓	✓
Guar. Lag FEs			✓	✓	✓	✓
Cohort FEs				✓		
N (mill.)	14.5	14.5	14.5	14.5	14.5	14.5

Table A-5: Regression estimates of the relationship between refinancing and monthly debt accumulation by OLS. For borrower i in month t, the refinancing indicator is turned on if she has completed a refinance in some month $\tau \leq t$. Outcomes are the first difference in debt balances. "Net Change" is the simple difference, while the column labeled "Positive" censors negative changes to 0, and the column labeled "Negative" censors positive changes to 0. "Observables" include 10 equally-sized bins for each CLTV (lagged 3 months and at origination), credit score (lagged 3 months and at origination), credit utilization (lagged 3 months and at origination), initial mortgage rate, remaining principal balance, and initial debt balances. We also include indicators for mortgage "purpose" (e.g. purchase, cash-out refi, etc.). Standard errors (in parentheses) are clustered at the county level.

	Pre-201206	Post-201206	Pre/post ratio
Credit Score (≤ 675 omitted)			
676-725	1.63	1.60	1.02
	(0.08)	(0.05)	(0.04)
726-775	1.92	1.76	1.09
	(0.09)	(0.05)	(0.04)
776-800	2.10	1.89	1.11
	(0.10)	(0.06)	(0.04)
> 800	2.07	1.78	1.16
	(0.10)	(0.06)	(0.04)
Credit Utilization (1^{st} quartile omitted)			
2^{nd} quartile	1.06	1.04	1.02
	(0.02)	(0.02)	(0.02)
3^{rd} quartile	0.99	1.01	0.98
	(0.02)	(0.02)	(0.02)
4^{th} quartile	0.90	0.94	0.96
	(0.02)	(0.02)	(0.03)
Cred. Lim. = 0	0.85	0.86	0.99
	(0.02)	(0.03)	(0.05)
CLTV ($\leq 85\%$ omitted)			
85%-90%	0.80	0.91	0.89
	(0.01)	(0.02)	(0.02)
90%-95%	0.68	0.85	0.80
	(0.01)	(0.02)	(0.02)
$> 95\%$	0.57	0.81	0.71
	(0.02)	(0.02)	(0.02)

Table A-6: Time-varying effects of key observables on refinancing take-up for HARP-eligible borrowers. Tables shows results from a Cox proportional hazard model analogous to column (5) of Table 8, but allowing the effects of credit score, credit utilization, and CLTV to vary for the two subperiods $t < 201206$ and $t \geq 201206$. We use June 2012 to partition the sample period because it marks the beginning of HARP 2.0 and the approximate start of the refinancing boom. Estimates are reported in terms of the hazard ratio. The final column shows the ratio between the two hazard ratios for the subperiods. Standard errors (in parentheses) are clustered at the county level.

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Guarantee Lag (months)				Eligible for HARP (p.p.)	
Init. LTV ($< 60\%$ omitted)						
60-65%	-0.016	-0.018	-0.006	-0.011	0.334	0.184
	(0.004)	(0.004)	(0.006)	(0.005)	(0.092)	(0.099)
65-70%	-0.014	-0.024	-0.001	-0.014	0.111	0.052
	(0.005)	(0.004)	(0.006)	(0.005)	(0.101)	(0.108)
70-75%	0.008	-0.009	0.023	0.002	-0.151	-0.175
	(0.006)	(0.005)	(0.007)	(0.006)	(0.096)	(0.111)
75-80%	0.022	-0.005	0.047	0.014	-0.654	-0.648
	(0.006)	(0.005)	(0.008)	(0.007)	(0.094)	(0.098)
80-85%	-0.030	-0.064	-0.051	-0.090	1.000	1.102
	(0.008)	(0.008)	(0.009)	(0.009)	(0.133)	(0.150)
85-90%	0.017	-0.051	0.030	-0.036	-0.826	-0.378
	(0.008)	(0.009)	(0.010)	(0.011)	(0.148)	(0.160)
$\geq$ 90%	0.020	-0.050	0.014	-0.056	0.076	0.469
	(0.014)	(0.013)	(0.015)	(0.014)	(0.271)	(0.271)
Init. LTV $\times$ MaxLag						
60-65%			-0.007	-0.005		
			(0.002)	(0.002)		
65-70%			-0.009	-0.007		
			(0.002)	(0.002)		
70-75%			-0.010	-0.007		
			(0.002)	(0.002)		
75-80%			-0.018	-0.014		
			(0.002)	(0.002)		
80-85%			0.015	0.018		
			(0.002)	(0.002)		
85-90%			-0.010	-0.011		
			(0.004)	(0.004)		
$\geq$ 90%			0.007	0.007		
			(0.004)	(0.004)		
N (mill.)	1.5	1.5	1.5	1.5	1.5	1.5
R^2	0.015	0.087	0.015	0.087	0.675	0.688
Cohort FEs	✓	✓	✓	✓	✓	✓
Other controls		✓		✓		✓

