Mortgage Market Design: Lessons from the Great Recession

Tomasz Piskorski and Amit Seru*

July, 2018

Abstract

The rigidity of mortgage contracts and a variety of frictions in the design of the market and the intermediation sector hindered efforts to restructure or refinance household debt in the aftermath of the financial crisis. In this paper, we focus on understanding the *design* and *implementation* challenges of ex ante and ex post debt relief solutions that are aimed at a more efficient sharing of aggregate risk between borrowers and lenders. Using a simple framework that builds on the mortgage design literature, we illustrate that ex ante–designed, automatically indexed mortgages and policies can facilitate a quick implementation of debt relief during a crisis. However, the welfare benefits of such solutions are substantially reduced if there are errors in understanding the underlying structure of income and housing risk and their relation to the indexes on which these solutions are based. Empirical evidence reveals significant spatial heterogeneity and the time-varying nature of the distribution of economic conditions, which pose a significant challenge to the effective ex ante design of such solutions. The design of ex post debt relief policies can be more easily fine-tuned to the specific realization of economic risk. However, the presence of various implementation frictions and their spatial heterogeneity can significantly hamper their effectiveness. Consequently, we argue that effective mortgage market design will likely involve a combination of ex ante and ex post debt relief solutions, with state contingencies. We conclude by discussing the potential gains—which can be large, given significant regional heterogeneity—from tying mortgage terms and policies to local indicators, as well as mechanisms that may alleviate the adverse effects of ex post implementation frictions.

* This paper was prepared for the Spring 2018 Edition of Brookings Papers on Economic Activity. Piskorski: Columbia Graduate School of Business and NBER; Seru: Stanford Graduate School of Business, Hoover Institution, SIEPR and NBER. Piskorski and Seru thank the National Science Foundation Award (1628895) on "The Transmission from Households to the Real Economy: Evidence from Mortgage and Consumer Credit Markets" for financial support. We thank Jan Eberly and Jim Stock for guidance and Anthony DeFusco and Paul Willen for helpful comments and discussion. We also thank Erik Hurst, Arvind Krishnamurthy, Karen Pence, and Richard Stanton for useful comments. We thank Monica Clodius, Susan Cherry, Sam Liu, Jonathan Wallen, and Yang Zhang for outstanding research assistance.

1. Introduction

The recent U.S. housing boom saw an unprecedented increase in household mortgage debt (Keys et al. 2013). This buildup of mortgage debt held by vulnerable households has been partly seen as having particularly exacerbated the severity of the aftermath (Mian and Sufi 2009, 2011, 2014b).[1] However, the characteristics of borrowers and loans originated before the crisis is not the only key factor that affected the severity of the housing market downturn during the Great Recession. A series of papers have argued that a number of factors related to the rigidity of contract terms, along with a variety of frictions in the design of the mortgage market and the intermediation sector, hindered efforts to restructure or refinance household debt, exacerbating the foreclosure crisis (Piskorski, Seru, and Vig 2010; Mayer et al. 2014; Di Maggio et al. 2017; Fuster and Willen 2017).

In response, the Federal Reserve altered its monetary policy by lowering short-term interest rates to historic lows. Also, the administration passed two unprecedented, large-scale debt relief programs: the Home Affordable Refinance Program (HARP), which aimed to stimulate mortgage refinancing activity for up to 8 million heavily indebted borrowers; and the Home Affordable Modification Program (HAMP), which aimed to stimulate a mortgage restructuring effort for up to 4 million borrowers at risk of foreclosure. Research suggests that the implementation of the low-interest-rate policy and these debt relief programs had mixed success (Agarwal and others 2017a, 2017b; Di Maggio and others 2017).

What can we learn from extant research for the potential design of more effective debt relief solutions in the future? In this paper, we focus on understanding the *design* and *implementation* challenges of ex ante and ex post debt relief solutions. In doing so, we also analyze the benefits of indexing such solutions to *local* economic conditions relative to *aggregate* indicators. The objective of this paper is to draw on lessons from prior research and provide evidence-based guidance on both these issues.

We start by discussing the literature that documents various frictions that hindered efforts to refinance or restructure mortgages during the Great Recession. The main frictions that have been documented center on (i) contract rigidity, due to which most contracts that were fixed-rate mortgages were locked in at high rates; (ii) equity refinancing constraints, due to which refinancing mortgages was not feasible for many distressed borrowers with insufficient equity; (iii) intermediary organizational constraints, due to which refinancing or debt relief was not passed on to borrowers; (iv) agency conflicts in servicing mortgages that were largely securitized, which

[1] For recent quantitative equilibrium models of housing booms and busts, see Landvoigt, Piazzesi, and Schneider (2015); Kaplan, Mitman, and Violante (2017); Guerrieri and Uhlig (2016); and Favilukis, Ludvingson, and Van Nieuwerburgh (2017). For an alternative view of the reasons behind the housing boom and bust, see Adelino, Schoar, and Severino (2016).

prevented restructuring; (v) a lack of competition in the refinancing market that blunted the extent of the pass-through to borrowers, lowering their incentives to refinance; and (vi) the ex post moral hazard concerns of intermediaries, whereby offering debt relief to distressed borrowers could alter the incentives of many solvent borrowers to continue making payments.

There is a large body of literature showing that these frictions, each in part, might have prevented debt relief from reaching distressed households. Consequently, there is an ongoing debate regarding the reform of the mortgage market to alleviate the impact of such frictions in the future. At the center of this debate are a variety of proposals concerning the redesign of mortgage contracts, as well as future debt relief policies. These proposals start from the premise that the current risk-sharing arrangement between borrowers and lenders in the mortgage market particularly relies on an option to default that can induce a large number of foreclosures during the crisis, with significant associated deadweight losses. In essence, these proposals argue for more efficient risk-sharing between borrowers and lenders to lower the incidence of costly foreclosures and the severity of future housing market downturns (Shiller 2008; Caplin and others 2008; Piskorski and Tchistyi 2011; Campbell 2013; Keys et al. 2013; Mian and Sufi 2014a; Eberly and Krishnamurthy 2014). Because we want to use the lessons from the literature to assess the design of future mortgages and debt relief policies, we start with the theoretical insights from the research on mortgage design. This allows us to think about various economic forces that should be in the consideration set as we make our assessment.

The main collective insights from this work (Piskorski and Tchistyi 2010, 2011, 2017; Eberly and Krishnamurthy 2014; Guren et al. 2017; Greenwald et al. 2018; Campbell et al. 2018) are as follows. In general, contracts or policies that temporarily reduce mortgage payments during recessions can potentially result in significant welfare gains by preventing costly foreclosures and providing consumption-smoothing benefits to households. This is especially the case for borrowers who face more income variability and can afford only a small down payment. To the extent possible, it would therefore be beneficial to design mortgages or debt relief programs that index mortgage payments to measures that capture the state of the local housing and labor markets. This would allow mortgage payments to be lower in states of the world when local labor markets and housing markets experience a downturn. Such indexation programs need to take into account their impact on the market equilibrium, including the incentives of households to borrow and repay their debt. Empirically relevant informational asymmetries and other frictions may limit the set of state-contingent contracts that are sustainable in market equilibrium. Risk aversion and other constraints may also curtail the ability of financial intermediaries to insure the aggregate risk, limiting the effectiveness of state-contingent mortgages or debt relief policies. Finally, contracts or debt relief policies based on other indexes—for example, interest rate indexation, in the case of adjustable-rate mortgages (ARMs)—may perform quite well in providing household debt relief during downturns, as long as such indexes closely co-move with home prices and borrowers' incomes.

Next, we use a simple framework that builds on these insights to illustrate how automatically indexed mortgage contracts or debt relief policies can lead to significant welfare gains for borrowers. The main channel, as mentioned above, is by reducing the debt burden during economic downturns and lowering the incidence of costly foreclosures. Using this framework, we illustrate two points. First, and very intuitively, a mortgage contract or debt relief policy contingent on some index is more efficient if the index is highly correlated with variables capturing relevant (for example, local) economic conditions for borrowers, and if these variables co-move with each other. Second, we show that the benefits of such solutions are substantially reduced if there are errors in understanding the underlying structure of income and housing risk and their relation to the indexes on which such contracts or policies are based.

Although the main insight behind why such contracts or types of debt relief might be efficient seems relatively straightforward, we spend the next section of the paper on explaining the design and implementation challenges of ex ante and ex post debt relief solutions in practice. A key insight of our framework is that successful implementations of ex ante debt relief solutions rely on a correct understanding of the underlying structure of income and housing risk and its relation to the indexes on which such contracts or policies will be based.[2] This observation is also consistent with the quantitative life cycle models of households' decisions, which emphasize the importance of recognizing a specific nature of household risk for an appropriate mortgage contract choice (Campbell and Cocco 2003, 2015).

