No Job, No Money, No Refi: Frictions to Refinancing in a Recession

ANTHONY A. DEFUSCO and JOHN MONDRAGON*

ABSTRACT

We study how employment documentation requirements and out-of-pocket closing costs constrain mortgage refinancing. These frictions, which bind most severely during recessions, may significantly inhibit monetary policy pass-through. To study their effects on refinancing, we exploit a Federal Housing Administration policy change that excluded unemployed borrowers from refinancing and increased others' out-of-pocket costs substantially. These changes dramatically reduced refinancing rates, particularly among the likely unemployed and those facing new out-of-pocket costs. Our results imply that unemployed and liquidity-constrained borrowers have a high latent demand for refinancing. Cyclical variation in these factors may therefore affect both the aggregate and distributional consequences of monetary policy.

FRICTIONS THAT PREVENT HOUSEHOLDS FROM refinancing their debt during times of economic distress can significantly inhibit policy efforts aimed at curtailing the costs of recessions. This was particularly true during the Great Recession, when frictions in the U.S. mortgage market held back a broad array of policies directed at providing debt relief and economic stimulus to households. These frictions ranged from widespread levels of negative equity, which limited the ability of many households to benefit from accommodative monetary policy

(Beraja et al. ([2019](#))), to competitive barriers in the mortgage market, which suppressed the take-up of federal mortgage modification and refinancing programs (Agarwal et al. ([2017a](#), [2017b](#))).[1]

In this paper, we study how two previously overlooked but important frictions may contribute to a lack of refinancing during recessions. To refinance a mortgage, borrowers typically need to both document that they are employed and pay upfront, out-of-pocket closing costs. Although always present, these constraints may be especially binding during recessions, when unemployment is high, income risk is elevated, and cash-on-hand is low. They are also likely to have significant distributional implications. The households who are most affected—the unemployed and the liquidity-constrained—are precisely those who would benefit most from refinancing into a lower interest rate. Yet, despite their potential importance, little is known about the extent to which these constraints actually bind in practice.

To quantify the effect of these frictions on refinancing in a recession, we exploit a sharp policy change introduced by the Federal Housing Administration (FHA) during the height of the Great Recession. Prior to late 2009, borrowers with an FHA mortgage were typically *not* constrained by out-of-pocket closing costs or employment documentation requirements. Instead, these borrowers were allowed to roll all closing costs into their new mortgage and were not required to provide any income or employment documentation so long as they refinanced into a new FHA mortgage through the FHA's Streamline Refinance (SLR) Program. However, in response to the general deterioration in the mortgage market, the FHA eliminated both of these provisions from the SLR Program in late 2009. Under the revised program, borrowers with negative equity had to pay for any upfront refinancing fees out-of-pocket, and unemployed borrowers were prohibited from refinancing altogether.[2] Changes in refinancing rates among FHA borrowers following the policy change should therefore be informative about both the latent demand for refinancing among the unemployed and the extent to which upfront costs inhibit refinancing during a recession.

To identify the combined effect of these changes to the SLR Program, we begin with a simple event study that exploits the sharp timing of the policy change. Graphical analysis reveals that refinancing rates among FHA borrowers experienced an exceptionally large and discrete fall in precisely the month that the policy changes took effect. This drop in refinancing persists even after controlling flexibly for time trends and a large set of borrower- and loan-level observables. Our baseline estimates imply that the policy reduced the monthly probability that an FHA borrower refinanced her mortgage by a full percentage

---

[1] See Piskorski and Seru ([2018](#)) for a comprehensive review of the literature studying how mortgage market frictions interacted with household debt relief and restructuring attempts during the Great Recession.

[2] Crucially, the FHA did not change its policy on home equity and refinancing. FHA borrowers with negative equity were still permitted to refinance through the SLR Program as long as they could pay for the closing costs and prove that they were employed.

point, which amounts to a decline of roughly 80% relative to the pre-shock average.

Although these results strongly suggest that the policy change had a negative effect on refinancing, the event study approach cannot completely rule out the possibility that the drop in FHA refinancing was driven by concurrent macroeconomic shocks. To address this concern, we estimate difference-in-differences specifications that use the unaffected conventional (non-FHA) market as a control group. This approach is motivated by a similar graphical analysis of refinancing in the conventional market, which does not reveal any discrete changes around the time of the policy change. Including conventional borrowers as a control group allows us to fully and nonparametrically control for aggregate trends in refinancing rates and yields results that are similar to the event study analysis. Across a range of specifications, we estimate that the policy led to a reduction in the monthly FHA refinancing rate of roughly 0.7 percentage points, which is more than 50% of the baseline rate. Finally, to further support our approach, we estimate flexible specifications that allow the effect on FHA refinancing to vary by month and find that the differential fall in refinancing among FHA borrowers coincides exactly with the implementation of the policy change. Taken together, these results provide strong evidence that the policy changes had a large negative effect on FHA refinancing rates.

Having documented the combined effect of the new employment documentation and closing cost requirements on refinancing rates among FHA borrowers, we then turn to examining the effects of these two provisions separately. We identify these effects using a triple-differences approach that compares how the post-policy fall in FHA refinancing relative to conventional refinancing varies across groups of borrowers who are more or less likely to be affected by each of the two constraints.

To isolate the effect of the employment documentation requirement, we use variation in the likelihood that a borrower is unemployed based on changes in county-level unemployment rates. Specifically, we take the difference in refinancing rates between borrowers in high- and low-unemployment counties, before and after the policy, and across FHA and conventional borrowers. Our estimates show that the post-policy fall in refinancing among FHA borrowers was substantially larger in high- relative to low-unemployment counties, but that there was no differential change in refinancing behavior among conventional borrowers across these two groups of counties. Our baseline estimate suggests that raising the county-level unemployment rate by one percentage point reduces the monthly probability that an FHA borrower refinances by about 0.05 percentage points following the policy change. These estimates are robust to our full set of controls, and the timing of the effect is consistent with the change in FHA policy. Importantly, these results are not caused by the simultaneously imposed new requirement that negative-equity FHA borrowers pay for closing costs out-of-pocket. We estimate all of our specifications only in the subsample of borrowers with positive equity and, in the preferred specification, allow for the level of home equity to have a fully flexible and time-varying independent effect on relative refinancing rates of FHA borrowers.

Taken together, our estimates imply that unemployed borrowers have a high demand for refinancing that is constrained by the standard mortgage underwriting process. Back-of-the-envelope calculations suggest that allowing unemployed borrowers to refinance at the rates implied by our estimates would have generated approximately 2.4 million additional refinances in 2009. Together, these refinances would have saved unemployed borrowers roughly $7.2 billion that year. This represents almost 10% of the entire amount paid out to unemployed people through the unemployment insurance system in 2009. These implied payment savings are also of roughly the same order of magnitude as what was achieved by the Home Affordable Modification Program (HAMP) and Home Affordable Refinancing Program (HARP), both of which notably excluded unemployed borrowers from participating (Agarwal et al. (2017a, 2017b)). Our results on employment documentation requirements therefore highlight a quantitatively important and previously underemphasized barrier to refinancing during a recession.

In the last part of the paper, we study the effects of the change in how upfront costs are financed. Following the policy change, borrowers with low or even negative levels of equity could still refinance their loans through the SLR Program. However, if there was insufficient equity to roll the upfront costs into the new loan, borrowers would have to pay these costs out-of-pocket. To the extent that paying the closing costs upfront was either unaffordable or suboptimal, this change could lead to a meaningful reduction in FHA refinancing even among employed borrowers. To measure this effect, we identify borrowers who likely have insufficient equity based on their initial loan-to-value (LTV) ratios and changes in local house prices. We then estimate a similar triple-difference model, taking the difference between borrowers with high and low equity levels, before and after the policy, and across the FHA and conventional markets. To avoid conflating this effect with the employment documentation requirements, we estimate these regressions only in the subsample of counties with below-median increases in unemployment and, in some specifications, also allow for county-level changes in unemployment to have a flexible and time-varying effect on FHA refinancing. We find that the inability of low-equity borrowers to roll the closing costs into the loan had very large negative effects. Our preferred estimate suggests that this friction reduced monthly refinancing rates among FHA borrowers by at least 0.5 percentage points. Comprehensive data on closing costs for FHA streamlines are not generally available, but estimates of the average range from $2,000 to $3,000 depending on the state (Woodward (2008)). Survey evidence suggests that many households would have difficulty accessing this amount of cash even in an emergency, which may explain why we find such large effects (Lusardi, Schneider, and Tufano (2011)).

Forcing households to pay for closing costs out-of-pocket could also reduce refinancing even among those with sufficient liquid assets. In particular, increases in up-front costs can push the refinancing option out of the money for households who discount cash flows at a rate higher than that at which they are able to borrow. To separate this mechanism from the liquidity effect, we construct estimates of the optimality of the refinancing option for each

borrower and in each month based on the model in Agarwal, Driscoll, and Laibson (2013). We then re-estimate our effects on the sample of borrowers for whom the refinancing option is still optimal even after having to pay for closing costs. The results for this subset of borrowers are similar to those in our full sample, which suggests that lack of liquidity is the dominant driver of the drop in refinancing following the policy change.

Related Literature

Our paper is closely related to a growing body of work studying the relationship between household financial frictions and monetary policy. Bernanke and Gertler (1995) were among the first to emphasize the household balance sheet channel as a way of understanding how monetary policy affects the real economy. Caplin, Freeman, and Tracy (1997) and more recently Beraja et al. (2019) emphasize the role that home equity plays in amplifying or mediating interest rate changes through the mortgage refinancing channel. Similarly, Di Maggio, Kermani, and Palmer (2016) show that the effects of quantitative easing during the Great Recession operated in large part through mortgage refinancing. We build on this work by quantifying the effects of both employment documentation and closing costs on refinancing, each of which likely become more important in typical recessions. The frictions we document, because they limit the pass-through of interest rate changes to households with fixed-interest mortgages, help quantify how much less economic stimulus is being affected through both conventional and unconventional monetary policy. Our work is also related to the mechanisms highlighted by Greenwald (2018), who emphasizes how payment-to-income restrictions affect the ability of interest rate changes to affect credit growth. Agarwal et al. (2018) and Scharfstein and Sunderam (2016) examine how frictions arising from market structure, lack of competition, and bank incentives affect the pass-through of monetary policy to households through several channels including refinancing.[3] Similarly, Bond et al. (2017) show that the presence of second mortgages can inhibit the ability of borrowers to refinance into to a lower interest rate on their first mortgage when junior creditors refuse to be re-subordinated. Finally, Auclert (2017) and Coibion et al. (2017) argue that monetary policy can have heterogeneous effects on households due to variation in wealth and income. We document that variation in income and liquidity can lead to large differences in mortgage refinancing, highlighting another channel through which differences across households interact with the transmission of monetary policy.

Our work is also related to the vast literature studying households' mortgage refinancing decisions. Much of this literature documents that households do not refinance at the frequency that would be expected in a fully rational frictionless benchmark model (Andersen et al. (2019), Agarwal, Rosen, and Yao

---

[3] Agarwal et al. (2017b) similarly show that lack of competition in the refinancing market inhibited the effectiveness of HARP by both lowering take-up rates and reducing interest rate savings on refinances completed through the program.

(2015), Campbell (2006), Chang and Yavas (2009), Deng and Quigley (2012), Deng, Quigley, and Order (2000), Green and LaCour-Little (1999), Johnson, Meier, and Toubia (2015), Keys, Pope, and Pope (2016)). We contribute to this literature by quantifying the effects of two important frictions that can help explain part of the observed sluggish refinancing behavior. Our results on the effects of closing costs provide empirical support for the results in, among many others, Agarwal, Driscoll, and Laibson (2013), Dunn and Spatt (2005), and Stanton (1995), who demonstrate the important role of upfront costs on refinancing behavior. We also rely on the characterization of refinancing optimality from Agarwal, Driscoll, and Laibson (2013) to test for the relative importance of liquidity and upfront costs. Our emphasis on the role of income and employment documentation relates to Archer, Ling, and McGill (1996), who emphasize the role of payment-to-income constraints in reducing refinancing as well as Campbell and Dietrich (1983), Dickinson and Heuson (1994), and Pavlov (2001). Our emphasis on both refinancing and the FHA SLR Program relates our paper to Ehrlich and Perry (2015), who also study the SLR Program, but focus on quasi-experimental variation in premiums to show how reduced payments affect mortgage performance.

Finally, the SLR Program presents an interesting complement to mortgage modification programs, which have been emphasized in the wake of the 2009 financial crisis (Adelino, Gerardi, and Willen (2013), Agarwal et al. (2011, 2017a, 2017b), Eberly and Krishnamurthy (2014), Ganong and Noel (2017), Haughwout, Okah, and Tracy (2016), Mayer et al. (2014)). Our work suggests that streamlined refinancing may be a useful alternative to modification programs, which potentially suffer from competitive and moral hazard frictions that restrict uptake. The benefits of the streamline program in reducing payments quickly, irrespective of property valuations and incomes, potentially apply to the Government-Sponsored Enterprise (GSE) market as well since explicit guarantees against credit risk are also made by Fannie Mae and Freddie Mac when those loans are securitized. As such, our results are directly informative about the large-scale refinancing programs proposed by Lucas, Moore, and Remy (2011) and Boyce et al. (2012), both of which advocate for a relaxation of refinancing standards in the conventional market along the lines of the FHA SLR Program.

The rest of the paper proceeds as follows. Section I describes the institutional background for our analysis and the details of the policy shock we examine. Section II describes the data and sample we use. Section III provides estimates of the overall effect of the policy on FHA refinancing rates. Section IV presents results on the two mechanisms, unemployment and upfront costs. Section V concludes.

I. Institutional Background

The FHA was founded in 1934 to help stabilize the mortgage market during the Great Depression. Now regulated by the Department of Housing and Urban Development (HUD), one of the primary functions of the FHA is to provide

access to homeownership for households unlikely to satisfy conventional mortgage underwriting requirements. To accomplish this goal, the FHA provides insurance to originators of FHA loans that fully protects against any principal losses associated with borrower default. To pay for the default insurance, the FHA charges borrowers a mortgage insurance premium (MIP). One part of the MIP is collected upfront (UFMIP) and often rolled into the mortgage, while a second part is added to the interest rate and collected monthly throughout the life of the loan. As a result, FHA mortgages typically have higher interest rates than comparable conventional loans but generally allow for higher LTVs and flexible income and credit requirements. In addition to purchase mortgages, the FHA also offers refinances, reverse mortgages, and cash-out refinances, along with both fixed and adjustable rates. During the period we study, the FHA was involved in financing nearly one out of every five new mortgages in the United States.[4]

A. The FHA SLR Program

When interest rates began to fall rapidly in 1981, the FHA faced new and substantial demand to refinance a large stock of high-interest loans. In response to this demand, the FHA created the SLR Program in October 1982. In its announcement of the program, the FHA outlined that "certain types of applications to refinance existing [FHA] mortgages need not contain a standard credit report and the regular verifications of deposit and employment."[5] Later, the FHA relaxed these standards even further by dropping the requirement that borrowers obtain an appraisal for the property being financed. From the FHA's perspective, the justification for a refinancing program that does not require either employment documentation or positive equity is relatively straightforward. If a borrower has an FHA mortgage, then the FHA has already insured that mortgage against default. By allowing the borrower to refinance and reduce their payment, the FHA has weakly reduced the probability of default.

The SLR Program quickly became a standard and popular option for FHA borrowers looking to refinance. For example, during the refinancing boom from 2001 to 2003, just over 80% of all FHA refinances went through the SLR Program. In 2009, which is when the policy change that we study occurred, FHA SLRs represented roughly 6% of all refinances in the United States.[6]

---

[4] See table 3 of https://www.hud.gov/sites/documents/FHA_SF_MARKETSHARE_2016Q2.PDF, which indicates that FHA loans constituted 21.1% and 17.5% of all new mortgages issued in 2009 and 2010, respectively.

[5] See https://www.hud.gov/sites/documents/82-23ML.TXT.

[6] These figures are based on the authors' calculations (available on request) using data from HUD's FHA Single-Family Outlook Reports and the Federal Reserve Board's report on the 2009 HMDA data. The majority (60%) of the FHA refinances that did not go through the SLR Program were cash-out refinances, which do not qualify for SLR. Although the FHA does not report statistics for the remaining 40% of these loans, they likely consisted primarily of conventional-to-FHA refinances, which are also excluded from the SLR Program.

To use the SLR Program, borrowers need to be refinancing an existing FHA mortgage and they cannot receive more than $500 cash back, which is typically used to cover small discrepancies in prepayments or estimated escrow costs. SLRs must also lower the borrower's payment unless there is a substantial reduction in the term of the mortgage. Prior to the policy change we study, lenders participating in the program were not required to document any cash that might be needed for closing nor were there any limits on the borrower's combined loan-to-value (CLTV) ratio so long as all subordinate financing retained its junior lien position.