Table A-7: Regressions to test whether HARP's eligibility requirement induced servicers to guarantee high-LTV mortgages more quickly. Columns (1)-(4) use guarantee lag in months, winsorized at 6, as the left-hand side variable (its sample average is 1.45 months). Columns (3) and (4) interact the LTV bins with MaxLag, the maximum number of months a loan can wait to be guaranteed while still maintaining HARP eligibility (e.g. for the April cohort, MaxLag $= 1$). Columns (5) and (6) use HARP eligibility—an indicator for whether the guarantee occurred before June 2009—as the left-hand side variable (55.95% of the loans are eligible). The sample includes all GSE FRMs that were originated between January and June 2009. "Other controls" include 10 equally-sized bins for credit score, credit utilization, interest rate and mortgage balance at origination, mortgage "purpose" (e.g. purchase, cash-out refi, etc.), and ZIP-code FEs. Standard errors (in parentheses) are clustered at the county level.

Time Since Refi:	$\leq$ 12 months	$>$ 12 months	Diff
	A. Debt Default		
Mortgage Default	0.45	-2.73	3.17
	(2.74)	(0.81)	(2.95)
Non-Mtg Default	-3.14	-3.36	0.22
	(3.57)	(1.35)	(3.69)
	B. Debt Accumulation		
Auto	45	16	29
	(24)	(9)	(28)
HELOC	-14	43	-57
	(20)	(9)	(25)
Bank Card	16	-10	26
	(14)	(6)	(15)
Student	-18	-1	-17
	(14)	(5)	(15)
Retail Consumer Debt	34	-5	40
	(12)	(5)	(14)

Table A-8: Regression estimates of the LATE of refinancing on monthly debt accumulation and the likelihood of default, allowing the effects to differ between the first 12 months since refinancing and after. Outcomes in Panel A are default indicators (effect reported in bp), and outcomes in Panel B are the first difference in debt balances (effect reported in $). Estimates result from instrumenting for the refinance indicators (one indicating whether borrower i refinanced in the past 12 months, the other indicating whether borrower i refinanced more than 12 months ago) with HARP eligibility interacted with a full set of quarter indicators. Controls include 10 equally-sized bins for each of: CLTV (lagged 3 months and at origination), credit score (lagged 3 months and at origination), credit utilization (lagged 3 months and at origination), initial mortgage rate, remaining principal balance, and initial debt balances. We also include indicators for mortgage "purpose" (e.g. purchase, cash-out refi, etc.) and fixed effects for quarter, ZIP code by quarter, and guarantee lag. Standard errors (in parentheses) are clustered at the county level.

	High Credit Score	High Credit Utilization	High CLTV
Refied (pp)	-0.58	1.11	10.34
	(0.31)	(0.29)	(0.46)
Mortgage Default (bp)	1.79	-0.90	-0.63
	(0.66)	(0.66)	(0.75)
Non-Mtg Default (bp)	0.77	-1.08	-0.63
	(0.79)	(0.77)	(0.87)
Auto ($)	-15	4	2
	(6)	(6)	(6)
HELOC ($)	-5	-5	4
	(5)	(5)	(5)
Bank Card ($)	9	-5	-6
	(4)	(4)	(4)
Student ($)	2	0	-1
	(3)	(3)	(3)
Retail Consumer Debt ($)	0	-5	-1
	(3)	(3)	(3)

Table A-9: Heterogeneous impact of HARP eligibility on balance sheet variables. Each cell reports the coefficient on a "triple difference" from separate regressions (ν_{31} from Equation A-4); the 3 variables being interacted are HARP eligibility, an indicator for being after June 2012, and an indicator for being above the median value of the characteristic (e.g. credit score). We control for the following variables: quarter FEs, ZIP-code FEs, observables, and guarantee lag FEs. Additional controls include 10 equally-sized bins for each of: CLTV (lagged 3 months and at origination), credit score (lagged 3 months and at origination), credit utilization (lagged 3 months and at origination), initial mortgage rate, remaining principal balance, and initial debt balances. We also include indicators for mortgage "purpose" (e.g. purchase, cash-out refi, etc.). Standard errors (in parentheses) are clustered at the county level.