To better explain this aspect, we analyze simple measures of housing and income risk and their co-movements across time, regions, and borrowers. We document empirical evidence pointing to significant spatial heterogeneity and to the time-varying nature of the distribution of economic conditions. Our spatial analysis starts at the state level and shows that states' local business cycles have quite different frequencies. For instance, using principal component analysis, we find that a national economic factor explains, on average, about 52 percent of the variation in the time series of a state's economic factor, and that this association varies substantially across states. Moreover, though we find that all state economic factors decline sharply during the Great Recession, substantial dispersion remains. Consistent with this observation, we find that the state-level economic variables are on average more correlated with the local economic factor than the national ones. A direct implication of this analysis is that spatial heterogeneity may limit the effectiveness of mortgage contracts or debt relief policies based on the national-level indexes.

[2] In particular, even the best-designed automatically indexed mortgage contract can perform quite poorly ex post if the lenders or policymakers have incorrect understanding of the true distribution of relevant risk.

Next, we zoom in to more granular geographical regions and conduct an analysis at the county level, with variables that both capture the risk of regions and that are available at high frequency. We find that, as within states, there are large spatial variations in delinquency rates and the equity positions of borrowers. At one end, even during the depths of the Great Recession, many counties have sizable housing equity on average and relatively low levels of unemployment and mortgage delinquencies. At the other end, some counties consist of a severely distressed pool of borrowers with depleted home equity.

We also consider the stability of relationships between county-level variables. We find that county-level mortgage default rates are positively related to increases in the unemployment rate and are negatively related to house price growth. This is not surprising, because the extensive empirical literature identifies these two factors as key drivers of mortgage defaults (Foote, Gerardi, and Willen 2008). However, we also find that the strength of these associations varies substantially over time. Moreover, the strength of the relationship between housing and income risk does not appear to be stable over time, pointing to a time-varying distribution of these variables. This evidence is also broadly consistent with research by Hurst et al. (2016) and by Beraja et al. (2017), who argue that regional shocks are an important feature of the U.S. economy and that the regional distribution of housing equity and income varies substantially over time.

Zooming in further, we show that within a county, there is significant heterogeneity at the ZIP code and individual levels. For instance, we find that there is again a large degree of heterogeneity in the distribution of negative equity and defaults in the U.S. population across time. It is particularly important that this evidence also shows that, even during the crisis, there was a large variation among borrowers within counties in delinquency and their home equity positions.

To investigate this issue more formally, we analyze how much variation in local variables—which might be used in ex ante and ex post policies—can be explained by variables at different levels of geographic granularity. The first exercise we undertake is a simple statistical analysis of what fraction of local variation can be explained at various levels of aggregation by considering an upper bound to the informativeness of various economic variables by their level of geographic aggregation. In our analysis, we focus on house prices, combined loan-to-value ratios, debt-to-income ratios, delinquency rates, and foreclosures. We show that explained variation monotonically decreases as we consider coarser geographic areas. For example, the fraction of ZIP code–level mortgage delinquency and foreclosure rates that can be explained by the corresponding county-level variables is, respectively, about 43 and 35 percent. This pattern suggests a large local variation at the ZIP code level that is not captured by county, state, or national data. We also assess the actual association of various national-, county-, and ZIP code–level variables with ZIP code–level delinquency and foreclosure rates and find similar evidence. We also examine the predictability of local housing-related variables with corresponding lagged variables at different levels of geographic aggregation. We again find that predictability worsens as we consider coarser geographic areas.

Next, we ask what the evidence documented above implies for the design of mortgage contracts and debt relief policies. Recall that for solutions such as automatically indexed contracts or debt relief policies to be effective, one needs to have a good ex ante understanding of the underlying distribution of the relevant economic risk and its relation to indexes on which such contracts or policies are based. Given the evidence of significant heterogeneity in space and time, along with limited data on crisis episodes, this can be quite challenging. Moreover, a major change in the nature of mortgage contracts or housing policy is likely to significantly alter market equilibrium, including future joint distribution of such economic outcomes as house prices, housing supply, homeownership rates, and household debt levels (Piskorski and Tchistyi 2017; Guren, Krishnamurthy, and McQuade 2017; Greenwald, Landvoigt, and Van Nieuwerburgh 2018). This further complicates an effective usage of historical data in the design and parametrization of future contracts or policies.[3]

Ex post debt relief policies have the advantage of being more fine-tuned to the specific realization of economic risk, and hence they can alleviate the ex-ante design challenges discussed above. However, various implementation frictions can hamper the effectiveness of ex post solutions. We provide evidence that there is significant spatial heterogeneity of frictions that can differentially affect the pass-through of ex post debt relief policies implemented by financial intermediaries. The presence of such factors and the difficulty of identifying them ex ante pose a significant challenge for implementing effective ex post debt relief policies. For instance, though HARP was largely indexed to the local economic conditions of the borrower, because it was based on the current loan-to-value ratio, it was not as effective as anticipated. In particular, because the implementation was through intermediaries, its effectiveness was hampered by intermediary frictions—such as capacity constraints—and also by market design, such as competition in the refinancing market (Agarwal and others 2017b; Fuster, Lo, and Willen 2017). Similar observations apply to HAMP, which based its eligibility criteria on the current debt-to-income ratio of the borrower, yet performed below its potential, due to the limited ability of intermediaries to conduct loan modifications (Agarwal et al. 2017a).

Finally, our empirical analysis also sheds light on the benefits of indexing ex ante and ex post debt relief solutions to local economic indicators. In particular, our evidence of significant spatial heterogeneity suggests that there might be substantial gains from fine-tuning debt relief solutions to more granular regional conditions and that one-size-fits-all policies might not be that efficient.

[3] See also Rajan, Seru, and Vig 2015, who illustrate that the changed nature of intermediation in the mortgage market (Keys and others 2010; Purnanandam 2011) may alter the stability of statistical relationships between key variables.

For instance, ignoring the heterogeneity in space, though ARM contracts indexed on national interest rate indexes might be helpful during periods of low interest rates, they may also exacerbate distress during periods of higher interest rates, as was the case in the late 2006, early 2008 period. Indexing policies and contracts to variables capturing local components of housing market risk (for example, ZIP code–level house price indexes and other local variables) could be more effective than policies based on national indexes. We note, however, that a full assessment of the relative benefits of such programs also requires a careful consideration of their implementation costs relative to more traditional contracts and policies.

Overall, our paper highlights an important trade-off between design and implementation when thinking about debt relief policies in the future. The pre-crisis designed, automatically indexed mortgage contracts or policies have the advantage of circumventing financial intermediary and other fictions by facilitating a quick implementation of debt relief during economic downturns. However, for such solutions to be cost-effective, lenders, policymakers, and borrowers may need to have a good ex ante understanding of the underlying distribution of the relevant economic risk and its relation to the indexes on which such contracts are based. Given the evidence we have discussed above, this can be challenging.

Ex post debt relief solutions, conversely, have the advantage of being more fine-tuned to the specific realization of economic risk. In other words, unlike pre-crisis designed contracts or policies, ex post policy interventions do not need to rely as much on a good ex ante understanding of the underlying distribution of the relevant economic risk and frictions and their relation to the severity of the crisis. However, ex post policy interventions can also delay debt relief and subject it to various implementation frictions that could hinder their effectiveness.

Consequently, we conclude that effective mortgage market design will likely involve a combination of ex ante and ex post debt relief solutions, with state contingencies. Finally, given our evidence, both types of solutions (ex-ante and ex post) may benefit from conditioning on more granular conditions (regional or individual), as opposed to one-size-fits-all indicators.

2. Frictions to Mortgage Debt Relief: Evidence from the Great Recession

The recent literature has documented how several frictions had an impact on the effectiveness of debt relief, thereby exacerbating the foreclosure crisis. The first such friction relates to mortgage contract rigidity—that is, the notion that most mortgage contracts were fixed-rate mortgages (FRMs) that were locked in at high rates. Di Maggio et al. (2017) and Fuster and Willen (2017) show that as interest rates reached historic lows during the Great Recession, borrowers with certain types of ARMs received automatic debt relief.[4] This experiment is useful for quantifying the effects of debt relief because it was received by every borrower with certain types of ARM contracts, regardless of any other frictions in the market that potentially could have hindered the extent of this debt relief.