Within the SLR Program there are two primary types of refinance: noncredit qualifying with appraisal, and noncredit qualifying without appraisal.[7] The most important distinction between these two options involves restrictions on the size of the new mortgage. In the first column of Table I, we provide a detailed layout of the maximum loan amounts that were permitted under both types of streamline before the policy change that we study. Without an appraisal, a borrower could finance all closing costs as well as any discount points so long as the new mortgage amount did not exceed the original principal balance of the mortgage being paid off. This was true regardless of whether the borrower's current house value placed them in positive or negative equity. If the borrower did get an appraisal, then the new mortgage was allowed to exceed the original principal balance up to a maximum of 97.75% of the newly appraised value, which could also be used to pay for any closing costs associated with the loan.[8] Neither type of streamline required lenders to check income or employment.

B. Major Changes to the SLR Program

On September 18, 2009, HUD announced sweeping changes to the streamline program, taking effect 60 days later.[9] We focus on the two major changes to the program that fundamentally altered access for unemployed borrowers and for borrowers with low levels of equity. First, lenders had to begin certifying that the borrower was employed with an income before extending an SLR.[10] Although no strict income limits were imposed, this new requirement explicitly excluded any borrower that was unemployed or had income that was

---

[7] In addition to the two noncredit qualifying options, there is also a third category of SLR referred to as credit-qualifying. Unlike the noncredit qualifying options, credit-qualifying SLRs require documentation of income, a minimum FICO score of 620, and underwriting to income ratios. This refinance represents a small share of FHA business and is primarily used when deleting a borrower from the mortgage or if the new refinance has substantially larger payments (e.g., due to a term reduction).

[8] These maximum LTVs were imposed starting in early 2009. See https://www.nclc.org/images/pdf/foreclosure_mortgage/loan_mod/hope/lmp_hope_refinance_transactions.pdf.

[9] For the full text of the announcement, see http://portal.hud.gov/hudportal/documents/huddoc?id=09-32ml.doc.

[10] In practice, lenders now had to fill out and certify the income sections of the Uniform Residential Loan Application (URLA).

Table I

Maximum Loan Amounts for FHA Streamline Refinances before and after the Policy Change

This table presents the FHA-mandated loan calculations for streamline refinances with and without an appraisal, before and after the policy changes we study. Net UFMIP refers to the upfront mortgage insurance payment required by the FHA minus any refund due to the borrower for the UFMIP on the original loan. The interest payoff reflects the fact that interest on FHA loans is guaranteed at the beginning of each new month. This means that if a loan is refinanced in the middle of a month, the borrower is still responsible for the remaining days' interest on the old loan. Closing costs include origination fees and other underwriting costs in addition to title, attorney, and recording fees. Net prepaid expenses will reflect any prepayment of interest, insurance, or taxes minus any amount still due for those costs. Discount points are upfront payments the borrower may make to the lender in exchange for a reduced rate, which the FHA mandated be "reasonable." Items with an "a" superscript were eliminated from the loan calculations by the policy change.

Pre-Policy Change (1)	Post-Policy Change (2)
Streamline without Appraisal	
Minimum of:	Minimum of:
Original principal balance	Original principal balance
+ Net UFMIP	+ Net UFMIP
or	or
Unpaid Principal balance	Unpaid principal balance
+ Net UFMIP	+ Net UFMIP
+ Interest payoff	+ Interest payoff
+ Closing costs[a]	
+ Net prepaid expenses[a]	
+ Discount points[a]	
Streamline with Appraisal	
Minimum of:	Minimum of:
97.75% of appraised value	97.75% of appraised value
+ Net UFMIP	+ Net UFMIP
or	or
Unpaid principal balance	Unpaid principal balance
+ Net UFMIP	+ Net UFMIP
+ Interest payoff	+ Interest payoff
+ Closing costs	+ Closing costs
+ Net prepaid expenses	+ Net prepaid expenses
+ Discount points[a]	

difficult to document from refinancing their mortgage, irrespective of the borrower's equity or credit score.[11]

The second change we examine prevented borrowers with low levels of equity from rolling closing costs into the new mortgage. This resulted from a change in

---

[11] Unemployed borrowers were likely the largest class of borrowers shut out from the market by the new policy. However, those whose income is difficult to verify and who may also have been affected include self-employed borrowers without a W-2, retirees drawing income from a retirement account, and borrowers with a significant fraction of income coming from tips or gratuities.

the treatment of refinances without appraisals. Prior to the policy change, the loan amount for SLRs without an appraisal was allowed to increase dollar for dollar with any increase in closing costs up to the original principal balance of the loan being paid off. This meant that a borrower would be able to finance her closing costs even if she had negative equity since the maximum loan amount was determined based on the amount of the original loan and not the value of the house. The change in policy eliminated this option entirely. As shown in column (2) of Table I, the maximum loan amount for streamlines without appraisals was reduced such that no closing costs could be rolled into the new mortgage.[12] In contrast, streamlines *with* an appraisal were still allowed to roll closing costs into the mortgage up to a maximum of 97.75% of the newly appraised value. Therefore, if a borrower wanted to finance closing costs using the new loan, she would have to order an appraisal and that appraisal would need to indicate that the house was worth more than the remaining unpaid balance. That is, she would need to have positive equity.

To summarize, the policy change completely eliminated the ability for unemployed FHA borrowers to refinance through the SLR Program and increased the out-of-pocket costs of refinancing substantially for borrowers with insufficient equity. After these changes were announced, lenders in the FHA market noted that the employment and appraisal changes would likely be very important. One lender stated that these changes were a "landscape shifter," and summarized the effects as "No job? No money? No FHA loan."[13]

It is somewhat difficult to document the exact rationale for why the FHA implemented these policy changes. In the press release describing the changes, the FHA indicated that they were intended to "bring documentation standards in line with other FHA loan origination guidelines" and would "[ensure] the borrower's capacity to repay the new mortgage."[14] However, the actual economic benefits of the policy changes to the FHA are unclear. While preventing underwater borrowers from rolling closing costs into the new mortgage may potentially reduce the FHA's liability, precluding these borrowers as well as those who are unemployed from refinancing into lower interest rates would likely increase that liability. The reason for this is simple: lower interest rates will at least weakly reduce default risk, and any borrower using the SLR Program would have already had that risk insured by the FHA. The SLR changes were announced two months after the swearing in of a new FHA commissioner, who was forced to address the serious deterioration in the insurance fund's capital reserves. It is possible that these changes were instituted as part of this broader effort. In particular, the FHA began treating SLR originations as new loans for the purposes of its actuarial review around this time and may therefore have needed to harmonize underwriting guidelines between the SLR

---

[12] In addition to this change, HUD also began requiring that any funds needed for closing be directly verified by the lender.

[13] This quote was originally available at https://themortgagereports.com/3231/fha-streamline-refi-changes, but an archived version is housed at https://web.archive.org/web/20120604014910/https://themortgagereports.com/3231/fha-streamline-refi-changes.

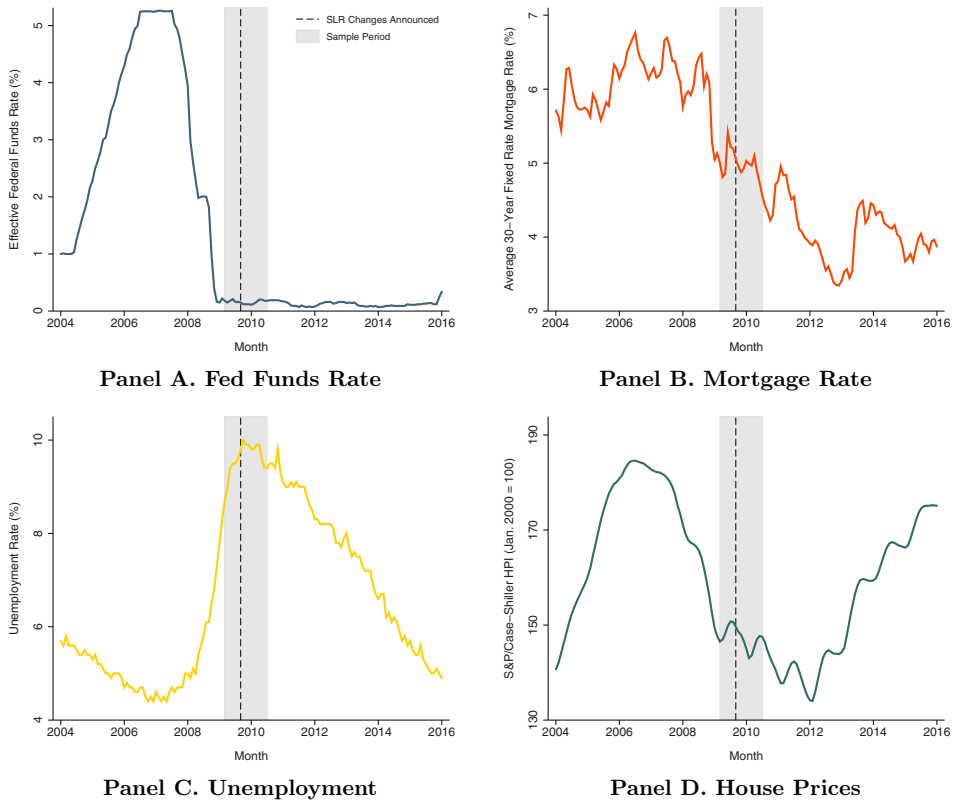[14] See https://archives.hud.gov/news/2009/pr09-177.cfm.

Figure 1. Aggregate economic conditions at the time of the FHA policy changes. This figure plots monthly time-series trends for several key measures of aggregate economic conditions around the time of the FHA SLR policy changes. The vertically dashed line in each panel marks the month of the policy announcement (September 2009). The shaded gray area marks our primary sample period, which covers the window of time spanning six months prior to the policy announcement and six months after its implementation (March 2009 to July 2010). Panel A plots the effective federal funds rate and is sourced from the Board of Governors' H.15 Release on Selected Interest Rates. Panel B plots the average offered interest rate on new 30-year fixed-rate mortgages from Freddie Mac's Primary Mortgage Market Survey. Panel C plots the Civilian Unemployment Rate from the Bureau of Labor Statistics' monthly Employment Situation Summary. Panel D plots the S&P/Case-Shiller U.S. National House Price Index and is normalized to a value of 100 in January of 2000. (Color figure can be viewed at wileyonlinelibrary.com)

Program and new originations as a result.[15] To the best of our knowledge, there is no clear documentation explicitly outlining why these steps were taken at this particular point in time, apart from statements like those described above.

Regardless of the exact reasons, these changes were likely to have important effects. In Figure 1, we plot trends in several key measures of aggregate

---

[15] See Aragon et al. (2010) for a discussion of how SLRs were treated in the 2009 actuarial review.

economic conditions at the time. These trends indicate that the policy changes were indeed likely to have had a significant effect on refinancing in the FHA market. Panels A and B, which plot the federal funds rate and the going rate on new 30-year fixed-rate mortgages (FRMs), indicate that the incentive to refinance and the associated benefits were likely to have been quite large. The SLR changes were announced shortly after the Federal Reserve began the first round of quantitative easing and aggressively lowered its key policy rate to near 0%, which was accompanied by a sharp one-percentage-point drop and continuing downward trend in mortgage rates. Yet as Panels C and D indicate, many borrowers in the FHA market would be shut out from these benefits by the new policy changes. The national unemployment rate had just reached its peak of nearly 10% and house prices were continuing their precipitous fall from their peak in 2006. All of these unemployed borrowers would be excluded from the SLR Program and those with limited cash-on-hand living in parts of the country experiencing the largest house price declines would very likely have found it prohibitively costly to refinance through the program due to the new need to pay for closing costs out-of-pocket.

C. Other Changes to SLR

In addition to the major changes outlined above, there were several other small changes to the SLR Program that were announced at the same time but are unlikely to affect our results. These changes were directed primarily at reducing the extent of refinance "churning," a practice by which mortgage originators would aggressively market refinances to existing borrowers to capture new origination fees despite generating no real benefit for the borrower.

To mitigate this practice, HUD began imposing requirements limiting the set of outstanding FHA loans that were eligible for a streamline based on both the age of the loan and the potential benefits to the borrower. In particular, following the policy change, only loans that were at least six months old and for which the refinance would lead to a "net tangible benefit" for the borrower were eligible for the SLR Program. The net tangible benefit requirement varied somewhat based on both the type of loan that was being refinanced (fixed-rate versus variable) and the type of loan that would be replacing it. However, for the vast majority of SLR transactions, which are fixed-to-fixed refinances, the net tangible benefit standard only required that the new monthly payment be at least 5% lower than the payment on the current loan.[16] Estimates from various sources suggest that almost all FHA refinances would have satisfied this requirement (Ehrlich and Perry (2015); Agarwal et al. (2017b); Lambie-Hanson

---

[16] The 5% threshold used to determine whether a loan meets the net tangible benefit standard would include any increase in payments due to rolled-in closing costs. If refinancing from an adjustable-rate mortgage (ARM) to an FRM, then the new rate could not be more than 200 bps greater than the current rate on one-year FHA ARMs. Refinances from ARMs to hybrids required that the payment not increase by more than 20%. Finally, FRMs refinancing into ARMs required a rate that was at least 200 bps less than the rate on the current loan.

and Reid (2017)). However, to limit the effect of these changes on our analysis, we restrict our sample to include only FRMs that had been outstanding for at least six months as of the date of the policy announcement.

In addition to the changes targeting refinance churning, HUD also started requiring that borrowers have satisfactory payment histories to qualify for an SLR. In particular, if the loan was less than 12 months old at the time of application, then the borrower was required to have made all payments on time to participate in SLR. If the loan was older than 12 months, then all payments in the last three months must have been on time and no more than one payment in the last year may have been 30 days late. In our analysis, we also restrict our sample to include only loans that met these requirements.

II. Data and Sample

A. Data Sources

We rely on the Loan-Level Market Analytics (LLMA) data from Corelogic for our primary analysis. The data are collected from large mortgage servicers and cover about 60% of first liens originated over the period we examine in both the agency and nonagency markets. We rely on three distinct files from the data set. The first is a static file containing information recorded at the time of origination, including borrower characteristics (e.g., FICO, debt-to-income [DTI], occupancy status), loan characteristics (e.g., loan amount, interest rate, LTV), property characteristics (e.g., ZIP code, property type), and an indicator for whether or not the loan is FHA insured. The second file is dynamic and records monthly performance information over the life of the loan. The performance data allow us to observe when a loan is delinquent or paid off, but do not distinguish between payoffs resulting from sales versus refinances. To address this issue we rely on the Supplemental Loan Analytics file, which uses merges (conducted by Corelogic) of the originations and performance data to public deeds records. Using these merged data, Corelogic is able to determine whether a loan that has been paid off is a refinance or a sale. Our sample is restricted to loans for which we are able to determine the payoff reason.

To construct estimates of a borrower's current equity, we use the reported LTV at origination along with the house price appreciation implied by Zillow's county-level house price indices. In each month, we impute the current value of the borrower's home and subtract the observed remaining balance to determine the borrower's level of equity. This estimate will suffer from error for at least two reasons. First, if the borrower's home has experienced idiosyncratic (with respect to the county) appreciation or depreciation, this will not be reflected in the county-level price changes. Second, if the borrower has taken out a junior lien against the house after origination of the first loan, this additional debt will not be reflected in the performance data. This means that we will tend to overstate the level of equity. These issues should, if anything, attenuate our estimates. Finally, we use estimates of county-level annual

unemployment rates available from the American Community Survey (ACS) to measure differences in the likelihood that a borrower is unemployed.

B. Sample Selection and Description

We restrict our attention to a set of relatively standard mortgages that were active at some point during the window of time spanning six months prior to the policy announcement and six months after the implementation date (March 2009 to July 2010). To construct our analysis sample, we begin with an initial 20% random sample of all first-lien, fixed-rate, 30-year, single-family, owner-occupied mortgages active during this period for which CoreLogic reports a nonmissing interest rate, LTV, and FICO score at origination.[17] From this sample, we then exclude any loans that would be prevented from participating in the updated SLR Program for mechanical reasons related to payment history. This includes loans that are less than six months old and loans with recent payment histories that do not meet the requirements described in Section I or with insufficient information to determine whether these requirements would be met. Finally, we keep only loans in counties for which both Zillow house price indices and ACS unemployment rates are available. The full details of our sample selection procedure, including the number of observations dropped at each step, are described in Section I of the Internet Appendix.[18]

Our final analysis sample contains approximately 1.3 million loans that are observed for a total of 15.6 million loan-months. Table II reports summary statistics for this sample. In Panel A, the unit of analysis is the individual loan and summary statistics are measured in the month of origination. In Panel B the unit of analysis is the loan-month and summary statistics are measured across all months for which a loan appears in our sample. The first column pools across all loans, whereas columns (2) and (3) report statistics separately for conventional and FHA loans, respectively. As expected, FHA loans tend to be smaller and have lower FICO scores and higher LTVs than conventional loans. However, FHA and conventional borrowers have roughly similar DTI ratios and (unconditional) interest rates. The average monthly probability of refinancing in the full sample is 0.84 percentage points and is slightly lower for FHA borrowers, who constitute approximately 17% of the sample.