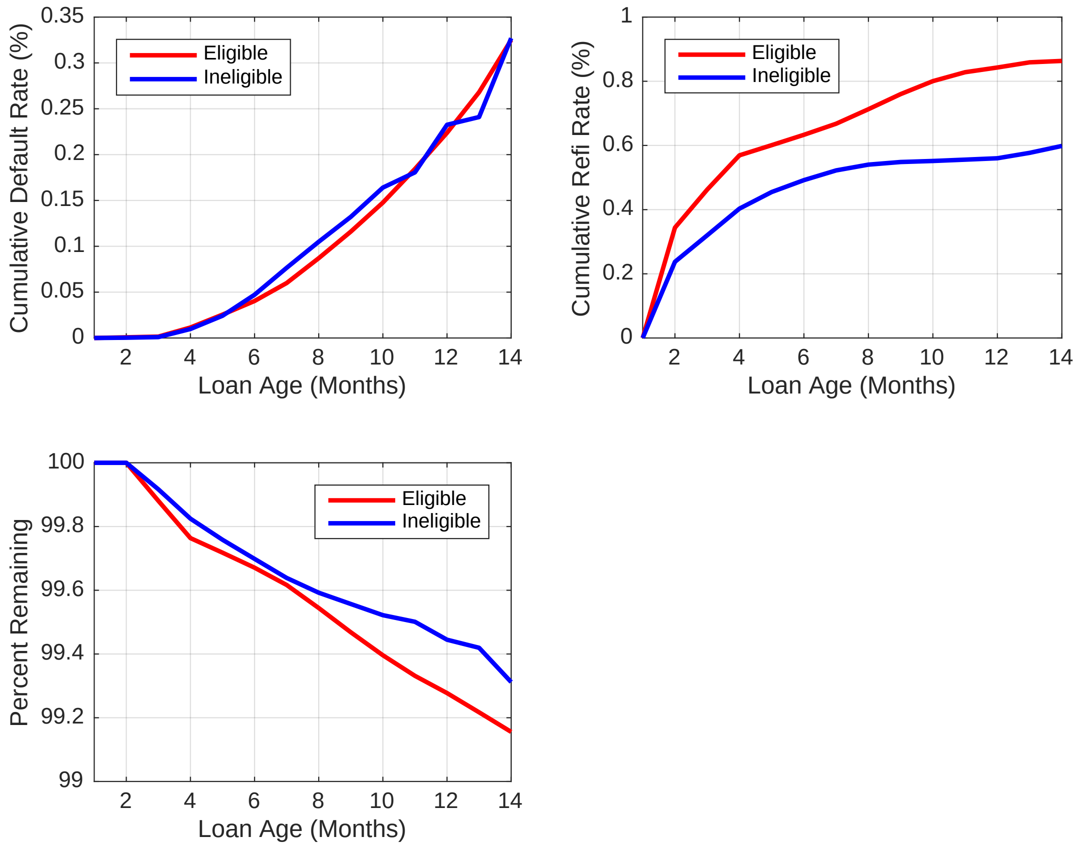

Figure A-1: Cumulative refinance and default rates for FRM GSE loans originated between January and June 2009 with initial LTVs above 70%. This is to check whether attrition from the data prior to sample selection (which occurs in March 2010) could affect the results. Only about 1% of the loans exited from the sample within 14 months of originating their mortgage (we choose 14 because that is the number of months between the first originations—January 2009—and the March 2010 sample selection date), so any induced selection is likely to be minimal. Furthermore, HARP eligibility was not very predictive of this form of attrition—default rates across the eligible and ineligible groups were nearly identical. While the eligible group was more likely to refinance in this period (and thus drop from the sample), this is entirely due to time effects (a regression that controls for month effects finds no significant predictive value of HARP eligibility on the probability of refinancing in this period). Thus, our decision to wait ≈ 1 year before selecting our sample seems unlikely to have caused any meaningful selection problems due to attrition.

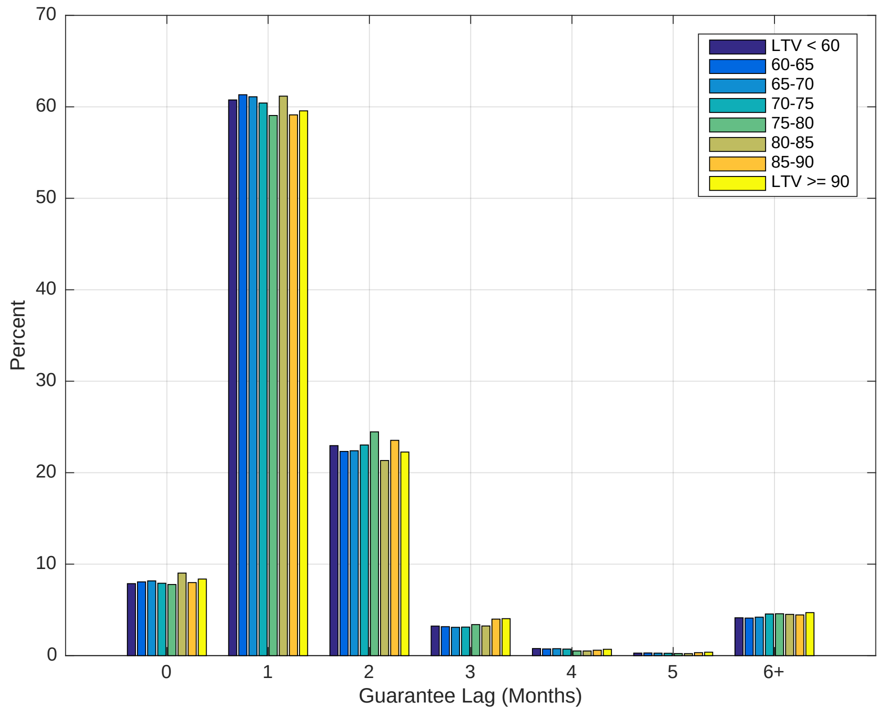

Figure A-2: Histogram of guarantee lag in CRISM. Guarantee lag is defined as the number of months between the mortgage being originated and being purchased by a GSE. We take the set of FRMs originated between January and June 2009 and split them between the bins shown.