In particular, exploiting variation in the timing of rate resets of ARMs during the aftermath of the recent crisis, Di Maggio et al. (2017) find that a sizable decline in mortgage payments (up to 50 percent) induces a significant increase in car purchases (up to 35 percent) and a decline in mortgage defaults. Borrowers with lower incomes and less housing wealth have a significantly higher marginal propensity to consume. Areas with a larger share of ARMs were more responsive to lower interest rates and saw a relative decline in defaults and an increase in house prices, car purchases, and employment. Di Maggio et al. (2017) evidence, along with that of Fuster and Willen (2017), highlights the importance of contract rigidity—that is, rigid FRMs versus flexible contracts, such as ARMs—for understanding the pass-through of debt relief to the real economy during periods of low interest rates.[5]

The next friction that hampers debt relief relates to equity refinancing constraints—that is, the notion that the refinancing of mortgages may not be feasible because many distressed borrowers may not have enough equity to refinance. This friction is particularly important for FRMs, the predominant financial obligation of U.S. households.[6] For such borrowers, automatic debt relief, such as that provided to ARM borrowers, is not feasible. Instead, refinancing constitutes one of the main direct channels through which households can get debt relief from the low-interest-rate environment induced by monetary policy.

Agarwal et al. (2017b) study how this constraint hampered the effectiveness of debt relief by examining the effects of HARP—again, a government program that allowed for the refinancing of insufficiently collateralized agency mortgages with government credit guarantees. The authors find that relaxing the equity constraint for refinancing led more than 3 million borrowers to refinance their loans, and that they experienced more than $3,000 in annual savings on average. Many of these borrowers subsequently increased their purchases of durable goods, such as automobiles, with larger effects among more indebted borrowers. A life cycle model of refinancing quantitatively rationalizes these patterns and produces significant welfare gains for borrowers from

[4] We note that subprime ARM contracts featuring the rate-adjustment floors limited the extent of debt relief received by these borrowers.
[5] This evidence is also consistent with Auclert (2017), who provides a model evaluating the role of redistribution in the transmission mechanism of monetary policy to consumption and predicts that if all U.S. mortgages had adjustable rates, the effect of monetary policy shocks on consumer spending would be significantly higher.
[6] See Green and Wachter (2005) for a discussion of the historical evolution of U.S. mortgage contracts.

relaxing the housing equity eligibility constraint during a crisis.[7] There is, again, spatial heterogeneity in the effects. Regions more exposed to the program—based on the percentage of eligible borrowers in the region—saw a relative increase in consumer spending, a decline in foreclosure rates, and a faster recovery in house prices. This evidence is also consistent with the work of Beraja et al. (2017), who document that before HARP, low interest rates mainly benefited borrowers in regions with relatively high housing equity, exacerbating regional economic inequality (see also Di Maggio, Kermani, and Palmer 2016).

Agarwal et al. (2017b) also illustrate that a lack of competition in the refinancing market blunted the extent of pass-through to borrowers, lowering their incentives to refinance. These frictions reduced the take-up rate among eligible borrowers by 10 to 20 percent and cost borrowers who refinanced their loans between $400 and $800 in annual savings from relief. Strikingly, the largest effects were among the most indebted borrowers—the primary target of HARP—where competitive frictions had the most bite. As before, there was spatial variation in these effects, depending on the degree of competitiveness in the refinancing market. These findings resonate well with those of Scharfstein and Sunderam (2016)—and also with those of Drechsler, Savov, and Schnabl (2017)—who show that the extent of the pass-through of low interest rates in the refinancing and bank deposit market is affected by the degree of competition. They are also broadly connected with the findings of Agarwal et al. (2018)—and of Benmelech, Meisenzahl, and Ramcharan (2017)—who demonstrate the importance of financial intermediaries for the pass-through of interest rate shocks in the credit card and auto loan markets.

Directly restructuring borrower debt through loan renegotiation is another feasible channel for offering debt relief. Despite the surge in distressed borrowers, the U.S. economy experienced limited loan restructuring activity early in the crisis, significantly exacerbating the high number of foreclosures. Research attributes this limited restructuring activity to institutional frictions due to securitization, which prevented renegotiation (Piskorski, Seru, and Vig 2010; Agarwal et al. 2011; Kruger 2017; Maturana 2017) and to lender concerns about strategic defaults, an inability to evaluate the repayment ability of borrowers, and concerns about the adverse impact of wide-scale renegotiations on future repayment incentives (Mayer et al. 2014; Adelino, Gerardi, and Willen 2014). Motivated by such frictions and perceived negative externalities of debt overhang and foreclosures (Campbell, Giglio, and Pathak 2011; Melzer 2017; Gupta 2018), the federal government implemented HAMP. In brief, the program provided substantial financial incentives to financial intermediaries (servicers) for renegotiating loans.

[7] For recent quantitative models emphasizing the importance of refinancing for household consumption, see, among others, Chen, Michaux, and Roussanov (2013); Wong (2018); Greenwald (2018); Beraja and others (2017); and Guren, Krishnamurthy, and McQuade (2017).

Agarwal and others (2017a) study the effects of this program and find that, when employed, the debt relief due to these renegotiations led to a lower rate of delinquencies and foreclosures for borrowers and higher consumer spending and house prices in more exposed regions. Ganong and Noel (2017) further show that temporary mortgage interest rate reductions induced by HAMP played the major role in explaining these effects. Of particular importance, Agarwal and others (2017a) show that the program reached just one-third of the eligible 3 to 4 million indebted households and that there is large heterogeneity across the financial intermediaries in the implementation of debt relief. These differences strongly correlate with banks' organizational design before the program was introduced: Banks that previously had fewer loans per employee, more training for staff, and shorter waiting times for telephone calls took more advantage of HAMP. Because about 75 percent of loans were serviced by banks with a low capability to restructure loans, the program's impact was severely curtailed. Finally, as before, there was significant spatial variation in the implementation of debt relief that relates to the regional share of loans handled by banks with more conducive organization design. These findings also resonate well with those of Fuster et al. (2013) and Fuster, Lo, and Willen (2017), who argue that intermediary capacity constraints had an impact on the extent of the pass-through of debt relief through lower interest rates in the refinancing market.[8]

To summarize, a large body of literature shows that several frictions, each in part, might have prevented debt relief from reaching distressed households, thereby significantly exacerbating the foreclosure crisis. These frictions pertain to both the rigid nature of mortgage designs and to various frictions in the implementation of debt relief policies, including intermediary constraints. Moreover, there is significant regional variation in how much debt relief was passed through to borrowers. We next turn to explaining the key forces that should drive such policies in order to make them more effective. We use insights from the theoretical literature on mortgage design and build a simple illustrative framework.

3. The Mortgage Design Literature and a Simple Framework

A key lesson of the research discussed so far is that the rigidity of mortgage contract terms, along with a variety of other frictions, prevented effective renegotiation or refinancing of distressed borrowers' loans during the recent crisis. Consequently, there is an ongoing debate regarding the reform of the mortgage market to alleviate the impact of such frictions in the future. At the center of this debate are a variety of proposals concerning the redesign of mortgage contracts and debt relief policies that would allow for a more efficient sharing of risk between borrowers and lenders.

[8] We note that the demand-driven factors, such as borrower inertia and inattention, can also limit the extent of interest rate pass-through through mortgage refinancing. For recent evidence on these factors, see Keys, Pope, and Pope (2016); and Andersen et al. (2014).

The hope is that the new mechanisms will lower the incidence of costly foreclosures and the severity of future housing market downturns (Shiller 2008; Caplin et al. 2008; Piskorski and Tchistyi 2011; Campbell 2013; Keys et al. 2013; Mian and Sufi 2014a; Eberly and Krishnamurthy 2014). There are also lessons related to the design and implementation of debt relief policies that require the intermediary sector for implementation. We now discuss implications that emerge from this literature and then use these insights to develop a framework that allows us to highlight the benefits of the automatically indexed mortgage contracts or debt relief policies relative to simple FRMs.

3.1 Implications from Mortgage Design Literature

The debate on the first issue is informed by the growing body of literature that addresses the questions of mortgage contract design and mortgage choice, and their implications for the broader economy. In particular, Piskorski and Tchistyi (2010, 2011) characterize optimal long-term mortgage contracts for borrowers with risky and hard-to-verify incomes in settings with costly foreclosure and stochastic interest rates, house prices, and employment. They show that efficient contracts should generally depend on house price and income indexes in a manner that reduces debt payments during economic downturns.[9] This can be done in a way that does not erode borrowers' incentives to repay their debts. Piskorski and Tchistyi (2010) show that when interest rate indexes are a good measure of a relevant risk ("state of the economy"), the optimal contract takes the form of a flexible payment ARM (the so-called option ARM), whereby the borrower can decide to make only minimum payments with the unpaid interest being added to the principal loan balance until it reaches a certain limit.[10] They also show that such solutions benefit most the borrowers who can afford only a small down payment and face substantial income risk.