III. The Combined Effect of SLR Policy Changes

The changes to the SLR Program announced by the FHA in 2009 may have led to a reduction in refinancing among FHA borrowers for two primary

---

[17] We limit the sample to owner-occupied loans secured by single-family homes as the FHA program has distinct procedures for condos and investor or second homes that also changed over this period.

[18] The Internet Appendix is available in the online version of this article on *The Journal of Finance* website.

Table II
Summary Statistics

This table reports descriptive statistics for loans in the final analysis sample. All table entries represent sample means or, in parentheses, standard deviations. Summary statistics are presented pooling across all loan types (column (1)) as well as separately for conventional (column (2)) and FHA loans (column 3). In Panel A, the unit of analysis is the individual loan and summary statistics are measured in the month of origination. In Panel B, the unit of analysis is the loan-month and summary statistics are measured across all months for which a loan appears in our sample.

	All Loans (1)	Conventional (2)	FHA (3)
Panel A: Loan-Level Characteristics at Origination			
Loan Amount ($1000's)	216.91	226.01	171.17
	(139.75)	(145.86)	(90.80)
FICO Score	722.82	732.21	675.56
	(61.48)	(55.90)	(66.29)
Loan-to-Value	74.29	70.15	95.10
	(18.66)	(17.41)	(7.48)
Back-End Debt-to-Income	37.52	37.14	39.68
	(15.49)	(16.08)	(11.37)
Interest Rate	5.90	5.91	5.82
	(0.81)	(0.83)	(0.71)
Panel B: Loan-Month Characteristics			
Current Balance ($1000's)	195.57	201.46	155.45
	(134.71)	(139.52)	(85.50)
Estimated Equity ($1000's)	96.05	109.42	4.98
	(174.82)	(183.06)	(31.03)
Current Loan-to-Value	75.41	72.20	97.30
	(25.97)	(25.43)	(17.74)
Loan Age (Years)	3.79	3.99	2.44
	(2.33)	(2.26)	(2.32)
Percent Refinanced	0.84	0.86	0.70
	—	—	—
Number of Loans	1,309,393	1,092,163	217,230
Number of Loan-Months	15,645,645	13,643,184	2,002,461

reasons. First, the new requirement that lenders document income explicitly excluded unemployed borrowers from refinancing through the program. Second, the reduction in the maximum loan amount for streamlines without an appraisal meant that underwater FHA borrowers who wanted to refinance would now need to pay for any upfront closing costs out-of-pocket. In this section, we estimate the combined effect of these two policy changes on FHA refinancing rates. In Section IV, we will examine the importance of each of these two channels separately.

A. Empirical Strategy

A.1. Event Study

To estimate the overall average effect of the policy changes, we use two alternative empirical strategies that leverage different aspects of our data. The first is a simple event study that compares refinancing behavior before and after the policy change for FHA borrowers while controlling for aggregate trends in refinancing as well as a broad set of loan-level and time-varying observables that are typically considered to be important inputs into a household's decision to refinance. This approach exploits the discrete timing of the policy change as the primary source of identification. The key identifying assumption is that the probability an FHA loan refinances would have evolved smoothly over time in the absence of the policy change. We will provide direct graphical evidence in support of this assumption below by showing that FHA refinancing rates tended to evolve smoothly in all months during our sample period except the month that the policy went into effect, when there was a large and discrete drop.

To implement this approach, we estimate versions of the following monthly, loan-level panel regression:

$$Refinance_{it} = \alpha + X'_{it}\gamma + \beta_0 \cdot Post_t + \delta_0(t - \tau) + \delta_1(t - \tau) \cdot Post_t + \epsilon_{it}, \quad (1)$$

where $Refinance_{it}$ is an indicator variable denoting whether or not loan i refinances in month t and X_{it} is a vector of loan-level and possibly time-varying observables. The indicator variable $Post_{it}$ takes the value one if month t falls on or after January 2010, the first month after the policy change.[19] The coefficient of interest is β_0, which measures the change in the average rate of refinancing among FHA borrowers after the policy has taken effect. To ensure that this coefficient will reflect only the discontinuous change in refinancing induced by the policy, we also include linear time trends that we allow to differ before and after the date of the policy change (τ = January 2010). These trends control for general changes in the likelihood of refinancing over time. If income documentation requirements or the need to pay for closing costs out-of-pocket are important barriers to refinancing, then we should expect to find $\beta_0 < 0$. Standard errors are clustered by core-based statistical area (CBSA) in all specifications.

One potential issue with this specification is that it does not allow for any anticipation effects. The policy changes were announced in late September 2009, which was a full two months before they took effect. There is some anecdotal evidence that lenders were aware of this and took efforts to notify potential clients of the need to refinance ahead of the changes.[20] To the extent that this

---

[19] Although December 2009 was the first full month when SLR applications had to abide by the new rules, due to the amount of time it takes for loans to close, many of the loans with applications prior to the deadline would likely not be recorded as refinanced until 30 or more days later. Therefore, we will always treat January 2010 as the first "post-policy" month.

[20] See, for example, https://themortgagereports.com/3231/fha-streamline-refi-changes.

behavior was widespread and borrowers decided to refinance early, this could lead us to overestimate the effect of the policy since it would generate a higher refinancing rate in the pre-period. To account for this, we also estimate specifications that include an additional indicator variable marking periods of time subsequent to the announcement of the policy. In particular, we estimate the following modified version of equation (1):

$$Refinance_{it} = \alpha + X'_{it}\gamma + \beta_0 \cdot Post_t + \delta_0(t - \tau) + \delta_1(t - \tau) \cdot Post_t$$
$$+ \beta_1 \cdot Post_t^{News} + \delta_2(t - \tau^{News}) \cdot Post_t^{News} + \epsilon_{it}, \qquad (2)$$

where $Post_t^{News}$ is an indicator variable that takes the value of 1 if month t falls on or after September 2009. As in the baseline specification, we allow the linear time trend to differ for months following the policy announcement ($\tau^{News} =$ September 2009). A small estimate of β_1 would suggest that there was relatively little anticipation of the policy.

A.2. Difference-in-Differences

One disadvantage of the event study approach is that it cannot account for sharp changes in outcomes that would have occurred even in the absence of the policy change. This is an especially important concern in our context because refinancing probabilities often exhibit large changes when interest rates begin to rise or fall. To address this issue, we also provide estimates based on a difference-in-differences strategy leveraging the fact that the policy changes had no effect on refinancing options for borrowers with conventional (non-FHA) mortgages. If movements in household expectations about interest rates or other macroeconomic factors caused a large change in refinancing at the same time as the policy change, this effect should manifest itself similarly among both conventional and FHA borrowers. Therefore, by netting out any changes in refinancing among conventional borrowers, we will be able to isolate the effect of the policy change alone.

The baseline specification that we use to implement this approach is a standard differences-in-differences regression estimated at a monthly frequency using the full sample of both conventional and FHA loans. Specifically, we estimate regressions of the following form:

$$Refinance_{it} = \alpha + \delta_t + X'_{it}\gamma + \beta_0 \cdot FHA_i + \beta_1 \cdot FHA_i \times Post_t + \epsilon_{it}, \qquad (3)$$

where δ_t is a vector of fixed effects for the month of observation and FHA_i is an indicator for whether or not loan i is FHA insured. The coefficient of interest is β_1, which measures the differential change in refinancing among FHA borrowers relative to conventional borrowers following the implementation of the SLR policy changes. This difference is conditional on a broad set of loan and borrower characteristics as well as time and geographic-specific factors.

The standard identifying assumption in this framework is that trends in FHA and conventional refinancing would have evolved in parallel in the absence of the policy change. In our context, the interpretation of this assumption requires some care. The nature of the policy change that we study was to make underwriting standards in the FHA market more similar to those in the conventional market. Prior to the policy change, FHA borrowers had easier access to refinancing than conventional borrowers. In particular, during the pre-period unemployed and underwater conventional borrowers would have typically been shut out of the market, whereas FHA borrowers would have still been able to refinance through the SLR Program. Because employment, house prices, and interest rates were all falling, this may have led to a decline in refinancing among conventional borrowers relative to FHA borrowers. This would violate the parallel trends assumption and lead us to *underestimate* any relative decline in FHA refinancing subsequent to the policy change. To account for this possibility, our set of control variables will always include a linear time trend for FHA borrowers. As in the event study analysis, this trend will be allowed to vary freely before and after the policy change. Below, we will provide graphical evidence showing that, conditional on these trends and the other controls that we include, refinancing rates in the two market segments evolved in parallel prior to the policy change.

B. Results

B.1. Graphical Evidence

As motivation for our empirical strategy, we begin by presenting simple graphical evidence indicating that the refinancing rates of FHA borrowers experienced a discontinuous and dramatic decline in exactly the month that the SLR policy changes went into effect. In Figure 2, we plot the raw unconditional probability that a loan refinanced during each month leading up to and after the policy changes. These refinancing rates are plotted separately for FHA loans (Panel A) and conventional loans (Panel B). The vertically dashed gray line in January 2010 marks the first post-policy month. In this figure and throughout the paper, we multiply all refinancing rates by 100, so that a value of 1 implies a 1% probability of refinancing in a given month.

Panel A of the figure shows that FHA refinancing rates fluctuated between roughly 0.6% and 1.5% prior to the policy, but then dropped sharply in January 2010 to 0.25%. For visual reference, the orange dashed lines plot the fitted values from a regression of the monthly refinancing probabilities on a linear time trend fit separately on either side of the policy change. These trends indicate that the refinancing rate among FHA borrowers fell by roughly one percentage point in precisely the month that the new restrictions to the SLR Program went into effect and remained low for the remainder of the sample period. The large and discontinuous nature of this drop provides strong evidence in support of our event study approach. In Panel B, we plot the analogous figure for
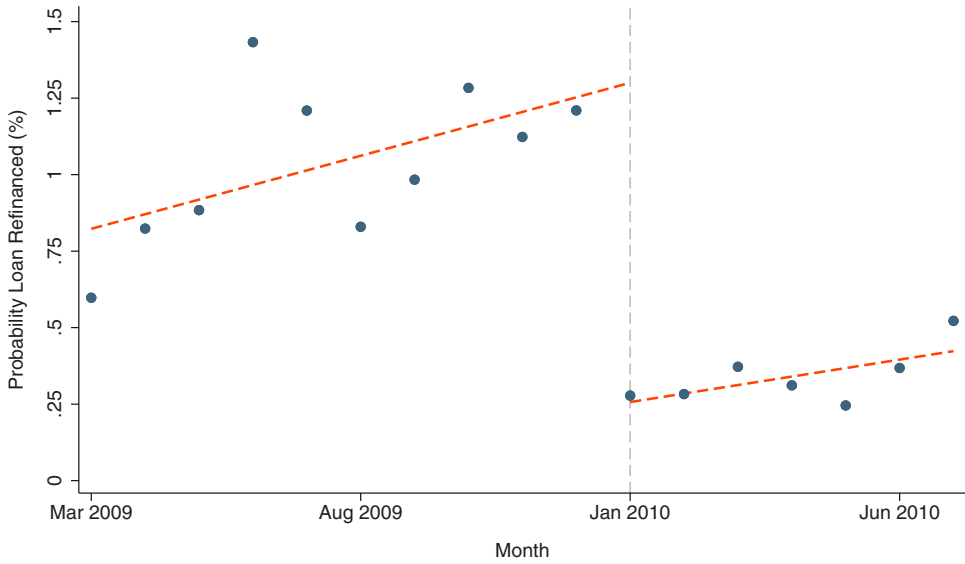
Panel A. FHA Loans
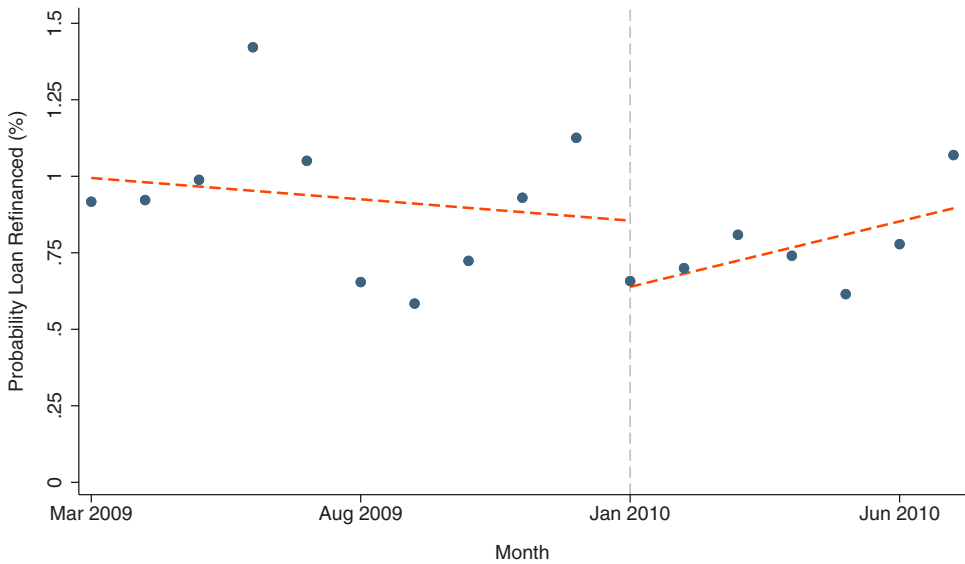
Panel B. Conventional Loans

Figure 2. FHA and conventional refinancing trends. This figure plots monthly unconditional refinancing rates between March 2009 and July 2010. Each dot represents the percentage of outstanding mortgages of a given type that refinanced in the indicated month. Refinancing rates are calculated separately for FHA (Panel A) and conventional loans (Panel B). The vertically dashed gray line in January 2010 marks the first month that the SLR policy changes went into effect. The dashed orange lines are the predicted values from a regression of the plotted refinancing rates on a linear time trend fit separately on either side of the cutoff date. (Color figure can be viewed at wileyonlinelibrary.com)

Table III

The Effect of the Policy Changes on FHA Refinancing: Event Study

This table reports event study estimates of the effect of the change in FHA policies on the monthly probability that an FHA loan refinances. Each column reports the estimated coefficients from a separate regression where the dependent variable is an indicator for whether or not a loan refinances in the month of observation. The outcome is multiplied by 100, so that all coefficients can be interpreted as percentage-point changes. In all specifications, coefficients are reported for the *Post* dummy denoting whether the month of observation is after the implementation of the policy changes (January 2010). The specifications in columns (3) and (4) also include an indicator for whether the month of observation was after September 2009, which was the month that the policy changes were announced (*Post News*). All specifications include fixed effects for the CBSA of the property as well as linear time trends that are allowed to differ on either side of the policy implementation date. In columns (3) and (4), an additional linear time trend is included for the period of time subsequent to the announcement date. Columns (2) and (4) include fixed effects for the current loan age (one-year bins), interest rate (one-percentage-point bins), and the borrower's estimated home equity ($10,000 bins), as well as the full pairwise interaction between the borrower's LTV (10-point bins) and FICO score (50-point bins) at origination. Standard errors are reported in parentheses and are clustered at the CBSA level. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Post	-1.041^{***}	-0.960^{***}	-1.026^{***}	-1.013^{***}
	(0.077)	(0.071)	(0.073)	(0.073)
Post News			-0.157^{***}	-0.112^{***}
			(0.039)	(0.040)
Time Trends	X	X	X	X
CBSA FEs	X	X	X	X
Loan Age FEs		X		X
Interest Rate FEs		X		X
LTV × FICO FEs		X		X
Equity FEs		X		X
Number of Observations	2,002,461	2,002,461	2,002,461	2,002,461

conventional loans. Although there is a slight difference in pre-trends between the two groups of loans, both appear to evolve roughly linearly prior to the policy change and there is no evidence of a drop in refinancing among conventional borrowers. Because we will always allow for separate linear trends between FHA and conventional loans, these results also lend support for the difference-in-differences strategy.

B.2. Event Study Results

Table III presents our main results from the event study analysis. The first two columns report estimates from the baseline specification given by equation (1). In the first column, we include only CBSA fixed effects and the linear trends. The coefficient on the *Post* dummy indicates that the change in policy reduced the monthly probability that an FHA loan refinanced by 1.04 percentage points. This estimate lines up closely with the raw averages reported in

Figure 2 and is large relative to the pre-period refinancing rate of roughly 1.25%. In the second column, we control nonparametrically for a host of loan and borrower characteristics that may also be important determinants of the likelihood of refinancing. To control for time-varying drivers of the demand for refinancing, we include fixed effects for the current loan age (one-year bins), interest rate (one-percentage-point bins), and estimated home equity ($10,000 bins) of the borrower. To control for differences in borrower characteristics at origination, we further include a full set of 50-point FICO score bins, 10-point LTV bins, and the pairwise interaction between the two. Including these controls has no meaningful effect on the result. The estimate reported in column (2) remains statistically significant at the 1% level and implies that the policy changes led to a reduction in FHA refinancing rates of 0.96 percentage points.

In columns (3) and (4), we report analogous estimates from the modified event study specification given by equation (2). This specification allows for the possibility that borrowers may have tried to front-run the policy changes by refinancing early in response to the news that was released several months before changes actually took effect. We find limited evidence of this type of anticipation effect. The coefficient on the $Post^{News}$ dummy is negative, and small in absolute value in both specifications. Moreover, including this coefficient and allowing for a separate linear time trend during the period between the announcement and implementation of the policy changes has essentially no effect on the magnitude of the main coefficient reported in the top row. Taken together, the estimates reported in this table suggest that the new constraints introduced by the SLR policy changes led to a reduction in refinancing among existing FHA borrowers of roughly one percentage point per month.

B.3. Difference-in-Differences Results

The event study results are largely confirmed by our difference-in-differences analysis, which compares not only how refinancing behavior changes following the implementation of the policy, but also whether the change in behavior is differential across FHA and conventional borrowers. In the first column of Table IV, we report estimates from a baseline version of the difference-in-differences specification given by equation (3). In this baseline regression, we control only for the month of observation, the CBSA of the property, and a linear time trend for FHA borrowers that is allowed to vary before and after the policy change. The coefficient of interest is reported in the second row and implies that the changes to the SLR Program reduced FHA refinancing rates by 0.8 percentage points, similar to the one-percentage-point reduction from the event study analysis. This effect is also large enough to more than offset the gap in refinancing rates that existed between FHA and conventional borrowers just prior to the policy change as indicated by coefficient estimate on the *FHA* dummy reported in the first row.

FHA and conventional borrowers differ along a broad set of observables. Because of this, one concern might be that differences in these observables would lead to large differences in refinancing rates that could confound our

Table IV

The Effect of the Policy Changes on FHA Refinancing: Difference-in-Differences

This table reports difference-in-differences estimates of the effect of the change in FHA policies on the monthly probability that an FHA loan refinances. Each column reports estimates from a separate regression where the dependent variable is an indicator for whether or not a loan refinances in the month of observation. The outcome is multiplied by 100, so that all coefficients can be interpreted as percentage-point changes. Coefficients are reported for the *FHA* "treatment" dummy as well as its interaction with an indicator for whether the month of observation was after the implementation of the policy changes (*Post*), which occurred in January 2010. All specifications include fixed effects for the month of observation and the CBSA of the property as well as linear time trends for FHA borrowers that are allowed to vary freely before and after the policy change. Column (2) adds fixed effects for the current loan age (one-year bins), interest rate (one-percentage-point bins), and the borrower's estimated home equity ($10,000 bins), as well as the full pairwise interaction between the borrower's LTV (10-point bins) and FICO score (50-point bins) at origination. Column (3) further interacts all of the additional fixed effects contained in column (2) with the *Post* dummy. Column (4) adds an additional set of interactions between these fixed effects and the *FHA* dummy. Standard errors are reported in parentheses and are clustered at the CBSA level. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
FHA	0.511***	0.827***	0.813***	2.041***
	(0.062)	(0.062)	(0.055)	(0.315)
FHA × Post	−0.804***	−0.727***	−0.709***	−0.708***
	(0.065)	(0.062)	(0.049)	(0.045)
Month FEs	X	X	X	X
CBSA FEs	X	X	X	X
FHA Time Trends	X	X	X	X
Loan Age FEs		X	X	X
Interest Rate FEs		X	X	X
LTV × FICO FEs		X	X	X
Equity FEs		X	X	X
Controls × Post			X	X
Controls × FHA				X
Number of Observations	15,645,645	15,645,645	15,645,645	15,645,645

estimates. In column (2), we begin to address this issue by controlling flexibly for all of the same characteristics included in our event study analysis (loan age, interest rate, current equity, LTV, and FICO). When we include these controls, the resulting estimate falls slightly but is statistically indistinguishable from the baseline effect reported in the first column. In column (3), we further interact all of the controls added in column (2) with the *Post* dummy. This allows for each borrower or loan characteristic to have a separate and time-varying effect on the likelihood of refinancing. Allowing for this additional flexibility does not affect the size of the implied drop in FHA refinancing caused by the policy change. Finally, in column (4), we further interact all of the borrower and loan-level controls with the *FHA* indicator. This allows for the possibility that FICO scores, for example, are differentially informative about

refinancing behavior for FHA borrowers relative to conventional borrowers. Allowing for these observables to vary with the type of loan gives an almost identical estimate. Across all of the specifications, we find robust evidence that borrowers with FHA loans are much less likely to refinance after the policy change relative to conventional borrowers. The size of this gap is large and indicates that the change in SLR policies led to a reduction in FHA refinancing of roughly 0.7 to 0.8 percentage points.

Finally, to give a sense of the dynamics of this effect, we estimate a more flexible version of the difference-in-differences specification that allows for the effect to vary by month. Specifically, we estimate a regression of the following form:

$$Refinance_{it} = \alpha + \delta_t + X'_{it}\gamma + \sum_\tau [\beta_\tau \cdot FHA_i \times \mathbb{1}_{t=\tau}] + \epsilon_{it}, \tag{4}$$

where $\mathbb{1}_{t=\tau}$ is an indicator variable taking the value of 1 if month t is equal to τ (e.g., December 2009). The β_τ coefficients from this regression provide a nonparametric measure of the differential trend in refinancing rates among FHA borrowers relative to conventional borrowers. We normalize the coefficient for December 2009 to zero, so that all estimates can be interpreted as the difference in refinancing rates between FHA and conventional borrowers in a given month relative to the corresponding difference in the month just prior to the policy changes. We include all of the same controls as in column (4) of Table IV, but instead of interacting these controls with just a single *Post* dummy, we allow for a full set of interactions with each of the month fixed effects.[21] If these observables are able to effectively control for any differences in pre-trends, then we should expect to find $\beta_\tau = 0$ in all months prior to December 2009.

In Figure 3, we plot these coefficients along with their 95% confidence intervals. The figure shows that, conditional on the controls, trends in refinancing rates between FHA and conventional borrowers evolved in rough parallel up until the month of the policy change.[22] However, starting in immediately the month of the policy change, there is a discrete drop in refinancing among FHA borrowers. The magnitude of this drop is roughly 0.7 percentage points, which is economically large and similar to the estimates from Table IV. Compared to the 1.25% rate of refinancing in the month prior to the policy change, this result suggests that the new restrictions to the SLR Program reduced FHA refinancing rates by just over 50% of the baseline.

B.4. The Cost of Not Refinancing

Taken together, the results above imply that imposing employment documentation requirements and requiring borrowers to pay for closing costs

---

[21] In Section II of the Internet Appendix, we also report results from specifications that are analogous to those from columns (1) to (3) of Table IV.

[22] Although there is some evidence of a differential pre-trend in the earlier months, this effect dissipates quickly, leading the trends to be roughly identical in the six months leading up to the policy change.
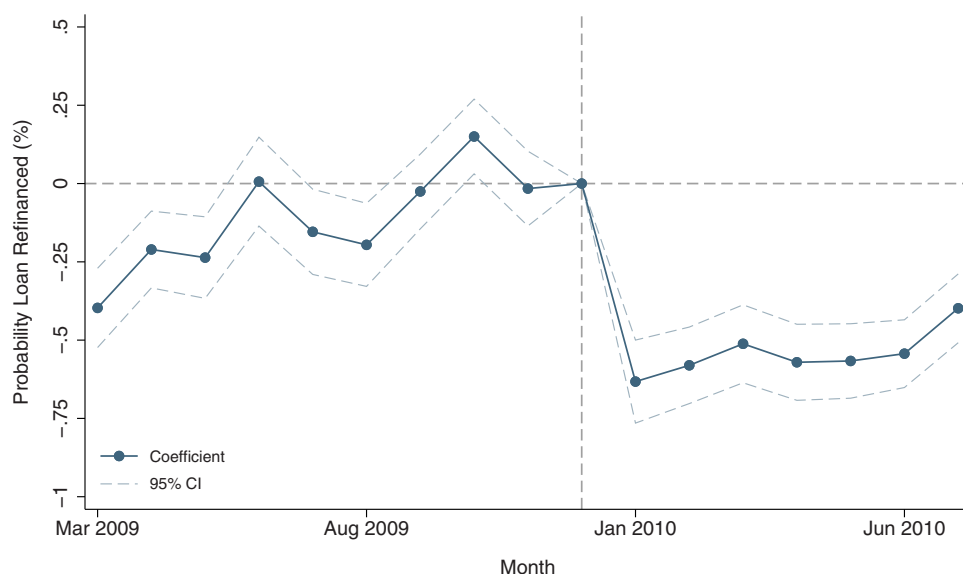
Figure 3. Flexible difference-in-differences estimates of the effect of the FHA policy changes on refinancing. This figure reports estimates of the effect of the change in FHA policies on FHA refinancing derived from a flexible difference-in-differences specification that allows the effect to vary freely by month of observation. Estimates were constructed by regressing an indicator for whether or not a loan refinances in a given month on a dummy variable denoting whether the loan was FHA insured and the interaction of that FHA dummy with a series of dummy variables indicating the month of observation. The coefficient for December 2009 is normalized to zero, so that all estimates can be interpreted as the change in the monthly probability of refinancing relative to the month prior to when the policy changes went into effect, which is marked by the vertically dashed gray line. The regression also included fixed effects for the CBSA of the property, the current loan age (one-year bins), interest rate (one-percentage-point bins), and the borrower's estimated home equity ($10,000 bins), as well as the full pairwise interaction between the borrower's LTV (10-point bins) and FICO score (50-point bins) at origination. With the exception of the CBSA fixed effects, all of these controls were also separately interacted with the FHA dummy and with the dummies for the month of observation. The 95% confidence intervals are based on standard errors that are clustered at the CBSA level. (Color figure can be viewed at wileyonlinelibrary.com)

out-of-pocket significantly constrains mortgage refinancing. The economic consequences of this reduced refinancing depend in part on how beneficial refinancing would have been for the individual borrowers who are prevented from doing so. If all of the borrowers excluded from the market by these constraints would have only marginally benefited from being able to refinance, then both the aggregate costs and the distributional consequences of these constraints are likely to be small. However, if a large fraction of borrowers prevented from refinancing would have meaningfully benefited from access to a lower interest rate, then the consequences of these constraints are potentially substantial. In this section, we sort borrowers according to their potential benefits from refinancing and document that the fall in refinancing among FHA borrowers

was driven in large part by borrowers who would have greatly benefited from being able to do so.

One way to measure the potential benefits of refinancing for an individual borrower would be to simply compare the interest rate they are currently paying to an estimate of the rate they would pay if they were to refinance their mortgage today. However, this comparison fails to take into account the potentially significant option value of waiting to refinance later. Among other things, this option value depends critically on expectations regarding future interest rates and the remaining life of the loan. To account for this, we follow Keys, Pope, and Pope (2016) and rely on the model of refinancing behavior provided by Agarwal, Driscoll, and Laibson (2013). This model takes standard loan and borrower characteristics as inputs and produces a threshold for the differential between a borrower's current rate and the rate on a new loan at which it would be optimal to refinance. If the gap between the prevailing rate and the borrower's current rate exceeds this threshold, it is optimal for the borrower to refinance in the sense that doing so would reduce the expected net present value of her obligations to the lender. We calculate these thresholds in each month and for each borrower in our sample using the same baseline calibration of the model used by Agarwal, Driscoll, and Laibson (2013) and Keys, Pope, and Pope (2016), which takes a conservative view on how many households should refinance.[23] For simplicity, we ignore the role of closing costs in the calibration and return to this in Section IV when discussing mechanisms.

One of the more critical inputs into this model is the "prevailing" rate to which a borrower could refinance if she chose to. We construct monthly estimates of this potential rate for each borrower in our sample using the set of observed refinances in each month among borrowers with similar characteristics. Specifically, we categorize observed refinances into FICO, LTV, state, month, and FHA cells and calculate the median observed interest rate on new loans originated within each of these cells.[24] We then assume that the medians within each cell give the potential rate for a borrower with the same observables.[25] Although clearly a simplification, this procedure helps to address the substantial dispersion in observed interest rates across borrowers and incorporates important correlations between borrower characteristics and potential rates. With these potential rates in hand, we can construct a measure of how beneficial it would be for a borrower to refinance at a given point in time by

---

[23] The specific calibration we use assumes that the real discount rate used by households is 5% (annual), the marginal tax rate is 28%, the relocation rate is 10%, the rate of events with full deductibility of expenses is 20%, the standard deviation of mortgage rates is 0.0109, the inflation rate is 3%, and we assume all borrowers are refinancing into 30-year FRMs. Actuarial data about the streamline program suggest that almost all SLRs are into fixed-rate loans.

[24] We bin FICO scores into standard categorizations used by lenders: ≥ 800, (800,740], (740,670], (670, 580], and ≤ 580. Similarly, for LTV we group loans into the following bins: > 90, [90, 80), [80, 70), [70, 60), and ≤ 60.

[25] If there is no observed rate for a given cell (no refinances were originated with those characteristics in that month and state), we assume that there is no potential rate for that kind of borrower and omit them from our analysis.
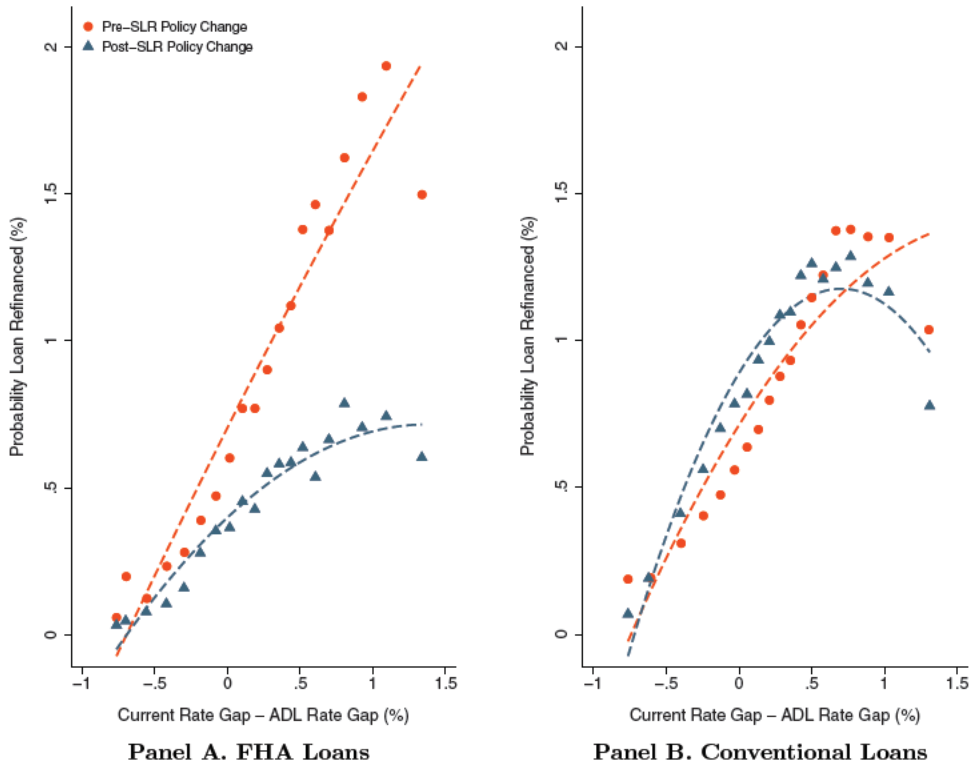
Figure 4. The relationship between refinancing and the benefits to refinancing. This figure plots the relationship between refinancing and the estimated benefits to refinancing for FHA and conventional borrowers before and after the SLR policy changes. The benefits to refinancing are proxied using the difference between the borrower's current rate gap and the rate gap at which it would be optimal for them to refinance according to the model in Agarwal, Driscoll, and Laibson (2013) (the "ADL rate gap"). Each dot in the figure is the average refinancing rate among borrowers in a given five-percentile bin of the distribution of differences between current rate gaps and ADL rate gaps. The relationship is plotted separately for loan-month observations prior to the SLR policy change (orange circles) and afterward (blue triangles). The dashed lines are the predictions from a quadratic fit of the refinancing rate to the difference in rate gaps across bins in a given time period. Panel A plots these relationships for FHA loans, whereas Panel B plots the same relationships among conventional loans. See Section III for details on how the ADL rate gaps and potential interest rates were calculated. (Color figure can be viewed at wileyonlinelibrary.com)

comparing the observed gap between their current and potential interest rate with the rate gap at which it would be optimal for the borrower to refinance according to Agarwal, Driscoll, and Laibson (2013) (the "ADL rate gap"). The larger the difference between the borrower's current rate gap and the ADL rate gap, the more beneficial refinancing is for that borrower.