These findings underscore the importance of recognizing the interplay between mortgage contracts and the nature of labor income, house prices, and interest rate risk. In this regard, they are related to the research using quantitative life cycle models of mortgage contract choice, such as that by Campbell and Cocco (2003, 2015), which study the implications of such factors for contract choice, consumer welfare, and default patterns.

A number of recent papers extend this literature by studying the implications of state-contingent mortgage contracts in general equilibrium frameworks. Piskorski and Tchistyi (2017) develop a

[9] Such state-contingent contracts could be accompanied by refinancing penalties to enhance longer-term risk-sharing between borrowers and lenders; for analyses of the benefits of such solutions, see Dunn and Spatt (1985) and Mayer, Piskorski, and Tchistyi (2013).

[10] The option to pay less than the minimum monthly interest owed on the loan is valuable for borrowers with fluctuating incomes and provides them effectively with an embedded credit line feature. The fact that the loan is an ARM is valuable, because it reduces the chance of foreclosures when it is relatively more costly (for example, during recessions when interest rates and returns to capital are low).

tractable general equilibrium framework of the housing and mortgage markets with aggregate and idiosyncratic risks, costly liquidity and strategic defaults, empirically relevant informational asymmetries, and an endogenous mortgage design. They focus on the designs that could be sustained in a competitive market equilibrium. They show that though, in general, one would like to index mortgage payments to both labor and housing market conditions, the empirically relevant frictions—including the possibility of strategic defaults discussed in Section 2—may result in equilibrium contracts that only tie mortgage payments to house prices.[11] The adoption of such home equity insurance mortgages would require timely and accurate regional house price indexes. Alternatively, appropriately structured ARM contracts may preserve the benefits of such solutions as long as the interest rate indexes closely co-move with home prices and borrowers' income. Piskorski and Tchistyi (2017) also show that unrestricted competition in mortgage design may lead to the nonexistence of equilibrium in some cases, suggesting a potential role for public policy in implementing new mortgage designs (for example, through subsidies from the government-sponsored enterprises). We come back to this issue in Section 5.

The work discussed above is complemented by recent studies of mortgage contracts in quantitative dynamic equilibrium models of housing markets.[12] In particular, Guren, Krishnamurthy, and McQuade (2017) use a quantitative equilibrium life cycle model with aggregate shocks, long-term mortgages, and an equilibrium housing market, focusing on mortgage designs that index payments to interest rates. They find that the welfare benefits are quantitatively substantial; ARMs improve household welfare relative to FRMs by the equivalent of 1 percent of annual consumption if the central bank lowers interest rates during a bust. Their findings are consistent with research by Di Maggio et al. (2017) and Fuster and Willen (2017), who show that mortgage interest rate declines during the Great Recession due to ARM contracts resetting to a low rate had a positive impact on borrowers and regions exposed to such reductions. Guren et al. (2017) find that an FRM that is convertible to an ARM, a contract similar to the one proposed by Eberly and Krishnamurthy (2014), may perform better than more standard contracts. However, they also point out that an endogenous response by households to such designs can significantly reduce their benefits. In a related paper, Campbell et al. (2018) use a calibrated life-cycle model with competitive risk-averse lenders. Looking at the response of consumption and defaults to given income and house price shocks they find that an ARM contract with an option that during recessions allows borrowers to pay only interest on their loan and extend its maturity has several advantages.

[11] Piskorski and Tchistyi (2017) show that, though beneficial for most borrowers, there are cases when such contracts may decrease the homeownership rate and the welfare of marginal homebuyers.

[12] This line of work is also related to Kung (2015), who explores a number of counterfactuals related to credit availability and mortgage contract forms in a quantitative equilibrium model of the housing \market.

Greenwald, Landvoigt, and Van Nieuwerburgh (2018) study the implications of shared appreciation mortgages that feature mortgage payments that adjust with house prices in a quantitative general equilibrium model with financial intermediaries. They show that if financial intermediaries retain a significant share of mortgages on their balance sheets, the indexation of mortgage payments to aggregate house prices may increase financial fragility, reduce risk-sharing, and lead to expensive financial sector bailouts. In contrast, the indexation to local house prices can reduce financial fragility and improve risk-sharing. The two types of indexation have opposite implications for wealth inequality.

Taken together, a number of key lessons can be derived from this literature. In general, contracts or policies that temporarily reduce mortgage payments during recessions can potentially result in significant welfare gains. To the extent possible, it would be beneficial to index mortgage payments to measures capturing the state of the local labor and housing markets, with mortgage payments being lower in states when these markets experience a downturn. Such indexation programs need to take into account their impact on the market equilibrium, including the incentives of households to borrow and repay their debt. In addition, empirically relevant informational asymmetries and other frictions may limit the set of contracts that are sustainable in equilibrium. Risk aversion and other constraints may also limit the ability of financial intermediaries to insure against the aggregate risk, limiting the effectiveness of state-contingent mortgages or debt relief policies. Finally, contracts or policies based on indexes not directly tied to the housing or labor markets (for example, interest rate indexation, in the case of ARMs) may perform quite well in providing debt relief during downturns, as long as such indexes closely co-move with home prices and borrowers' incomes.

This discussion implies that one of the fundamental requirements for the successful implementation of new mortgage designs or debt relief policies is a thorough understanding of the underlying structure of the economic risk faced by borrowers. Moreover, one needs to understand how the relevant risk relates to a variety of possible indexes that can be used in the design of mortgage contracts or debt relief policies in practice. In the next subsection, we illustrate the importance of these factors in a simple, stylized, illustrative framework. In Section 4, we then provide empirical evidence on these issues as they relate to the actual design of mortgage contracts and debt relief.

3.2 A Simple Illustrative Framework

3.2.1 Setup

We now discuss a simple illustrative framework that draws on insights from the literature we discussed above and will allow us to highlight the benefits of the automatically indexed mortgage contracts or debt relief policies relative to simple FRMs. The key benefit of the indexed mortgage contracts in our setting is that they can reduce the incidence of costly foreclosures due to their state-contingent repayment rates without eroding lenders' ability to break even on their loans. We use this framework to explore two issues. First, we illustrate, through a few numerical examples, how the benefits of such indexed contracts or policies relate to the type of index used by lenders or policymakers and its relation to the underlying structure of economic risk. Second, we investigate how the benefits of such solutions change if there are errors in understanding the underlying structure of income and housing risk and their relation to the indexes on which such contracts or policies are based. Although our framework has a number of important limitations, we believe that the key insights we develop here will also be applicable in much richer settings.[13]

We consider a simple, stylized, partial equilibrium mortgage lending framework, where a risk-neutral borrower with linear utility buys a home worth P_0 by borrowing $P_0 - D$ from a risk-neutral lender; hence, the down payment is D. If $D = 0$, then the borrower pays zero down payment. The borrower can down-pay D equal to his or her initial personal wealth, W_0, upon buying the house. For simplicity, we normalize the discount factor and risk-free rate to be 1. We first consider an FRM, the most commonly used residential mortgage contract in the United States. Under the terms of an FRM, the borrower faces a fixed mortgage interest rate of $\bar{r}$.

The borrower derives utility of θ from living in the home. During the next period, after the loan is made, the borrower realizes his or her income y drawn from a normal distribution f^y, with $y \sim N\left(\bar{y}, \sigma_y^2\right)$. Furthermore, he or she sees the updated home price P_1 drawn from a normal distribution f^P, with $P_1 \sim N(\bar{P}, \sigma_P^2)$. If the borrower sells his or her home at P_1 or defaults, he or she loses θ of utility. If the borrower defaults, the lender receives only $\delta \in (0,1)$ of P_1, where δ captures some liquidation costs and the borrower suffers a utility cost of $\bar{v}$. We further assume that $\theta + \bar{v} > (1 + \bar{r})P_0$, implying that the borrower has an incentive to repay his or her debt. Given this setting, the borrower's optimal strategy can be described as follows:

(1) If realized income is such that $y < (1 + \bar{r})(P_0 - D)$, and the realized house price is such that $P_1 < (1 + \bar{r})(P_0 - D)$, then the borrower has no choice but to default. His or her realized lifetime utility will be $u(y, P_1) = y - \bar{v} - D$.

[13] Notably, among others, (i) we restrict the contract choice to a simple linear rule as a function of a given index; (ii) we only focus on liquidity-driven defaults, neglecting strategic defaults that also accounted for a substantial amount of defaults during the Great Recession; (iii) we do not incorporate empirically relevant informational asymmetries between borrowers and lenders; (iv) we do not model the long-term aspect of mortgage contracts and the possibility of loan refinancing; (v) we do not analyze the impact of borrower and lender risk aversion on consumer welfare and mortgage terms; (vi) we do not take into account general equilibrium effects of changes in contract terms, including the impact of indexation on house prices; and (vii) we set aside the question of what mortgage contracts would be sustainable in the competitive market equilibrium with relevant frictions and whether there is a scope of welfare-improving public policy intervention in such settings. The literature discussed in subsection 3.1 addresses mortgage contract design and its implications, capturing many such factors and complications.

(2) If realized $y < (1+\bar{r})(P_0 - D)$ and $P_1 \geq (1+\bar{r})(P_0 - D)$, then the borrower cannot repay the loan but can sell the home. His or her realized lifetime utility will be $u(y, P_1) = y + P_1 - (1+\bar{r})(P_0 - D) - D$.
(3) If realized $y \geq (1+\bar{r})(P_0 - D)$ and $\theta \geq P_1$, then the borrower repays the loan without selling the house. His or her realized lifetime utility will be $u(y, P_1) = y + \theta - (1+\bar{r})(P_0 - D) - D$.
(4) If realized $y \geq (1+\bar{r})(P_0 - D)$ and $\theta < P_1$, then the borrower sells the home, and his or her realized lifetime utility will be $u(y, P_1) = y + P_1 - (1+\bar{r})(P_0 - D) - D$.

We note that a default occurs if both house prices and income are sufficiently low, consistent with the "double trigger" notion in the literature (Foote, Gerardi, and Willen 2008). The competitive FRM mortgage interest rate will be the lowest $\bar{r}$—because the lower is $\bar{r}$, the higher is the borrower's utility—such that the lender breaks even.

Formally, we formulate this problem as follows. First, we define the distribution of income and house price as follows:

$$\mathbf{X} = \begin{bmatrix} y \\ P_1 \end{bmatrix}, \mu = \begin{bmatrix} \bar{y} \\ \bar{P} \end{bmatrix}, \mathbf{\Sigma} = \begin{bmatrix} \sigma_y^2 & \rho_{yP}\sigma_y\sigma_P \\ \rho_{yP}\sigma_y\sigma_P & \sigma_P^2 \end{bmatrix}, \text{ and } \mathbf{X} \sim N(\mu, \mathbf{\Sigma}).$$

Given the above discussion, under the FRM contract, the consumer's expected utility maximization problem, subject to the lender's break-even condition, can be formulated as a function of defaulting, selling, and paying states:

$$\begin{aligned} \max_{\bar{r}} \ & \Pr_{\text{def}} \times E(y - \bar{v}|\text{def}) + \Pr_{\text{sel}} \times E[y + P_1 - (1+\bar{r})(P_0 - D)|\text{sel}] + \Pr_{\text{pay}} \\ & \times E[y + \theta - (1+\bar{r})(P_0 - D)|\text{pay}] - D \\ \text{s.t.} \quad & P_0 - D = \Pr_{\text{def}} \times E(\delta P_1|\text{def}) + \Pr_{\text{sel}} \times [(1+\bar{r})(P_0 - D)] + \Pr_{\text{pay}} \times [(1+\bar{r})(P_0 - D)], \end{aligned}$$

where we define the probabilities given above as follows:

$$\begin{aligned} \Pr_{\text{def}} &= \Pr[y < (1+\bar{r})(P_0 - D), P_1 < (1+\bar{r})(P_0 - D)], \\ \Pr_{\text{sel}} &= \Pr[y < (1+\bar{r})(P_0 - D), P_1 \geq (1+\bar{r})(P_0 - D)] + \Pr[y \geq (1+\bar{r})(P_0 - D), \theta < P_1], \\ \Pr_{\text{pay}} &= \Pr[y \geq (1+\bar{r})(P_0 - D), \theta \geq P_1]. \end{aligned}$$

In the calculations given above, we assume that the borrower uses all his or her initial wealth for a down payment. It is worth noting that in our simple, stylized setting, the borrower will generally have an incentive to down-pay as much as possible because this reduces the expected mortgage cost, which is weakly higher than the riskless saving rate. In a later discussion, we focus on two particular cases: (i) a zero down payment ($D = 0$); and (ii) a 20 percent down payment ($D = 0.2P_0$). The former case is meant to represent highly indebted borrowers with very little initial housing equity, and the latter represents more creditworthy prime borrowers who can afford a substantial down payment.

We next consider an indexed-rate mortgage (IRM) contract of the form $r = \alpha_0 + \alpha_1 i$, where i is an index drawn from a standard normal distribution f^i, with $i \sim N(0,1)$. Hence, $r \sim N(\alpha_0, \alpha_1^2)$, and the overall distribution of stochastic variables is defined as follows:

$$\mathbf{X} = \begin{bmatrix} y \\ P_1 \\ r \end{bmatrix}, \mu = \begin{bmatrix} \bar{y} \\ \bar{P} \\ \alpha_0 \end{bmatrix}, \mathbf{\Sigma} = \begin{bmatrix} \sigma_y^2 & \rho_{yP}\sigma_y\sigma_P & \rho_{yi}\sigma_y\alpha_1 \\ \rho_{yP}\sigma_y\sigma_P & \sigma_p^2 & \rho_{Pi}\sigma_P\alpha_1 \\ \rho_{yi}\sigma_y\alpha_1 & \rho_{Pi}\sigma_P\alpha_1 & \alpha_1^2 \end{bmatrix}, \text{and } \mathbf{X} \sim N(\boldsymbol{\mu}, \mathbf{\Sigma}).$$

We note that the borrower's optimal behavior and lifetime realized utility are the same as described above for the FRM contract, except replacing $\bar{r}$ with the realization of $r \sim N(\alpha_0, \alpha_1^2)$.

In offering this contract, lenders optimally choose the parameters α_0 and α_1, while taking the distribution of the index as given. For example, we could think of an ARM contract as a special case of an IRM contract, where the index i is just some spread over realization of the interest rate index (for example, that of 1-year Treasuries or the London Interbank Offered Rate). We could also think of the IRM as representing a state-contingent debt relief policy that depends on the policy index i coupled with simpler contracts (for example, an FRM).[14]

We further assume that the introduction of IRM contracts may be subject to a certain up-front fixed cost c per borrower that is faced by lenders relative to a setting with FRM contracts. This cost represents some additional unmodeled cost of issuing more complex contracts or implementing a debt relief policy—such as the potential costs of educating borrowers, costs of unmodeled uncertainty about the actual distribution of the index, some additional hedging costs for the lender, or some administrative costs of implementing a debt relief policy.

Given the above-noted setup under any particular correlation schedule ρ_{yP}, ρ_{yi}, and ρ_{Pi}, the competitive equilibrium IRM contract maximizes the consumer's expected utility across the three states, subject to lender break-even condition:

$$\begin{aligned} \max_{\alpha_0,\alpha_1} \quad & \mathrm{Pr}_{\mathrm{def}} \times E(y - \bar{v}|\mathrm{def}) + \mathrm{Pr}_{\mathrm{sel}} \times E[y + P_1 - (1+r)(P_0 - D)|\mathrm{sel}] + \mathrm{Pr}_{\mathrm{pay}} \\ & \times E[y + \theta - (1+r)(P_0 - D)|\mathrm{pay}] - D \\ \text{s.t.} \quad & P_0 - D = \mathrm{Pr}_{\mathrm{def}} \times E(\delta P_1|\mathrm{def}) + \mathrm{Pr}_{\mathrm{sel}} \times E[(1+r)(P_0 - D)|\mathrm{sel}] + \mathrm{Pr}_{\mathrm{pay}} \\ & \times E[(1+r)(P_0 - D)|\mathrm{pay}] - c, \end{aligned}$$

where

[14] Implementation of such a debt relief policy with simpler contracts may require an ex ante commitment from policymakers, lenders, and borrowers.

$$\mathrm{Pr}_{\mathrm{def}} = \Pr[y < (1+r)(P_0 - D), P_1 < (1+r)(P_0 - D)],$$
$$\mathrm{Pr}_{\mathrm{sel}} = \Pr[y < (1+r)(P_0 - D), P_1 \geq (1+r)(P_0 - D)] + \Pr[y \geq (1+r)(P_0 - D), \theta < P_1],$$
$$\mathrm{Pr}_{\mathrm{pay}} = \Pr[y \geq (1+r)(P_0 - D), \theta \geq P_1],$$
$$r = \alpha_0 + \alpha_1 i.$$

This problem has no closed-form solution, so to gain insights, we focus on numerical solutions for a set of parameters given in Table 1. We note that the main insights from our illustrative framework are valid across a wide range of parameters.