In Figure 4, we sort borrowers according to this measure and plot their refinancing rates before and after the policy change. Each dot in the figure is the average refinancing rate among borrowers in a given five-percentile bin of the

distribution of differences between current rate gaps and ADL rate gaps. The relationship is plotted separately for loan-month observations prior to the SLR policy change (orange circles) and afterward (blue triangles). The dashed lines are the predictions from a quadratic fit of the refinancing rate to the difference in rate gaps across bins in a given time period. Negative values on the *x*-axis are borrowers for whom refinancing would be suboptimal according to ADL, whereas positive values are borrowers who should refinance.

Panel A plots the results for FHA borrowers. Two features of this figure are worth pointing out. First, refinancing rates fall after the policy change in every single bin. This is consistent with the large average fall in refinancing among FHA borrowers documented above. Second, and more importantly, the fall in refinancing is largest among precisely the set of borrowers who would benefit most from being able to refinance. Monthly refinancing rates fall from approximately 2% to 0.5% for borrowers with the largest potential benefits to refinancing. In contrast, the differences for borrowers who would either not benefit from refinancing or only marginally benefit, though positive, are negligible in economic magnitude. This implies that the majority of the overall fall in refinancing documented above is driven by borrowers with a substantial benefit to refinancing. For completeness, Panel B plots the same relationships for conventional borrowers. Consistent with the difference-in-differences analysis, there is no evidence of a systematic fall in refinancing after the policy change for these borrowers. Moreover, in this segment of the market the relationship between refinancing and the estimated benefits from doing so is generally upward-sloping and similar both before and after the policy change.[26]

As an alternative way to gauge the costs of these constraints for FHA borrowers, we can calculate the potential forgone monthly savings from not refinancing. Given their current interest rates, loan sizes, and estimated potential interest rates, FHA borrowers for whom the ADL model implies it would be optimal to refinance could have saved an average of $160 per month if they were to have refinanced into a new 30-year loan (a 17% reduction in monthly payments). This is a substantial increase in disposable income. The average FHA borrower only earns approximately $66,000 per year (Ehrlich and Perry (2015)) and most of the borrowers prevented from refinancing by these constraints likely earn far less given that the constraints apply specifically to *unemployed* borrowers. Similar differences in monthly payment burdens have been shown to have large effects on mortgage default rates and consumption in populations that likely include a larger fraction of employed households (Ehrlich and Perry (2015); Agarwal et al. (2017b); Di Maggio et al. (2017); Ganong and Noel (2017)). These effects would presumably be even larger among unemployed and liquidity-constrained households who have higher marginal propensities to consume and are most directly affected by the constraints we consider.

---

[26] The nonmonotonicity at high values of the benefits to refinancing in Panel B is likely driven by the fact that many underwater conventional borrowers were prevented from refinancing during this period. As interest rates continued to fall over time, these borrowers experienced growing potential benefits to refinancing that they were unable to act on.

IV. Mechanisms

Our results thus far indicate that the new employment documentation requirements and restrictions on financing closing costs collectively led to a large reduction in FHA refinancing that was likely to have been economically costly for the affected borrowers. In this section, we investigate heterogeneity in this response across borrowers to study how these two major program changes separately contributed to the drop in refinancing. Although occurring at the same time, these two changes affected observably distinct sets of borrowers, which allows us to plausibly trace out their respective effects. However, given that these two changes occurred simultaneously and likely interact, we are not able to decompose the overall effect into the portion that is driven by each restriction. Instead, we focus on establishing whether or not each mechanism had economically significant and independent effects on refinancing.

A. Employment Documentation Requirements

A.1. Graphical Evidence

The change in employment documentation requirements introduced in 2009 fundamentally altered the nature of the SLR Program. Prior to this change, FHA borrowers were able to qualify for an SLR regardless of their income or employment status. This meant that unemployed borrowers had the same access to refinancing as any other FHA borrower. However, the FHA changed this when it began requiring lenders to verify employment as a condition for receiving an SLR.

As evidence that this new constraint was binding, Figure 5 plots unconditional refinancing rates by month separately for FHA and conventional loans and across groups of borrowers that are more or less likely to be unemployed at the time of the policy change. To proxy for the likelihood that a borrower is unemployed, we use changes in county-level unemployment between 2006 and 2009 and categorize loans into "high-" and "low-" unemployment groups based on whether they fall into the top or bottom quarter of the distribution of these changes across counties.[27] In Panels A and B, we plot monthly refinancing rates for FHA borrowers in the high- and low-unemployment groups, respectively. Panels C and D plot the analogous refinancing rates for conventional borrowers.

---

[27] We use changes in unemployment as our proxy rather than levels to address the fact that county-level unemployment rates exhibit substantial noise that is highly correlated with the level of unemployment. For example, the correlation between the ACS reported margin of error and the level of unemployment in 2009 was $\rho = 0.42$. This type of heteroskedastic measurement error can exacerbate standard attenuation bias and may lead estimates based on the level of unemployment to be overly conservative. Taking the change alleviates this issue as the ACS-reported measurement error is only weakly correlated with the change in unemployment ($\rho = 0.16$). Moreover, when regressing the level of unemployment on the change, the coefficient on the change is effectively one. To further alleviate issues of measurement error, we also take two-year averages on either end of the difference. Formally, we calculate the 2006 to 2009 change in unemployment as $\Delta UR_{09-06} = (UR_{10} + UR_{09})/2 - (UR_{06} + UR_{05})/2$.
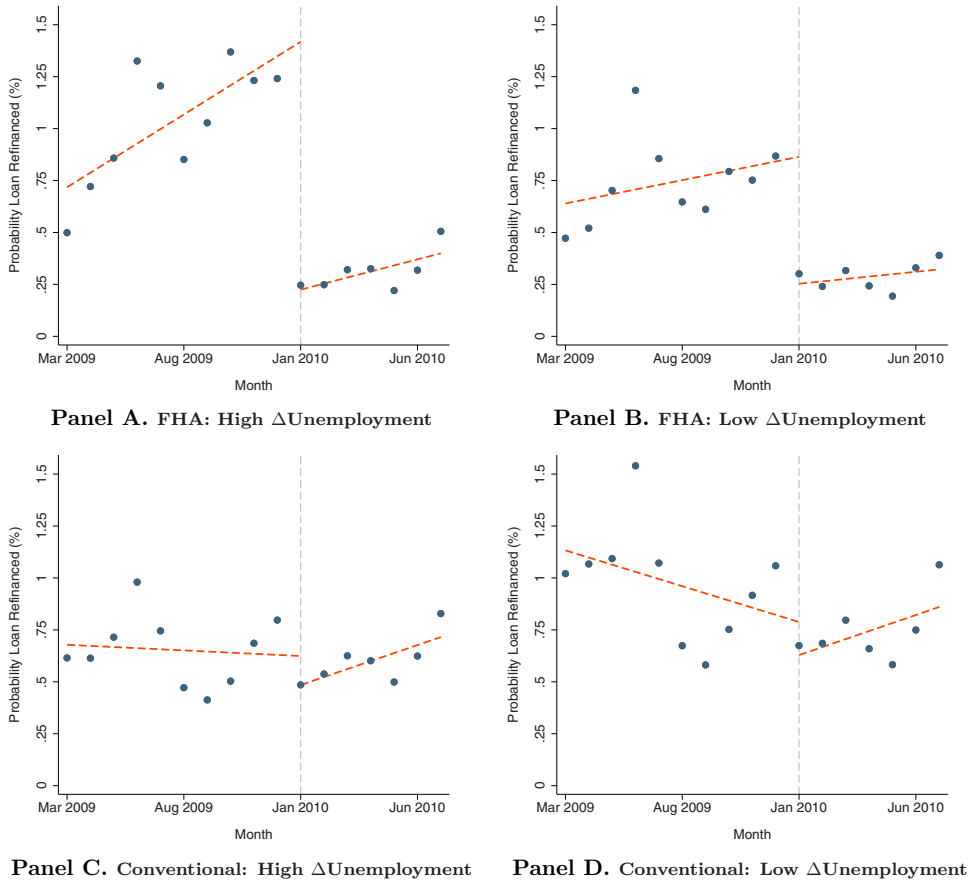
Panel A. FHA: High △Unemployment

Panel B. FHA: Low △Unemployment

Panel C. Conventional: High △Unemployment

Panel D. Conventional: Low △Unemployment

Figure 5. FHA and conventional refinancing trends by county unemployment change. This figure plots unconditional monthly refinancing rates between March 2009 and July 2010 among FHA and conventional borrowers and across counties experiencing differing changes in unemployment. Loans are categorized into "high-" and "low-" unemployment change groups based on whether they fall into the top or bottom quarter of the distribution of changes in county-level unemployment rates between 2006 and 2009. Each dot represents the percentage of outstanding mortgages in a given group that refinanced in the indicated month. Refinancing rates are calculated separately for FHA (Panels A and B) and conventional loans (Panels C and D). The vertically dashed gray line in January 2010 marks the first month that the SLR policy changes went into effect. The dashed orange lines are the predicted values from a regression of the plotted refinancing rates on a linear time trend fit separately on either side of the cutoff date. (Color figure can be viewed at wileyonlinelibrary.com)

Comparing across the panels in the top row shows that FHA borrowers in high-unemployment counties had substantially higher rates of refinancing than FHA borrowers in low-unemployment counties before the policy change. This suggests that unemployed borrowers refinance at a higher rate than their employed counterparts when they are able to do so. After the policy change, however, there is a discrete drop in refinancing and both groups of FHA

borrowers begin to refinance at similar rates. The fact that the drop in refinancing was roughly twice as large for the high-unemployment group is consistent with the idea that the new income documentation requirements were more binding for this group. In contrast, refinancing rates in the conventional market (Panels C and D) remain constant around the time of the policy change and are always somewhat higher in counties with smaller increases in unemployment. Although not conclusive, these results suggest that unemployed FHA borrowers had a high demand for refinancing during the pre-period that was substantially constrained by the new employment documentation requirements.

A.2. Empirical Strategy

To more formally analyze the effect of the change in employment documentation requirements, we use a triple-differences strategy that is directly motivated by the results in Figure 5 but that allows us to control for many other factors that are correlated with unemployment and also related to refinancing. The idea behind this strategy is to compare changes in refinancing behavior before and after the policy change across groups of FHA borrowers who are more or less likely to be unemployed while using similar changes in the conventional market as a counterfactual for what would have happened in the absence of the policy. As in Figure 5, we use changes in county-level unemployment as a proxy for the likelihood that a borrower is unemployed. The identifying assumption in this context is that, conditional on the controls we include, the differential change in FHA refinancing rates across counties that experienced high and low changes in unemployment would have paralleled that in the conventional market in the absence of the policy.

We implement this approach by estimating versions of the following triple-differences regression:

$$
\begin{aligned}
Refinance_{it} = {} & \alpha + \delta_t + X_{it}'\gamma + \beta_0 \cdot FHA_i + \beta_1 \cdot \Delta UR_i \\
& + \beta_2 \cdot FHA_i \times Post_t + \beta_3 \cdot \Delta UR_i \times Post_t + \beta_4 \cdot FHA_i \times \Delta UR_i \\
& + \beta_5 \cdot FHA_i \times \Delta UR_i \times Post_t + \epsilon_{it}.
\end{aligned}
\tag{5}
$$

In this specification, the variable ΔUR_i measures the change in the unemployment rate from 2006 to 2009 in borrower i's county, and all other terms are as previously defined. As in the difference-in-differences analysis above, the set of controls X_{it} always includes FHA-specific linear time trends that are allowed to differ before and after the policy change. These trends are included to adjust for the fact that FHA borrowers may have had a differential capacity to refinance in response to the deterioration in economic conditions leading up to the policy change. The coefficient of interest is β_5, which provides a measure of how much FHA refinancing rates fall relative to conventional loans following the policy change and as the likelihood of unemployment increases.

If employment documentation requirements were an important barrier to refinancing during this period, we should expect to find $\beta_5 < 0$.

One key confound that needs to be considered when estimating equation (5) is the borrower's level of home equity. Although FHA borrowers could qualify for an SLR regardless of home equity throughout the entire sample period, only borrowers with positive equity were able to roll closing costs into their loans subsequent to the policy change. Because house prices and unemployment are highly correlated during this period, a simple comparison that does not control effectively for home equity would risk conflating the effect of the income documentation requirements with the increased upfront costs for negative-equity borrowers. We address this concern in two ways. First, when estimating (5), we always focus only on the sample of borrowers who are in positive equity in a given month. This group of borrowers is able to finance their closing costs using the new loan both before and after the policy change and should therefore only be affected by the change in income documentation requirements. Second, because our estimate of home equity is measured with error, we also estimate versions of the specification that include controls for the complete interaction between the *FHA* indicator, the *Post* indicator, and a set of dummies for the borrower's current equity ($10,000 bins). This specification allows for home equity to have a separate effect on refinancing for FHA and conventional borrowers both before and after the policy change and should therefore control for any residual independent effect of the new closing cost requirements arising from the fact that some borrowers who we deem to be in positive equity are not.[28]

Finally, it is important to note that the results from these triple-difference regressions only pick up the *differential* fall in refinancing between FHA borrowers in high- relative to low-unemployment counties. If the introduction of the employment documentation requirements also reduced refinancing rates among employed borrowers, perhaps due to the increased psychological burden associated with providing documentation, then our results will not reflect this effect.[29] Unfortunately, we are not aware of any reliable proxy for cross-sectional variation in these types of costs that would allow for a direct measurement of the role that such hassle costs may have played. Although these costs are likely to be relatively small in comparison to the substantial potential benefits to refinancing discussed in Section III, below we discuss how their presence may affect the interpretation of our results.

---

[28] In Section II of the Internet Appendix, we also estimate an alternative to this specification in the full sample of loans allowing for both the employment documentation and closing costs to separately affect FHA refinancing rates. The results from this specification are qualitatively similar to those from our preferred estimates in the restricted sample of borrowers who are in positive equity.

[29] These effects would, however, be reflected in the difference-in-differences and event study estimates.

<div align="center">

Table V

Refinancing and Unemployment: Triple-Difference Estimates

</div>

This table reports estimates of the effect of the change in SLR income documentation requirements on the monthly probability that an FHA loan refinances. In all specifications the sample includes only loan-months for which the borrower is estimated to be in positive equity. The outcome variable is an indicator for whether or not a loan refinances in the month of observation and is multiplied by 100, so that all coefficients can be interpreted as percentage-point changes. For reference, column (1) reports estimates from a baseline difference-in-differences regression that is directly analogous to the specification in column (4) of Table IV, but that only includes loan-months in positive equity. Coefficients are reported for the *FHA* "treatment" dummy as well as its interaction with the *Post* dummy. Columns (2) to (5) report estimates from triple-difference regressions that further interact the *FHA* and *Post* dummies with the change in county-level employment, which is also included linearly in the regression. The *Post* dummy takes the value of 1 if the month of observation is after the implementation of the policy changes (January 2010). All specifications include fixed effects for the month of observation and the CBSA of the property as well as linear time trends for FHA borrowers that are allowed to vary freely before and after the policy change. Column (3) adds fixed effects for the current loan age (one-year bins), interest rate (one-percentage-point bins), and the borrower's estimated home equity ($10,000 bins), as well as the full pairwise interaction between the borrower's LTV (10-point bins) and FICO score (50-point bins) at origination. Column (4) further interacts all of the additional fixed effects contained in column (3) with the *Post* dummy and the *FHA* indicator. Column (5) allows for an unrestricted effect of home equity by interacting each $10,000 home equity bin fixed effect with the complete interaction between the *FHA* and *Post* indicators. Standard errors are reported in parentheses and are clustered at the CBSA level. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
FHA	1.498***	−0.054	0.353***	1.321***	1.208***
	(0.356)	(0.055)	(0.065)	(0.355)	(0.349)
FHA × Post	−0.476***	−0.280***	−0.231***	−0.315***	−0.349***
	(0.046)	(0.081)	(0.073)	(0.068)	(0.077)
FHA × ΔUR		0.046***	0.042***	0.047***	0.047***
		(0.011)	(0.012)	(0.012)	(0.012)
FHA × ΔUR × Post		−0.057***	−0.047***	−0.046***	−0.047***
		(0.015)	(0.013)	(0.014)	(0.014)
Month FEs	X	X	X	X	X
CBSA FEs	X	X	X	X	X
FHA Time Trends	X	X	X	X	X
Loan Age FEs	X		X	X	X
Interest Rate FEs	X		X	X	X
LTV × FICO FEs	X		X	X	X
Equity FEs	X		X	X	X
Controls × Post	X			X	X
Controls × FHA	X			X	X
Equity FEs × FHA × Post					X
Number of Observations	13,250,266	13,250,266	13,250,266	13,250,266	13,250,266

A.3. Results

Columns (2) to (5) of Table V present our triple-difference estimates of the effect of the change in employment documentation requirements on FHA refinancing. For reference, in column (1) we also report results from a difference-in-differences specification that is identical to the specification used to estimate

the overall effect in column (4) of Table IV but estimated only in the subsample of borrowers with positive equity. The effect of the policy in this subsample is roughly two-thirds as large as the corresponding effect in the full sample. This difference is exactly what would be expected given that the closing cost requirements are nonbinding for borrowers with positive equity.