It is worth noting, as will become clear shortly, that if the additional cost of issuing IRM contracts is equal to zero, the IRM loans will always be weakly better for borrowers than the FRMs. The reason is that the IRM contracts nest the FRM ones. As we illustrate below, with the positive fixed cost of issuing a more complex mortgage, whether such a mortgage will be better than an FRM depends on how closely i, y, and P_1 co-move with each other.

3.2.2 Benefits of Mortgage Debt Indexation

First, let us consider the case of no-fixed-cost index mortgages compared with FRMs. We start by showing the borrower's utility gain (in percentage terms) under an IRM compared with an FRM, assuming that (P, y) are perfectly correlated. The top panel of Figure 1 plots this result: On the horizontal axis, we have varying degrees of correlation between i with y. Because y and P are perfectly correlated, this is also the correlation between i with P.

The top panel of Figure 1 shows an important feature of our setting: An IRM without a fixed issuing cost would never do worse than an FRM, *provided that the lenders correctly assess the distribution of the underlying risk*. The reason is that by optimally choosing $r = \alpha_0 + \alpha_1 i$ in the contract, one can always reduce α_1 to zero when the index correlation with y or P is approaching zero. In this sense, the FRM would simply be a special case of an IRM. As soon as the index correlation with y or P turns positive, there is always some benefit from reducing the default probability. Hence, the optimal contract would also have $\alpha_1 > 0$, turning on the volatility of the index mortgage interest rate. Therefore, as is evident from Figure 1, the benefit of an index mortgage contract for avoiding costly foreclosures generally is larger when the correlation between the index and income or the house price is higher. We also note that in a setting with borrower risk aversion, state-contingent lending contracts may provide additional benefits to households by partially insuring their labor income risk and hence allowing them to better smooth their consumption profiles. This additional benefit should increase the value of state-contingent contracts relative to FRMs.

In reality, house prices and household incomes are not perfectly correlated. To explain how our insights might change due to this, we next consider two cases: (i) Corr(y, P) = .25 (low correlation);

and (ii) Corr(y, P) = .75 (high correlation). The results are shown in the top panels of Figure 2. The following results emerge: (i) An IRM is never worse than an FRM, because an FRM is a special case of the IRM contract when $\alpha_1 = 0$; (ii) generally, the higher are Corr(P, i) and Corr(y, i), the larger is the gain from an indexed loan relative to an FRM; and finally, (iii) the utility gains under IRMs are generally higher, given that Corr(P, y) is higher.

Next, we take into account the possibility of a down payment for a house purchase. As formulated in the model, we consider the case of a 20 percent down payment in both an FRM and IRM. The bottom panels of Figure 2 show the corresponding results. We see that for the case of a 20 percent down payment with significant positive home equity, the gain from indexed contracts is smaller. This is intuitive, because the down payment lowers the default probability and the associated deadweight losses from having a rigid contract.

Now we consider the case where issuing an IRM has a fixed cost for the lender—in particular, 1 percent of the initial house price. Again, we start by showing the utility gain (or loss) of an IRM compared with an FRM, assuming that house prices and incomes are perfectly correlated. The bottom panel of Figure 1 shows these results. Compared with the top panel, the bottom panel shows that for our parameters with a fixed cost of issuing an indexed loan, there is a range of correlations where utility under the indexed loan is lower than under the FRM. In general, this plot indicates that with the additional cost of issuing an IRM loan, there may be a range of correlations where utility under the indexed loan may be lower than under the FRM.

To shed more light on this issue, Figure 3 reproduces the analysis in Figure 2, but with an additional cost of indexation equal to 1 percent of the initial house price per borrower. This Figure consistently shows that an IRM contract is more likely to benefit consumers when the index correlation with income and house price is sufficiently high. When the index correlation with income and house price is not sufficient, there can be a utility loss compared with an FRM, due to the IRM's issuing cost.

Our simple framework shows that a successful implementation of indexed mortgages crucially relies on a correct understanding of the underlying structure of income and housing risk and its relation to the indexes on which such contracts or policies will be based. To illustrate this point, Figure 4 shows the borrower's utility (in percentage terms) under an IRM designed for an incorrectly projected high correlation between income and house prices (equal to .75) and a high projected correlation between the index and income and house prices (equal to .60). These are compared with scenarios of indexed mortgages that are correctly designed knowing that the actual correlation between income and house prices is low (equal to .25) and that the actual correlation between the index and income and house prices is as shown in Figure 4. The computation assumes no down payment and no indexation cost. As we observe, incorrect beliefs about the distribution of key economic variables result in a substantial decline in efficiency relative to a contract designed under the correct distribution of economic variables. This is because there are instances when the borrower faces a substantial increase in the interest rate during periods of low income and house prices, which increases the risk of a costly foreclosure. We note that more elaborate indexation programs—such as an FRM with an option to be converted to an ARM, a contract proposed by Eberly and Krishnamurthy (2014)—could partly alleviate the impact of such ex ante design errors.

Based on our numerical results, presented in Figures 1 through 4, we summarize the main insights of our simple framework:

(1) Without the additional cost of indexation, an IRM contract is always weakly better than an FRM contract.

(2) The higher the correlation between income and house prices with the index, the bigger the gain from an indexed loan relative to an FRM.

(3) The utility gains under an indexed loan or an indexed debt relief policy are generally higher when the correlation between house prices and income is higher.

(4) With the additional cost of issuing an indexed loan, there is a range of correlations where utility under the indexed loan is lower than under the FRM. Besides, it is possible that when the index is sufficiently correlated with income and house price, the indexed contract is better than the FRM.

(5) Gains from indexed contracts are much higher for borrowers who make only a small or no down payment (and thus have little housing equity).

(6) Benefits of indexed mortgages or debt relief policies crucially depend on a correct understanding of the underlying structure of income and housing risk and its relation to the indexes on which such contracts or policies will be based. In the case of incorrect beliefs about these relationships, the benefits of such solutions can decrease substantially.

Our simple framework highlights the importance of understanding the underlying structure of income and housing risk and its relation to the indexes on which contracts or debt relief policies will be based. More broadly, this includes an assessment of the expected degree of heterogeneity across regions and borrowers, the stability of such relations over time, and the relative value of policies based on national versus local indexes.

4: Spatial and Individual Variation in Income and Housing Risk

In this section, we analyze the structure of income and house price risk across regions and to assess their relation to mortgage defaults and the home equity positions of borrowers. We also discuss how this risk relates to possible indexes that could be used in future mortgage contracts or debt relief policies.

4.1 Evidence from U.S. States

To measure local economic conditions, we take a stance on variables that summarize business conditions. These variables include real GDP growth, personal income growth, unemployment, and house price growth. Real GDP growth measures the output of the economic area. Real personal income growth measures changes in the wealth of local consumers. Both GDP and income data are from the U.S. Bureau of Economic Analysis. We deflate using the CPI-U from FRED (a database maintained by the Federal Reserve Bank of St. Louis). Because unemployment is a permanent loss to income, we include the local unemployment rate. Unemployment data are from the U.S Bureau of Labor Statistics. To measure expectations about future economic conditions, we use changes in the market value of real estate. When available, we use data from Zillow; otherwise, we use data from the Freddie Mac House Price Index. For national housing data, we use the S&P/Case–Shiller U.S. National Home Price Index. Table 2 displays summary statistics for the national- and state-level economic series.

We assume that the local business cycle influences output, income, unemployment, and house prices. For each state, we extract this common component through a principal component analysis. The first component explains, on average, 60 percent of the variation in these four series. This component loads positively on output, income, and house prices, but negatively on unemployment. Table 3 displays the summary statistics for the weights of the first component and its explained variation. The large explained variation and loadings are consistent with a proxy for local economic conditions. Figure 5 plots its mean and the 10th–90th percentile range over time. Note that the economic factor declines sharply during the financial crisis of 2008, but dispersion remains rather stable.