Column (2) reports estimates from a baseline version of the triple-differences regression that includes only the FHA time trends and a set of month and CBSA fixed effects as controls. The coefficient estimate on the triple-interaction term in the bottom row implies that the policy-induced drop in refinancing rates for FHA borrowers increases by roughly 0.06 percentage points for each one-percentage-point increase in the county-level unemployment rate. The average loan in our sample was in a county that experienced an increase in unemployment of roughly four percentage points. At that change in unemployment rates, this effect is enough to account for a reduction in refinancing of 0.23 percentage points per month, which is roughly half the size of the overall effect reported in column (1).

In the remaining columns of the table, we add a series of control variables that increasingly restrict the nature of the variation being used to identify how the fall in FHA refinancing depends on local employment conditions. In column (3), we include the same detailed set of fixed effects for loan-level characteristics that were included in our analysis of the overall effect of the policy change (loan age, interest rate, current equity bin, and LTV-by-FICO bins). Column (4) further interacts these controls with the *FHA* dummy and the *Post* indicator. In both cases, the coefficient on the triple-interaction term remains negative, qualitatively similar to, and statistically indistinguishable from the baseline estimate in column (2). In column (5), we allow the controls for home equity to enter even more flexibly by interacting each equity bin fixed effect not only with the *FHA* and *Post* dummies, but also with their interaction. This specification explicitly controls for the differential effect that the change in FHA policy may have had on FHA borrowers through the importance of equity and its relation to closing costs. The estimate is nearly identical to that in column (4), which provides confidence that these results primarily reflect the effect of the change in employment documentation requirements.

To further explore the robustness of this relationship, in Panel A of Figure 6, we report estimates from an alternative and more flexible parameterization that allows for the effect to vary by month of observation. Specifically, we plot coefficient estimates from the following specification:

$$Refinance_{it} = \alpha + \delta_t + X'_{it}\gamma$$
$$+ \sum_\tau \mathbb{1}_{t=\tau} \cdot \left[\beta_{0\tau} \cdot FHA_i + \beta_{1\tau} \cdot \Delta UR_i + \beta_{2\tau} \cdot FHA_i \times \Delta UR_i\right] + \epsilon_{it}.$$

$$(6)$$

This specification allows for separate monthly coefficients on FHA status, changes in local unemployment, and the interaction between the two. In the figure, we plot the $\beta_{2\tau}$ coefficients, which measure how the gap between FHA
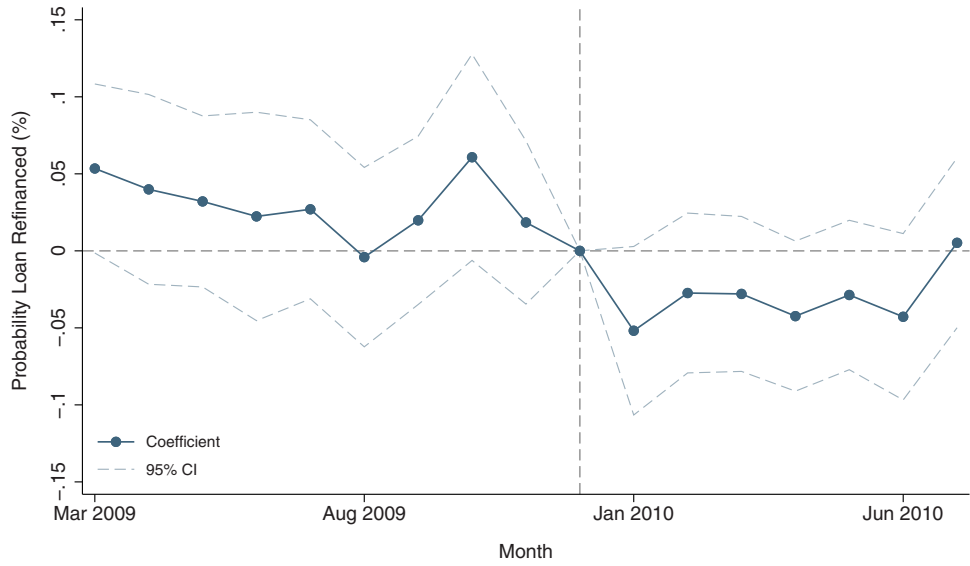
Panel A. Employment Documentation Effect
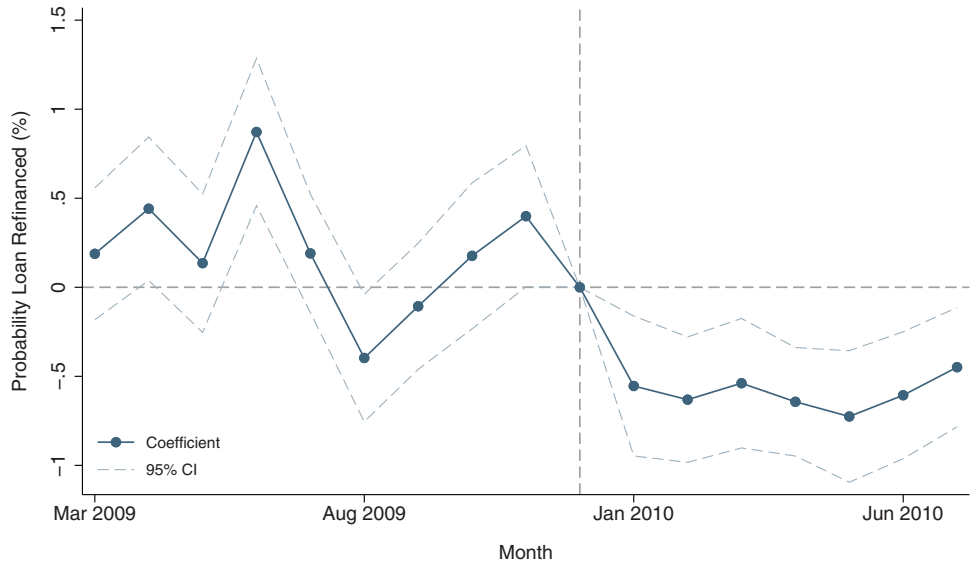
Panel B. Upfront Costs Effect

Figure 6. Triple-difference estimates by month of observation. This figure reports estimates of how the effect of the change in FHA policies on FHA refinancing varied across differences in county-level unemployment rates (Panel A) and borrower equity (Panel B). Estimates are derived from a flexible triple-difference specification that allows the effect to vary freely by month of observation. Panel A reports coefficient estimates from a regression of an indicator for refinancing on the interaction between the FHA dummy, county-level unemployment changes, and month fixed effects. Panel B reports estimates from a similar regression that instead interacts the FHA dummy and month fixed effects with an indicator for whether the borrower's estimated home

and conventional refinancing is related to local unemployment rates during each month in our sample period. As before, we normalize the coefficient for December 2009 to zero, so that each estimate can be interpreted as the effect relative to the month just prior to the policy change. Although the estimates have relatively wide confidence intervals, there is a noticeable level shift that occurs in precisely the month that the policy takes effect. This provides confidence that the results in Table V are driven directly by the policy change and not by some other omitted factor that is correlated with unemployment.

A.4. Interpreting the Magnitudes

Although our results provide compelling evidence that the new employment documentation requirements were a substantial barrier to refinancing for FHA borrowers during this period, it is important to note that the magnitude of these estimates may not extrapolate to the general population. In particular, given their demographic characteristics, it is likely that changes in local unemployment load more strongly on FHA borrowers relative to the average household. That is, a one-percentage-point increase in county-level unemployment may translate into a greater than one-percentage-point increase in unemployment among FHA borrowers. If true, this would lead us to overestimate the effect of employment documentation requirements for the typical borrower.

To provide a sense of how large this bias may be, we use data from the Survey of Consumer Finances (SCF) to measure how the change in unemployment between 2007 and 2009 among households with an outstanding FHA mortgage in 2007 compared to the same change for all households over that period. For FHA borrowers, the unemployment rate in the SCF increased by 7.8 percentage points over this period, whereas the increase among all households was only 6.2 percentage points. If we assume that changes in county-level unemployment load similarly on the two groups of borrowers, this would imply that a one-percentage-point increase in the local unemployment rate translates into an increase in unemployment for FHA borrowers of roughly 1.25 (7.8/6.2) percentage points.

Even with this scaling, however, the results in Table V imply substantial effects of employment documentation requirements on refinancing rates. For example, dividing the coefficient estimate in the bottom row of column (5) by 1.25 would yield an estimate of 0.0376. This implies that the likelihood of refinancing for the average borrower in our sample was roughly 0.15 percentage points lower than it otherwise would have been as a result of the fall in

---

equity is less than or equal to 0. Both regressions include fixed effects for the month of observation and the CBSA of the property. The coefficient for December 2009 is normalized to zero, so that all estimates can be interpreted as the differential change in the monthly probability of refinancing relative to the month prior to when the policy changes went into effect. The 95% confidence intervals are based on standard errors that are clustered at the CBSA level. (Color figure can be viewed at wileyonlinelibrary.com)

employment between 2006 and 2009.[30] This effect would be even larger in counties hit by above-average employment shocks, which has important implications for the distributional impact of stimulative policies, such as the first round of quantitative easing (QE1), that operate in large part through the mortgage refinancing channel (Di Maggio, Kermani, and Palmer (2016)).

As an alternative way to interpret our triple-difference estimates, we can also convert them into an implied rate at which unemployed borrowers would refinance if permitted. This rate provides a direct measure of the latent demand for refinancing among unemployed borrowers. For clarity of exposition, we work with a simplified expression for the probability of refinancing. In particular, let r_0^{FHA} and r_1^{FHA} denote the monthly probabilities that an FHA borrower refinances during the pre- and post-policy periods, respectively. If we assume that employed and unemployed borrowers each refinance at constant rates when permitted, then we can express these two FHA refinancing rates as weighted averages of the refinancing rates of employed (r_E) and unemployed (r_U) borrowers,

$$r_0^{FHA} = r_U \times UR + r_E \times (1 - UR)$$
$$r_1^{FHA} = r_E \times (1 - UR), \tag{7}$$

where UR denotes the unemployment rate and the second line follows from the fact that unemployed FHA borrowers are prohibited from refinancing after the policy change.

Using similar expressions for the refinancing probabilities in the conventional market, we can express our difference-in-differences estimate as

$$DiD \triangleq \left(r_1^{FHA} - r_0^{FHA} \right) - \left(r_1^{Conv} - r_0^{Conv} \right)$$
$$= \left(r_1^{FHA} - r_0^{FHA} \right) - 0$$
$$= -r_U \times UR, \tag{8}$$

where the second line follows from the assumption that the change in FHA policy had no effect on conventional refinancing rates.[31] The coefficient of interest reported in our regressions is a triple-differences estimate, which is simply the partial derivative of the difference-in-differences estimate with respect to the unemployment rate,

$$\frac{\partial DiD}{\partial UR} = -r_U = -0.047 \Rightarrow r_U = 0.047. \tag{9}$$

Rescaled by 100 to account for the units of the unemployment rate, our estimates imply that unemployed borrowers would refinance at a monthly rate of

---

[30] This calculation is based on a 4% average change in county-level unemployment: $\frac{-0.047}{1.25} \times 4.0 = -0.15$.

[31] For simplicity, we omit borrower characteristics as well as secular trends in refinancing and unemployment rates from this analysis as they would simply difference out in the end.

about 4.7% if permitted. This is roughly five times the refinancing rate of the average borrower during this period. Alternatively, if we use the same 1.25 scaling factor as above, the implied rate at which an unemployed borrower would refinance if permitted would be 3.76%.

The derivation of equation (9) above relies on the assumption that the refinancing rate for employed FHA borrowers is the same before and after the policy change. However, if employed borrowers also reduced their refinancing rates, perhaps due to the added psychological hassle associated with providing documentation, then our results would actually *underestimate* the true refinancing rate among unemployed borrowers. To see this, suppose that employed FHA refinancing falls from r_E to $r_E - \rho$ after the policy change, where ρ represents the fall in refinancing due to hassle costs. In this case, the second line of equation (7) would be $r_1^{FHA} = (r_E - \rho) \times (1 - UR)$, which would yield an alternative expression for the difference-in-differences estimate: $DiD \triangleq -r_U \times UR - \rho \times (1 - UR)$. Differentiating this expression with respect to the unemployment rate would imply that our triple-difference estimate is actually equal to $\frac{\partial DiD}{\partial UR} = -r_U + \rho = -0.047$, which underestimates the true refinancing rate for unemployed borrowers by exactly the amount of the (unobserved) fall in refinancing among employed borrowers due to hassle costs.

A.5. *Implied Aggregate Effects of Employment Documentation Requirements*

Thus far, our analysis has focused on the direct effect of the changes to the SLR Program on FHA borrowers themselves. However, the estimate from equation (9) also provides us with a way to evaluate how the standard practice of prohibiting unemployed borrowers from refinancing may have affected aggregate refinancing outside the FHA market during the Great Recession. In this section, we extrapolate our estimate to the general population to perform a simple back-of-the-envelope calculation demonstrating that this effect may indeed have been substantial.

As a starting point for this calculation, note that the Bureau of Labor Statistics (BLS) estimated that there were 15.1 million unemployed persons in the United States as of December 2009.[32] Using data from the 2009 SCF, we calculate that roughly 38% of these unemployed persons also held mortgages.[33] However, even if they were permitted to refinance, not all of these unemployed mortgage borrowers would have been able to do so. In particular, many of them were likely to be in negative equity, which itself was a binding constraint on refinancing outside the FHA market during this time. Using information on

---

[32] Source: U.S. Bureau of Labor Statistics, Unemployment Level [UNEMPLOY], retrieved from FRED, Federal Reserve Bank of St. Louis, https://fred.stlouisfed.org/series/UNEMPLOY, August 19, 2019.

[33] Strictly speaking, the SCF is a survey of households, not persons. We are, however, unaware of any other data set that would allow us to calculate the share of unemployed persons with a mortgage in 2009.

self-reported house values and the balance of outstanding mortgages reported in the SCF, we estimate that roughly 25% of unemployed borrowers were in negative equity as of 2009.[34] Together, this means that there were approximately 4.3 million unemployed mortgage borrowers in positive equity who were potentially prevented from refinancing in 2009 due to employment documentation requirements. Our estimates from equation (9) imply that these borrowers would have refinanced at a monthly rate of 4.7% if permitted. This means that there would have been roughly 2.4 million more refinances in 2009 had these unemployed borrowers been able to refinance. Our data on potential interest rates discussed in Section III indicate that the average borrower who refinanced in 2009 would have lowered their interest rate by roughly 104 bps, resulting in annual savings of approximately $3,000. Thus, over the course of 2009, our estimates imply that allowing unemployed borrowers to refinance would have provided these already cash-strapped borrowers with an additional $7.2 billion in disposable income. This represents almost 10% of the entire $79.5 billion paid out to unemployed people through the unemployment insurance system in 2009.[35]

As an alternative way of putting these numbers into context, we can also compare them to the increase in refinancing and mortgage modification that was achieved by the two largest household debt relief policies implemented during the financial crisis: HARP and HAMP. These policies aimed to provide debt relief to underwater and distressed borrowers by extending government-backed credit guarantees to cover refinances of negative-equity GSE loans and by providing incentives for private servicers to renegotiate terms on delinquent mortgages. Together, Agarwal et al. (2017a, 2017b) estimate that HARP and HAMP generated an additional 3 million refinances and 1 million mortgage modifications between March 2009 and December of 2012, and that these refinances and modifications saved borrowers between $3,000 to $4,000 annually. Importantly, both of these programs explicitly excluded unemployed borrowers from participating. Thus, our estimates suggest that implementing a broad-based streamline refinancing program aimed at easing employment documentation requirements could have potentially generated an additional increase in disposable income of roughly the same order of magnitude as that achieved by policies focused on relieving negative-equity constraints and agency conflicts in the secondary mortgage market.

---

[34] In performing this calculation we aggregate balances on both first and second mortgages outstanding and compare them to an adjusted estimate of the home value that accounts for the well-known upward bias in self-reported housing wealth. In particular, we multiply the self-reported house value by 0.9, reflecting the findings in Benítez-Silva et al. (2015) that homeowners over-report housing values by about 8% on average and in Chan et al. (2016) that this bias is larger during housing busts. Using this correction factor, the implied share of *all* mortgage borrowers (not just the unemployed) that were in negative equity in 2009 is 23.9%, which is very close to the aggregate estimates reported by CoreLogic at the time.