To further characterize the cross-sectional heterogeneity of state-level economic conditions, we regress the local economic factor on a constant and the national economic factor. The data cover all 50 states and the District of Columbia from 1980 to 2016. A national economic factor explains, on average, 52 percent of the variation in the time series of a state. However, this explanatory power varies substantially across states.[15] Local economic conditions in Alaska are least represented by the national economic factor, with an R^2 near 0, whereas Minnesota is most represented, with 82 percent of the variation explained. Online appendix Figure A1 illustrates the distribution of R^2. The fraction of variation explained is closely related to the correlation between local and national economic conditions. The heterogeneity in correlation is also illustrated by variation in the sensitivity of local economic conditions to that of the nation. An improvement in national economic conditions of 1 standard deviation, on average, improves local economic conditions by 0.70 standard deviation. However, this varies substantially, as is illustrated in online appendix Figure A1. For example, North Dakota has a beta of 0.08, while California has a beta of 0.94.

Other macroeconomic variables perform similarly in explaining the variation in state-level business cycles. We consider the underlying macroeconomic variables to the national factor (GDP growth, income growth, house price growth, and unemployment), macroprudential policy rates (the federal funds rate), interest rates (nominal and real 1-year Treasury rates), and the 30-year mortgage interest rate. The federal funds rates, Treasury rates, and mortgage interest rates are sourced from FRED. For each state, we regress local economic conditions on a constant and the underlying macroeconomic variable iteratively. All these national-level macroeconomic variables differ substantially in explanatory power and betas across states. Table 4 provides summary statistics detailing the variation.

Using local economic variables to explain local business conditions is both intuitive and more effective. For all economic series, the state-specific series are, on average, more correlated with the local economic factor. The average correlation between the state-level change in unemployment and the local economic factor is –.68, but the correlation between the national change in unemployment is, on average, –.50. For house prices, there is also a large gain: .67 for state–state correlation and .58 for state–national correlation. Figure 6 illustrates the cross-sectional distribution of correlations between state economic conditions and state economic variables (top left panel) and state economic conditions and national economic variables (top right panel). Notably, the distributions tend to be shifted toward 1 for real GDP growth, income growth, and house price growth, and toward –1 for the unemployment rate. Finally, the bottom panel of Figure 6 shows the substantial heterogeneity in correlations between changes in state economic conditions and national-level interest rate indexes.

Overall, this simple analysis illustrates that local economic conditions exhibit substantial heterogeneity, which is not that closely related to national macroeconomic conditions or interest

[15] This national economic factor is constructed similarly to the state-level economic factors. The national economic factor is the first component of a principal component analysis on real GDP growth, income growth, house price growth, and unemployment.

rate indexes. Furthermore, state-level economic conditions vary in their correlation and sensitivity to national conditions. This regional heterogeneity may limit the ability of national macroprudential policy or mortgage contracts based on national-level indexes to comprehensively and effectively respond to local economic conditions.

4.2 Evidence from U.S. Counties and Zip Codes

So far, we have shown that states exhibited heterogeneous business cycles from 1980 to 2016. Now we turn to the county level to show that counties also experience substantial heterogeneity. Our data come from a variety of sources. The county unemployment rate is from the U.S. Bureau of Labor Statistics, county income is from the U.S. Census Bureau, and county house prices come from Zillow's Home Value Index. We complement the county-level data with additional housing variables. County first mortgage serious delinquency rates and combined loan-to-value ratios (CLTVs) come from a 10 percent representative sample of the U.S. population provided by Equifax, covering the sample period 2005–16.[16] For each county, we focus on local economic variables—unemployment rate, change in unemployment rate, real income growth—and housing variables—house prices, CLTVs, and mortgage delinquency rates. We also complement our analysis by presenting evidence on foreclosure rates, VantageScores, and debt--to-income ratios (all from Equifax data).[17]

We begin by examining the means and standard deviations of real income growth and the unemployment rate. The top panels of Figure 7 show the means of these variables. Unsurprisingly, there are a sharp decrease in mean income growth and a sharp increase in mean unemployment in about 2008. Even more important, the bottom panels of Figure 7 show the standard deviations of both variables. There is considerable variation in both income and unemployment across counties for the entire time series, with spikes at 2008. Although the standard deviation of unemployment begins to decrease after 2010, the standard deviation of income growth remains at the elevated level.

Next, we examine the mean and standard deviation of housing variables. The left panels of Figure 8 show the means of house price growth, CLTVs, and delinquency rates. Again, the means vary substantially over time, with CLTVs and delinquency rates reaching their maxima in about 2010

[16] We note that the Equifax data we use do not have a direct measure of current CLTVs of mortgage borrowers. We compute this variable in a region (county or ZIP code) by diving the average combined mortgage debt level of borrowers with first mortgages on their credit files by the median house price in a region (from Zillow). We verified that our measure of average CLTV in a region is closely related to the CLTV measure from widely used Credit Risk Insight Servicing McDash (CRISM) data that cover approximately 70 percent of mortgage borrowers. We also note that our measure indicates slightly higher CLTV levels than do the CRISM data, likely due to the well-known underrepresentation of subprime borrowers in the CRISM data; see online appendix figure A2 for more details.

[17] The Equifax-based DTI should be interpreted with caution because the Equifax data do not report the actual income of the borrower and instead provide the estimated income based on credit variables.

and 2011, and with the house price index growth reaching its minimum in 2010 (see also Mayer et al. 2009). As the right panels of Figure 8 show, the standard deviations also fluctuate throughout the time series, with the volatility of all three variables reaching a peak during the period from 2009 to 2011. Overall, Figures 7 and 8 show that both the mean values of county variables and the variability of these values across counties vary significantly over time.

Another way to view heterogeneity spatially is by presenting heat maps of county-level variables before, during, and after the financial crisis. Figure 9 does so by plotting the unemployment rate, while Figure 10 plots house price growth. The top panels illustrate that even before the recession, there was some heterogeneity across counties. We can see from the middle panels that heterogeneity increased during the crisis. And the bottom panels show that most counties recover across these two variables, but some remain in a distressed state. These two figures illustrate the extent of the heterogeneity across counties in various periods across income and house price risk. This evidence is also consistent with the urban economics literature, which documents significant heterogeneity in local house price movements (Glaeser, Gyourko, and Saiz 2008; Sinai 2013).

Figure 11 similarly plots the heat map with CLTVs and delinquency, in 2010. We note that areas with high CLTV levels often correspond to the areas that experienced high house price growth before the crisis (see the top panel of Figure 10). This reflects, in part, a significant amount of home equity extraction in areas that experience rapid house price growth before the bust (Mian and Sufi 2011; Bhutta and Keys 2016). Figures 10 and 11 suggest that the heterogeneity in unemployment and house price growth implies a significant heterogeneity in housing equity and mortgage defaults during the peak of the Great Recession. Many counties have high CLTVs, delinquency rates, and unemployment rates and low house price growth in 2010, but other counties continue to perform quite well. This evidence is consistent with the work of Mian and Sufi (2014b), who show a strong link between household leverage and the extent of house price declines at the regional level, and the subsequent increase in unemployment during the Great Recession. At the same time, though many counties have high CLTVs, delinquency rates, and unemployment rates and low house price growth in 2010, other counties continue to perform quite well.

This heterogeneity exists across years, but especially so during the financial crisis. Figures 12 and 13 show similar evidence for U.S. ZIP codes.[18] Online appendix Figure A3 complements this evidence by showing similar heterogeneity in foreclosure rates, debt-to-income ratios, and VantageScores. Strikingly, at this more granular level, the evidence of the heterogeneity becomes even more pronounced. Overall, this evidence indicates that the Great Recession did not affect

[18] Our analysis of heterogeneity at the ZIP code level is limited because we do not have access to good unemployment data at this level.

regions uniformly, and that there is a substantial heterogeneity in housing equity and default that is also visible in the heterogeneity of unemployment and house price movements.

Thus far, we have visually examined heterogeneity in space and time through means and standard deviations. Next, we consider the stability of relationships between county-level variables. We regress the dependent variable on the independent variable interacted with annual dummy variables for each year. In Figure 14, we show the coefficients of such regressions, where we regress the change in the mortgage default rate on the change in unemployment rate (left panel) and on house price growth (right panel), respectively. Both panels include 95 percent confidence intervals.

Figure 14 confirms that the extent of mortgage defaults in a region is closely associated with changes in unemployment rates and house prices, with mortgage defaults being generally lower in areas experiencing lower levels of unemployment and higher house price growth. This is not surprising because the extensive empirical literature identifies these two factors as key drivers of mortgage defaults (Foote, Gerardi, and Willen 2008; Keys et al. 2013). This is also consistent with the predictions of our simple illustrative framework presented in subsection 3.2. Interestingly, though the relationships between these variables are quite strong, the strength of these relationships also varies over time. For example, the regression of mortgage defaults on unemployment rates is positive throughout the entire time series, but varies substantially (the left panel of Figure 14).