[35] Source: U.S. Department of Labor, Monthly Program and Financial Data, retrieved from https://oui.doleta.gov/unemploy/claimssum.asp, August 19, 2019.

A.6. *Employment Documentation Requirements versus Negative-Equity Constraints*

One of the key features of our empirical setting is that it allows us to isolate the importance of employment documentation requirements from that of other potentially correlated barriers to refinancing. Chief among these is negative equity. As noted by Beraja et al. (2019), negative equity was both highly correlated with unemployment during the Great Recession and a significant barrier to refinancing. Because underwater FHA borrowers were always permitted to refinance throughout our entire sample period, our estimates reveal the direct effect of employment documentation requirements among a subset of households who were unaffected by these negative-equity constraints. However, this does not directly answer the question of whether employment documentation requirements are an independently important barrier to refinancing in the general population. If all unemployed borrowers are in negative equity, then relaxing employment documentation requirements would have no effect on the refinancing rate outside the FHA market. The statistic we cite from the SCF above—that only 25% of households with a mortgage and an unemployed member were in negative equity in 2009—suggests that this is unlikely to be the case. In Figure 7, we provide further evidence indicating that employment documentation requirements are an independent constraint on refinancing that, in many recessions, may even be more important than negative equity.

Panel A of Figure 7 plots the cross-sectional relationship between county-level unemployment and negative equity for borrowers in our main analysis sample. For each county, we calculate the share of loans that were ever in negative equity at some point in our sample period and then group counties into bins according to these shares. Within these groups we then plot the cross-county distribution of changes in unemployment between 2006 and 2009. Each box plot in the figure reports the 5^{th}, 25^{th}, 50^{th}, 75^{th}, and 95^{th} percentiles of the distribution of unemployment rate changes among counties with negative-equity shares in the ranges indicated on the x-axis. The mean of this distribution is indicated by the orange diamond associated with each box plot. As is well known, the two measures are closely related on average; mean unemployment rises with the negative-equity share. However, the figure also makes clear that there is enormous variation around this trend. For example, unemployment rates increased by up to almost 10 percentage points even in counties where the share of negative-equity loans never exceeded the national average of 25%. Moreover, regardless of negative-equity share, there were always some counties experiencing increases in unemployment above the national average of 5%. Thus, consistent with the evidence from the SCF, there are many borrowers in our sample for whom employment documentation requirements were likely to bind even though negative equity constraints did not.

Panel B of Figure 7 further demonstrates that the positive relationship between unemployment and negative equity that existed during the Great Recession was atypical of recent U.S. experience. The figure plots time-series
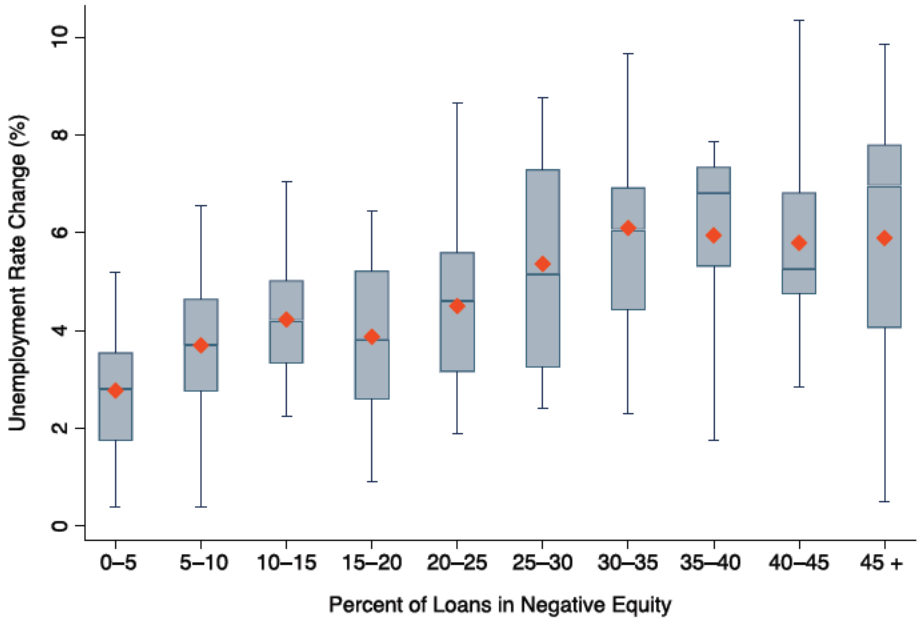
Panel A. Unemployment and Negative Equity in the Cross Section
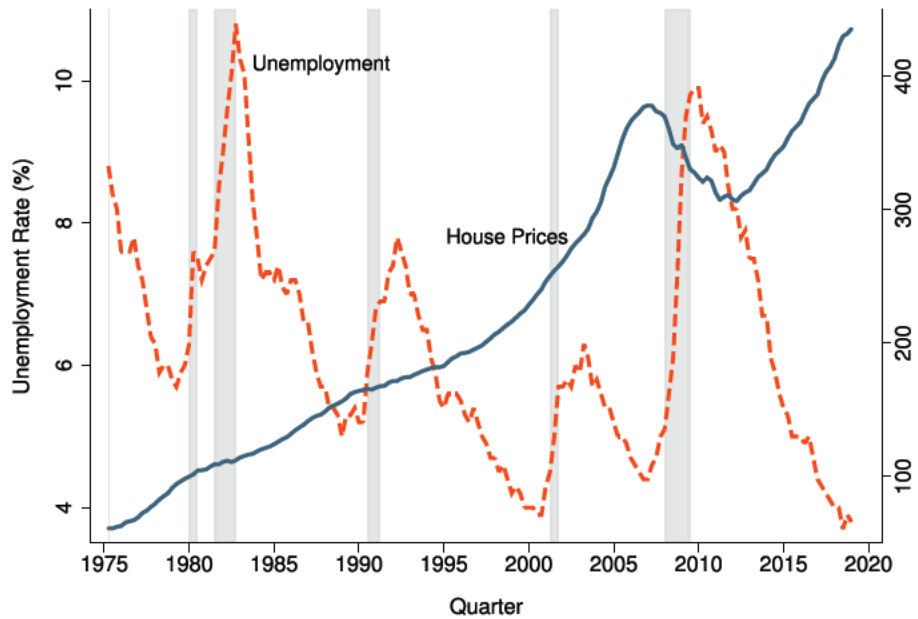
Panel B. Unemployment and House Prices in the Time Series

Figure 7. The relationship between unemployment and negative equity. This figure plots cross-sectional and time-series patterns in the relationship between unemployment, negative equity, and house prices. Panel A plots the relationship between county-level unemployment changes

patterns in unemployment and house prices at the national level over the longest horizon for which we can reliably measure aggregate house prices. The unemployment rate is plotted on the left axis using data from the BLS, and house prices are plotted on the right axis using the Federal Housing Finance Agency (FHFA) House Price Index. The time periods shaded in gray correspond to NBER-dated recessions. As expected, the unemployment rate spikes in all five recessions shown. The patterns for house prices, however, are remarkably different. It is only during the Great Recession that national house prices fall. In all four of the other recessions during which it is possible to measure aggregate house prices, there is no meaningful or sustained depreciation. This means that it was employment documentation requirements, not negative equity, that were more likely to be the binding constraint on refinancing during the "typical" recession. In that sense, we think that the estimates in this paper, which focus on employment documentation requirements, are a strong complement to existing evidence on the importance of negative equity as a barrier to refinancing during the Great Recession.

B. Upfront Costs

The second major change to the SLR Program was the reduction in maximum loan amounts for SLRs without an appraisal. As discussed in Section I, this change eliminated the ability for negative equity borrowers to roll the upfront closing costs of refinancing into their new loan. Instead, after the policy change, these borrowers would now have to pay for any upfront costs out-of-pocket. Thus, while negative-equity FHA borrowers were still fully permitted to refinance through the SLR Program, the out-of-pocket costs associated with doing so were significantly higher after the policy change. To study the effects of this change on refinancing rates, we proceed in the same manner as in our analysis of the employment documentation requirements. First, we present simple graphical evidence indicating that this new constraint appears to have had a larger effect on refinancing among borrowers that were more likely to have been affected by it. Second, we estimate triple-differences regressions that are motivated by this evidence and that allow us to more precisely quantify the extent to which the need to pay for closing costs out-of-pocket constrains refinancing.

---

and negative equity during the Great Recession. Counties in the primary analysis sample are grouped into bins based on the share of loans in the county that were ever in negative equity at some point during the sample period. Each box plot in the figure reports the 5th, 25th, 50th, 75th, and 95th percentiles of the distribution of unemployment rate changes between 2006 and 2009 among counties with negative-equity shares in the ranges indicated on the *x*-axis. The mean change in unemployment for counties in a given bin is indicated with an orange diamond. Panel B plots the quarterly time series of national unemployment from the BLS (left axis) and house prices as measured by the FHFA single-family house price index (right axis). The house price index is normalized to one in the first quarter of 1980. Shaded gray areas correspond to NBER-dated recessions. (Color figure can be viewed at wileyonlinelibrary.com)
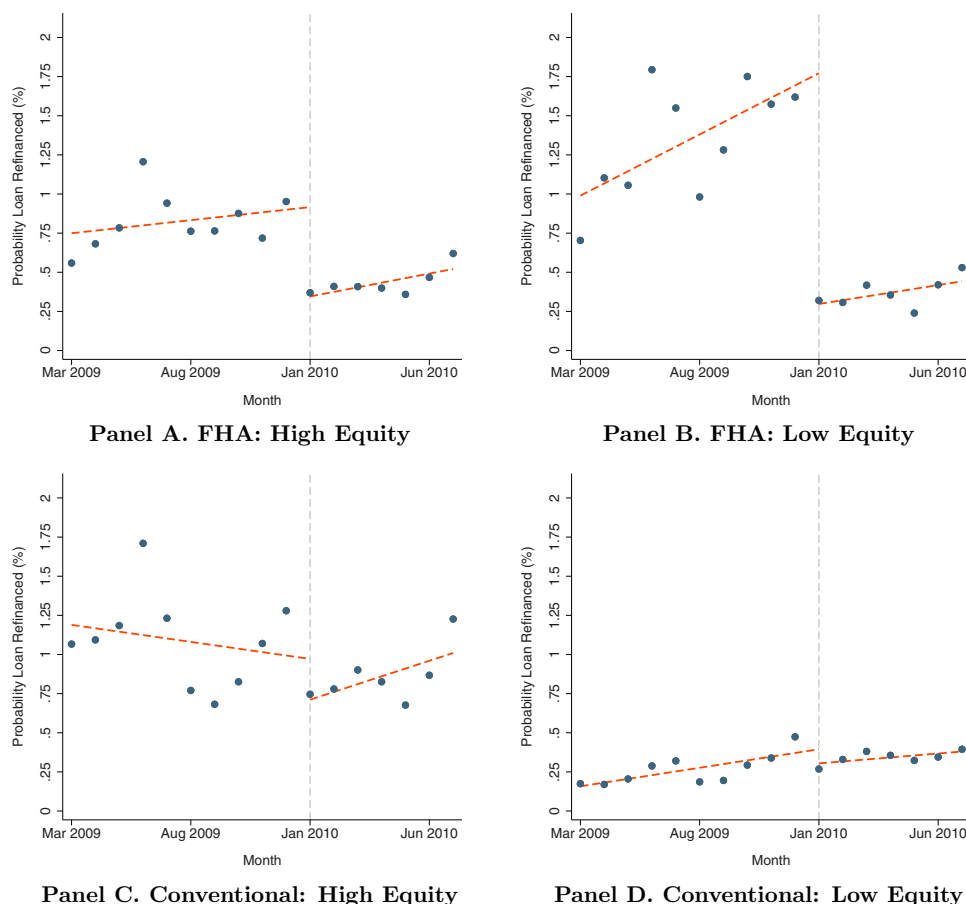
Panel A. FHA: High Equity

Panel B. FHA: Low Equity

Panel C. Conventional: High Equity

Panel D. Conventional: Low Equity

Figure 8. FHA and conventional refinancing trends by borrower's equity. This figure plots unconditional monthly refinancing rates between March 2009 and July 2010 among FHA and conventional borrowers with differing levels of home equity. Borrowers are categorized as "high" equity if their estimated level of home equity is greater than or equal to $20,000 and "low" equity if it is less than or equal to $0. Home equity is estimated by subtracting the borrower's current outstanding balance from an estimate of the current home value derived from the initial purchase price and subsequent growth implied by the relevant county-level Zillow home price index. Each dot represents the percentage of outstanding mortgages in a given group that refinanced in the indicated month. Refinancing rates are calculated separately for FHA (Panels A and B) and conventional loans (Panels C and D). The vertically dashed gray line in January 2010 marks the first month that the SLR policy changes went into effect. The dashed orange lines are the predicted values from a regression of the plotted refinancing rates on a linear time trend fit separately on either side of the cutoff date. (Color figure can be viewed at wileyonlinelibrary.com)

B.1. Graphical Evidence

Figure 8 presents unconditional monthly refinancing rates for FHA and conventional borrowers with differing levels of home equity. We categorize borrowers into "high-" and "low-" equity groups based on whether their estimated

home equity in the month of observation is greater than or equal to $20,000 or less than or equal to $0, respectively.

Panels A and B plot refinancing rates separately for high- and low-equity FHA borrowers. Comparing across these panels reveals that both positive- and negative-equity FHA borrowers experienced a discrete fall in refinancing in the month of the policy change. However, this fall was nearly three times as large for negative-equity borrowers. Importantly, this differential fall is not a direct result of negative equity itself, since both groups of borrowers were still permitted to refinance through the SLR Program provided that they could pay the upfront costs. Rather, it is consistent with the idea that the new need to pay for closing costs out-of-pocket, which only affected negative-equity borrowers, was a binding constraint on refinancing during this period.

Panels C and D of the figure plot the analogous refinancing rates for conventional borrowers and show that there was essentially no change in refinancing among either group around the date of the policy change. These figures also make clear that negative equity itself *was* a binding constraint in the conventional market. Unlike in the FHA market, in each month high-equity conventional borrowers refinance at substantially higher rates than their low-equity counterparts. Taken together, these results suggest that requiring negative-equity FHA borrowers to pay for closing costs out-of-pocket may have created a substantial barrier to refinancing.

B.2. Empirical Strategy

To more precisely quantify the magnitude of these effects, we use a triple-differences framework that directly parallels the approach we used to estimate the income documentation effects above. The idea is to compare changes in refinancing between high- and low-equity FHA borrowers relative to conventional borrowers following the policy change while controlling flexibly for other potential drivers of refinancing. To do this, we categorize borrowers into "high-" and "low-" equity groups and estimate the following triple-differences regression:

$$\begin{aligned} Refinance_{it} = {} & \alpha + \delta_t + X'_{it}\gamma + \beta_0 \cdot FHA_i + \beta_1 \cdot LowEquity_{it} + \beta_2 \cdot FHA_i \times Post_t \\ & + \beta_3 \cdot LowEquity_{it} \times Post_t + \beta_4 \cdot FHA_i \times LowEquity_{it} \\ & + \beta_5 \cdot FHA_i \times LowEquity_{it} \times Post_t + \epsilon_{it}. \end{aligned} \tag{10}$$

In this specification, $LowEquity_{it}$ is an indicator for whether borrower i's estimated level of home equity in month t is less than or equal to 0 and all other terms are as previously defined. The coefficient of interest is β_5, which measures the difference in refinancing probability for FHA borrowers with low equity relative to FHA borrowers with high equity after the policy is in effect, relative to the same difference in the conventional market. As before, the identifying assumption is that the change in refinancing rates for high- and low-equity FHA borrowers would have evolved in parallel with the same change in refinancing for conventional borrowers in the absence of the policy. To increase

the likelihood that this assumption holds, we continue to include FHA-specific linear time trends that are allowed to vary freely before and after the policy change among our set of controls.

One key concern, as with our results on the effects of unemployment, is that shocks to income and movements in house prices are correlated over this period. This could lead us to conflate the effect of the new closing cost requirements with the employment documentation effect. To reduce the likelihood that our estimates merely reflect variation in unemployment, when estimating equation (10), we always restrict attention to the subsample of counties with below-median increases in unemployment. We also estimate specifications that allow for the effect of unemployment to flexibly and differentially affect FHA borrowers before and after the policy change. In particular, our preferred specification includes controls for the complete interaction between the *FHA* dummy, the *Post* indicator, and a set of fixed effects for each decile of the distribution of county-level unemployment changes. These steps should reduce the likelihood that our analysis is conflating the need for negative-equity borrowers to pay for closing costs out-of-pocket with their potential inability to document employment.

B.3. Results

Table VI presents our estimates of the effect of requiring borrowers to pay for upfront costs out-of-pocket. For reference, column (1) reports baseline results from a difference-in-differences specification identical to our preferred specification used to measure the overall effect in Table IV but restricted to the set of counties with below-median increases in unemployment. As expected, the effect in this subsample is substantially smaller than in the full sample due to the fact that the employment documentation requirements are less likely to bind for borrowers in these counties.

The remaining columns present our main triple-difference estimates. Column (2), which controls only for month fixed effects, CBSA fixed effects, and FHA-specific linear time trends, suggests that FHA borrowers who were required to pay for closing costs out-of-pocket were 0.8 percentage points less likely to refinance subsequent to the policy change relative to their positive-equity counterparts. Column (3) controls flexibly for a detailed set of fixed effects for various loan-level characteristics (loan age, interest rate, and LTV-by-FICO bins) as well as a set of fixed effects denoting which decile of the distribution of county-level unemployment changes the loan falls into. Column (4) further interacts these controls with the *FHA* and *Post* indicators. The coefficient estimate falls to -0.56 when these controls are included but remains statistically significant and is still economically quite large. In column (5) we directly control for the independent effect of the employment documentation requirements by fully interacting each unemployment change decile fixed effect with the *FHA* and *Post* indicators. Including these controls has a minimal effect on the coefficient, which provides assurance that our sample restriction and other controls are effectively isolating the effect of the closing cost

Table VI
Refinancing and Upfront Costs: Triple-Difference Estimates

This table reports estimates of the effect of the change in SLR closing cost requirements on the monthly probability that an FHA loan refinances. In all specifications, the sample includes only counties for which the 2006 to 2009 change in unemployment was below the median across counties. The outcome variable is an indicator for whether or not a loan refinances in the month of observation and is multiplied by 100, so that all coefficients can be interpreted as percentage-point changes. For reference, column (1) reports estimates from a baseline difference-in-differences regression that is directly analogous to the specification in column (4) of Table IV, but that only includes counties with below median increases in unemployment rates. Coefficients are reported for the *FHA* "treatment" dummy as well as its interaction with the *Post* dummy. Columns (2) to (6) report estimates from triple-difference regressions that further interact the *FHA* and *Post* dummies with an indicator for whether the borrower's estimated home equity is less than 0 ("Low Equity"), which is also included in the regression. The *Post* dummy takes the value of 1 if the month of observation is after the implementation of the policy changes (January 2010). All specifications include fixed effects for the month of observation and the CBSA of the property as well as linear time trends for FHA borrowers that are allowed to vary freely before and after the policy change. Column (3) adds fixed effects for the current loan age (one-year bins), interest rate (one-percentage-point bins), and county-level unemployment change decile associated with the loan, as well as the full pairwise interaction between the borrower's LTV (10-point bins) and FICO score (50-point bins) at origination. Column (4) further interacts all of the additional fixed effects contained in column (3) with the *Post* dummy and the *FHA* indicator. Column (5) allows for an unrestricted effect of unemployment changes by interacting each unemployment change decile fixed effect with the complete interaction between the *FHA* and *Post* indicators. Column (6) drops borrowers for whom refinancing is not optimal if they have to pay closing costs. See the text for more details on how estimates of refinancing optimality are constructed. Standard errors are reported in parentheses and are clustered at the CBSA level. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
FHA	0.627	0.060	0.299***	0.650	0.522	1.073*
	(0.432)	(0.044)	(0.055)	(0.423)	(0.436)	(0.596)
FHA × Post	−0.551***	−0.403***	−0.357***	−0.458***	−0.272***	−1.032***
	(0.059)	(0.066)	(0.061)	(0.050)	(0.067)	(0.251)
FHA × Low Equity		1.089***	0.969***	0.899***	0.869***	1.080***
		(0.165)	(0.187)	(0.187)	(0.187)	(0.274)
FHA × Low Equity × Post		−0.800***	−0.746***	−0.560***	−0.510***	−0.816***
		(0.127)	(0.132)	(0.130)	(0.127)	(0.250)
Month FEs	X	X	X	X	X	X
CBSA FEs	X	X	X	X	X	X
FHA Time Trends	X	X	X	X	X	X
Loan Age FEs	X		X	X	X	X
Interest Rate FEs	X		X	X	X	X
LTV × FICO FEs	X		X	X	X	X
ΔUR FEs	X		X	X	X	X
Controls × Post	X			X	X	X
Controls × FHA	X			X	X	X
ΔUR FEs × FHA × Post					X	X
Optimal Refi Subsample						X
Number of Observations	5,441,498	5,441,498	5,441,498	5,441,498	5,441,498	964,318

requirement. Taken together, these results suggest that the restrictions on financing upfront costs for negative-equity borrowers posed a substantial barrier to refinancing during this period.

To check the robustness of these effects, in Panel B of Figure 6, we present results from a more flexible version of the triple-difference regression that allows for the effect to vary by month of observation. These results parallel those presented for unemployment changes in Panel A and are derived from a specification directly analogous to equation (6). In particular, we fully interact the indicator for low equity with month fixed effects and the *FHA* dummy, which allows for the effect of negative equity to vary freely over time. We omit the December 2009 month fixed effect so that all results can be interpreted relative to the month just prior to the policy change. Although the estimates from this specification are noisy, there is a visible level shift down in refinancing that occurs in precisely the month that the policy takes effect. All but one of the coefficient estimates in the pre-period are statistically indistinguishable from zero, whereas the post-period coefficients are all negative, significant, and of roughly the same magnitude as the pooled effects reported in Table VI.

B.4. Optimality versus Liquidity

Our results thus far show that having to pay for upfront costs out-of-pocket reduced refinancing rates for negative-equity FHA borrowers substantially. This decline may be due to two distinct mechanisms. First, there is a long literature arguing for the presence and importance of liquidity constraints (Deaton (1991), Carroll (2001), Zeldes (1989)). These constraints could prevent some borrowers from being able to pay the upfront costs needed to refinance.[36] Second, being forced to pay costs upfront instead of financing them with the loan may change the optimality of the refinancing decision even for borrowers with ample liquidity. This can happen when the subjective discount rate of the borrower differs from the actual interest rate on the loan. In this section, we attempt to distinguish between these two channels.

Although we cannot observe household liquidity or the upfront costs of refinancing directly, we are able to identify a group of households for whom the refinancing decision is still likely to be "optimal" even if they need to pay for closing costs out-of-pocket. Within this set of households, changes in refinancing behavior caused by the policy should be driven primarily by the liquidity effect. Therefore, by comparing the behavior of this group of households to that of the entire sample, we are able to gauge the relative importance of the liquidity effect.

To measure the optimality of the household's refinancing option, we draw on the same approach used in Section III and compute the implied rate gap at which it would be optimal for a borrower to refinance according to the model

---

[36] Borrowers may literally have the cash available to pay the costs, but if the borrower's precautionary motives imply that this level of liquidity would not leave a large enough liquidity buffer, then we would ascribe this failure to refinance to liquidity constraints.

in Agarwal, Driscoll, and Laibson (2013). However, for this exercise we also directly incorporate the role of closing costs. Specifically, whereas before we assumed that all closing costs could be rolled into the new loan when calculating the ADL rate gaps, here we require that borrowers pay a substantial sum out-of-pocket. We then re-calculate the ADL rate gaps and construct an indicator for whether it would be optimal for a borrower to refinance in each month given her current rate, potential rate, and the estimated upfront cost. Because there are no widely available data on closing costs, we follow Agarwal, Driscoll, and Laibson (2013) and calibrate this cost to be $2,000 plus 1% of the loan balance being refinanced.

In column (6) of Table VI, we use this measure to explore the extent to which the effect of forcing borrowers to pay for upfront costs is driven by changes in the optimality of refinancing relative to liquidity. Specifically, we re-estimate the specification from column (5) in the subsample of borrowers for whom it would still be optimal to refinance even if they were required to pay for closing costs out-of-pocket. Changes in refinancing behavior in this sample should be driven primarily by liquidity. The estimated effect in this subsample is still substantially negative and in fact larger than the analogous results for the full sample in column (5). The fact that refinancing rates fall even in this subsample suggests that lack of liquidity is an important driver of the fall in refinancing for negative-equity borrowers following the policy change.

V. Conclusion

Using large changes in the FHA SLR Program, we present evidence that requiring borrowers to document employment and pay upfront costs introduces economically meaningful frictions to mortgage refinancing. This suggests that the pass-through of monetary or other policies that aim to stimulate consumption through the mortgage refinancing channel may be less efficient in recessions, when unemployment is higher and households have less liquidity. Moreover, these frictions are likely to bind most for precisely the households whose expenditures may be most sensitive to reduced rates—those with little cash-on-hand or who recently experienced a negative income shock. This fact may exacerbate the already unequal impacts of recessions by limiting the extent to which reductions in interest rates or other policies that operate through mortgage refinancing benefit lower income households directly. Evaluating the feasibility and welfare impacts of a broader SLR Program that is accessible to conventional or private-label borrowers is well beyond the scope of this paper. However, our results suggest that there are a significant number of borrowers that would refinance their mortgages when lower rates are on offer but who cannot do so because of these large frictions in the mortgage market.

REFERENCES

Adelino, Manuel, Kristopher Gerardi, and Paul S. Willen, 2013, Why don't lenders renegotiate more home mortgages? Redefaults, self-cures and securitization, *Journal of Monetary Economics* 60, 835–853.

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, and Douglas D. Evanoff, 2011, The role of securitization in mortgage renegotiation, *Journal of Financial Economics* 102, 559–578.

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru, 2017a, Policy intervention in debt renegotiation: Evidence from the Home Affordable Modification Program, *Journal of Political Economy* 125, 654–712.

Agarwal, Sumit, Gene Amromin, Souphala Chomsisengphet, Tim Landvoigt, Tomasz Piskorski, Amit Seru, and Vincent Yao, 2017b, Mortgage refinancing, consumer spending, and competition: Evidence from the home affordable refinancing program, NBER Working Paper No. 21512.

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel, 2018, Do banks pass through credit expansions to consumers who want to borrow? *Quarterly Journal of Economics* 133, 129–190.

Agarwal, Sumit, John C. Driscoll, and David I. Laibson, 2013, Optimal mortgage refinancing: A closed-form solution, *Journal of Money, Credit and Banking* 45, 591–622.

Agarwal, Sumit, Richard J. Rosen, and Vincent Yao, 2015, Why do borrowers make mortgage refinancing mistakes? *Management Science* 62, 3494–3509.

Andersen, Steffen, John Y. Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai, 2019, Sources of inaction in household finance: Evidence from the danish mortgage market, NBER Working Paper.

Aragon, Diego, Andrew Caplin, Sumit Chopra, John V. Leahy, Yann LeCun, Marco Scoffier, and Joseph Tracy, 2010, Reassessing FHA risk, Technical Report, National Bureau of Economic Research.

Archer, Wayne R., David C. Ling, and Gary A. McGill, 1996, The effect of income and collateral constraints on residential mortgage terminations, *Regional Science and Urban Economics* 26, 235–261.

Auclert, Adrien, 2017, Monetary policy and the redistribution channel, NBER Working Paper No. 23451.

Benítez-Silva, Hugo, Selçuk Eren, Frank Heiland, and Sergi Jiménez-Martín, 2015, How well do individuals predict the selling prices of their homes, *Journal of Housing Economics* 29, 12–25.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra, 2019, Regional heterogeneity and the refinancing channel of monetary policy, *Quarterly Journal of Economics* 134, 109–183.

Bernanke, Ben S., and Mark Gertler, 1995, Inside the black box: The credit channel of monetary policy transmission, *Journal of Economic Perspectives* 9, 27–48.

Bond, Philip, Ronel Elul, Sharon Garyn-Tal, and David K. Musto, 2017, Does junior inherit? Refinancing and the blocking power of second mortgages, *Review of Financial Studies* 30, 211–244.

Boyce, Alan, Glenn Hubbard, Chris Mayer, and James Witkin, 2012, Streamlined refinancings for up to 14 million borrowers, Working Paper.

Campbell, John Y., 2006, Household finance, *Journal of Finance* 61, 1553–1604.

Campbell, Tim S., and J. Kimball Dietrich, 1983, The determinants of default on insured conventional residential mortgage loans, *Journal of Finance* 38, 1569–1581.

Caplin, Andrew, Charles Freeman, and Joseph Tracy, 1997, Collateral damage: Refinancing constraints and regional recessions, *Journal of Money, Credit and Banking* 29, 496–516.

Carroll, Christopher D., 2001, A theory of the consumption function, with and without liquidity constraints, *Journal of Economic Perspectives* 15, 23–45.

Chan, Sewin, Samuel Dastrup, and Ingrid Gould Ellen, 2016, Do homeowners mark to market? A comparison of self-reported and estimated market home values during the housing boom and bust, *Real Estate Economics* 44, 627–657.

Chang, Yan, and Abdullah Yavas, 2009, Do borrowers make rational choices on points and refinancing? *Real Estate Economics* 37, 635–658.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia, 2017, Innocent bystanders? Monetary policy and inequality in the U.S., *Journal of Monetary Economics* 88, 70–89.

Deaton, Angus, 1991, Saving and liquidity constraints, *Econometrica* 59, 1221–1248.

Deng, Yongheng, and John M. Quigley, 2012, Woodhead behavior and the pricing of residential mortgages, SSRN Working Paper.

Deng, Yongheng, John M. Quigley, and Robert Order, 2000, Mortgage terminations, heterogeneity and the exercise of mortgage options, *Econometrica* 68, 275–307.

Kermami, Maggio, Marco Di Amir, Benjamin Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao, 2017, Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging, *American Economic Review* 107, 3550–3588.

Di Maggio, Marco, Amir Kermani, and Christopher Palmer, 2016, How quantitative easing works: Evidence on the refinancing channel, NBER Working Paper.

Dickinson, Amy, and Andrea Heuson, 1994, Mortgage prepayments: Past and present, *Journal of Real Estate Literature* 2, 11–33.

Dunn, Kenneth B., and Chester S. Spatt, 2005, The effect of refinancing costs and market imperfections on the optimal call strategy and the pricing of debt contracts, *Real Estate Economics* 33, 595–617.

Eberly, Janice, and Arvind Krishnamurthy, 2014, Efficient credit policies in a housing debt crisis, *Brookings Papers on Economic Activity* 2014, 73–136.

Ehrlich, Gabriel, and Jeffrey Perry, 2015, Do large-scale refinancing programs reduce mortgage defaults? Evidence from a regression discontinuity design, CBO Working Paper 2015-06.

Ganong, Peter, and Pascal Noel, 2017, The effect of debt on default and consumption: Evidence from housing policy in the great recession, Working Paper.

Green, Richard K., and Michael LaCour-Little, 1999, Some truths about ostriches: Who doesn't prepay their mortgages and why they don't, *Journal of Housing Economics* 8, 233–248.

Greenwald, Daniel L., 2018, The mortgage credit channel of macroeconomic transmission, SSRN Working Paper.

Haughwout, Andrew, Ebiere Okah, and Joseph Tracy, 2016, Second chances: Subprime mortgage modification and redefault, *Journal of Money, Credit and Banking* 48, 771–793.

Johnson, Eric, Stephan Meier, and Olivier Toubia, 2015, Money left on the kitchen table: Exploring sluggish mortgage refinancing using administrative data, surveys, and field experiments, SSRN Working Paper.

Keys, Benjamin J., Devin G. Pope, and Jaren C. Pope, 2016, Failure to refinance, *Journal of Financial Economics* 122, 482–499.

Lambie-Hanson, Lauren, and Carolina Reid, 2017, Stuck in subprime? Examining the barriers to refinancing mortgage debt, FRB of Philadelphia Working Paper No. 17–39.

Lucas, Deborah, Damien Moore, and Mitchell Remy, 2011, An evaluation of large-scale mortgage refinancing programs, CBO Working Paper 2011–4.

Lusardi, Annamaria, Daniel Schneider, and Peter Tufano, 2011, Financially fragile households: Evidence and implications, *Brookings Papers on Economic Activity* 42, 83–151.

Mayer, Christopher, Edward Morrison, Tomasz Piskorski, and Arpit Gupta, 2014, Mortgage modification and strategic behavior: Evidence from a legal settlement with countrywide, *American Economic Review* 104, 2830–2857.

Pavlov, Andrey D., 2001, Competing risks of mortgage termination: Who refinances, who moves, and who defaults? *Journal of Real Estate Finance and Economics* 23, 185–211.

Piskorski, Tomasz, and Amit Seru, 2018, Mortgage market design: Lessons from the Great Recession, *Brookings Papers on Economic Activity* 49, 429–513.

Scharfstein, David, and Adi Sunderam, 2016, Market power in mortgage lending and the transmission of monetary policy, IMF Working Paper.

Stanton, Richard, 1995, Rational prepayment and the valuation of mortgage-backed securities, *Review of Financial Studies* 8, 677–708.

Woodward, Susan E., 2008, A study of closing costs for FHA mortgages, Technical Report, Department of Housing and Urban Development.

Zeldes, Stephen P., 1989, Consumption and liquidity constraints: An empirical investigation, *Journal of Political Economy* 97, 305–346.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.
Replication code.