Figure 15 sheds additional light on this question by examining the stability of the relationship between house price growth and the change in unemployment (left panel) and the change in CLTVs (adversely related to the change in housing equity) and the change in the unemployment rate (right panel). The evidence points to significant instability between these two key drivers of mortgage defaults. In other words, it appears that it is not always the case that regions experiencing a substantial increase in house prices (or housing equity) also experience a substantial *simultaneous* decrease in unemployment. For example, the regression of house price growth on the change in unemployment results in a strong negative relationship for most of the time series; but the strength of the relationship decreases from 2010 onward, and we even find positive results from 2013 to 2015. This evidence is also broadly consistent with that of Hurst et al. (2016) and Beraja et al. (2017), who show that regional shocks are an important feature of the U.S. economy and that the regional distribution of housing equity and income varies over time.

4.3 Relative Importance of Local Economic Indicators

The evidence discussed above suggests that local and regional economic conditions display considerable heterogeneity that is related to the state of the housing market. This suggests that indexing mortgage contract terms or debt relief policies to indexes capturing the local component of economic conditions may improve the efficiency of such solutions. To shed more light on this issue, we now more formally assess the association of various national-, county-, and ZIP code–level variables. In particular, we analyze how much variation in local variables—which might be used in ex ante and ex post policies—can be explained by variables at different levels of geographic granularity. As will become clear, doing so allows us to better explain how to assess and predict these local variables.

The first exercise we undertake is a simple statistical analysis of what fraction of local variation can be explained at various levels of aggregation. In our analysis, we will focus on five variables that we discussed earlier in the paper: house prices, CLTVs, debt-to-income ratios (DTIs), delinquency rates, and foreclosures. We analyze several levels of geographic granularity. At the most granular level, a local housing market is here defined by its ZIP code (out of 14,250 ZIP codes). Similarly, we also assess geographic granularity at the level of the city (7,600), county (1,000), metropolitan area (730), state (51), and nation (1). Our ZIP code analysis uses a sample that spans January 1997 to December 2017 for house prices and July 2005 to December 2017 for CLTVs, DTIs, delinquency rates, and foreclosures. We focus on variations in both the levels of these variables and their growth rates. Additionally, we demean the series by ZIP code to absorb time-invariant, cross-sectional heterogeneity. To make the analysis robust to outliers, we winsorize the tails at the 1 percent level. Online appendix Table A4 reports the summary statistics for the housing market variables on which we focus.

To characterize the importance of local economic variables in capturing the state of the local housing market, we estimate the fraction of variation that can be explained by geography-time fixed effects. Formally, we regress

$$Y_{i,t} = \sum_{\tau=1}^{T} \beta_{j(i),\tau} Geography_{j(i)} \times d_{\tau} + \varepsilon_{i,t} \ ,$$

where $Y_{i,t}$ is the housing market variable for ZIP code i in period t. In particular, the housing market variables we consider include real house prices, CLTVs, DTIs, delinquency rates, and foreclosures, as discussed above in the data section. The term d_{τ} is a time dummy variable taking the value of 1 when $\tau = t$ and 0 otherwise, and we measure $Geography_{j(i)}$ by different levels of aggregation: city, county, metropolitan area, state, and national levels. Hence, $\sum_{\tau=1}^{T} Geography_{j(i)} \times d_{\tau}$ is a series of geography-time fixed effects.

The setup for this regression puts an upper bound on the variation of $Y_{i,t}$ that can be explained by any economic variable at aggregation level j. For example, suppose we included ZIP code–time fixed effects. This specification soaks up all the variation in $Y_{i,t}$, yielding an R^2 of 1. Clearly, the upper bound on explanatory power is agnostic about which underlying economic variable explains variation in a given housing market variable. This effectively estimates an upper bound to R^2 that could be generated from a contemporaneously measured variable at the level of aggregation j. The regression results are reported in Table 5.

The top panel of Table 5 show the results for the variation in the levels of housing variables. It document how the fraction of variation explained at the local level increases with the granularity of the geographic area. In each cell, we report the unadjusted R^2 (expressed as a percentage) of the ZIP code–level housing variable regressed on different levels of geographic aggregation–time dummies, as in the equation given above. First, when geographic aggregation is at the ZIP code level, the R^2 is 100 percent, as ZIP code–time dummies span the full panel data set. Note that we report unadjusted R^2 because doing so provides a clear benchmark of 100 percent explained variation when we use the data with the highest level of granularity.

Next, consider aggregation at the city level to explain ZIP code–level variation in housing markets. Looking across columns 1 to 5 of the top panel of Table 5, note that a city-level variable can explain a significant amount of the variation (which ranges between 71 percent and 85 percent) for local ZIP code–level housing variables. For example, 71 percent of the local variation in the foreclosure rate can be explained by a city-level variable (column 5), while about 86 percent of the variation of ZIP code–level house prices can be explained by a city-level economic variable (column 1). However, note that at the city level, we have, on average, fewer than two ZIP codes grouped together in our data. Hence, this upper bound on R^2 would likely be tighter in the broader sample of ZIP codes.

Explanatory power monotonically decreases as we consider coarser geographic areas. For example, the fractions of the mortgage delinquency and foreclosure rates that can be explained by county-level aggregation are, respectively, about 44 percent and 35 percent. This pattern suggests a large local variation at the ZIP code level that is not captured by county-, state-, or national-level data.

Finally, we move to national-level aggregation, which means using a national time series trend to explain the ZIP code–level time series patterns for the housing market. We see that for the house price growth rate, this time series can explain about 34 percent of the local variation (column 1 of Table 5, top panel). Though a decent fraction, this still represents a significant drop from the county-level result (80 percent). For delinquency and foreclosure rates, the national-level time series pattern can merely explain 27 percent and 14 percent of the local pattern, respectively (columns 4 and 5).

As a robustness check, we also consider the heterogeneity of growth rates for housing variables. The bottom panel of Table 5 presents the corresponding results. As in the case of the level variables, the explanatory power substantially decreases as we consider coarser geographic areas. For example, the fraction of house price growth variation that can be explained by county-level aggregation is 56 percent. Explained variation decreased by 21 percentage points when moving from city-level aggregation (from 77 percent to 56 percent). The drop is more extreme for other housing variables. For DTIs, CLTVs, delinquency rates, and foreclosure rates, the variation explained drops significantly, from about 60 percent to slightly above 10 percent, when we move from the city to county levels. This interesting pattern suggests a large local variation at the city (or even finer) level that is not captured by county-, state-, or national-level data.

Finally, when we move to national-level aggregation, we see that this time series can explain about 25 percent of the local variation in the house price growth rate (column 1 of Table 5, bottom panel). Though a decent fraction, this still represents a significant drop from the city-level result. For all other variables, including delinquency and foreclosure rates, the national time series pattern can explain merely 0.5 to 4.5 percent of the local pattern (columns 2 through 5, bottom panel).

Overall, the upper bound to R^2 is uniformly lower for growth rates relative to levels of variables at all levels of aggregation. This decrease is likely because there is typically more variation in growth rates than level variables, which might be very persistent. This is especially true for variables such as delinquency and foreclosure rates. Finally, for robustness, online appendix Table A5 shows the results for the variables measured at monthly changes. We find very consistent evidence with our analysis given above for the level and growth rates of housing variables.

The statistical exercise above identified an upper bound on the informativeness of various economic variables by their level of geographic aggregation. We next assess the actual association of various national-, county-, and ZIP code–level variables with ZIP code–level delinquency and foreclosure rates. We begin by investigating the association between ZIP code–level delinquency rates and lagged national-level variables, including the average unemployment rate, house price growth, income growth, federal funds rate, CLTV, DTI, and VantageScore. We consider four lags for each independent variable. All the variables are measured at a quarterly frequency, with the exception of income growth. Income growth is only available on an annual basis, so each quarter is given the value of that year's annual income growth. That is, the lagged income growth for all four quarters of a year receives the value of annual income growth of the previous year.

More specifically, we run a regression of the following form for the quarterly delinquency rate in ZIP code *i*:

$$delinquency_{i,t} = \alpha + \sum_{n \in N} \sum_{j=1}^{4} \beta_{n,j} X_{n,t-j} + \varepsilon_t.$$

We consider two variants of this regression: one without and one with nonlinear (squared) terms for the independent variables.

The first row of column 1 in Table 6 shows the adjusted R^2 from this regression. Column 2 shows the corresponding results from the specification with nonlinear terms. As we observe, national economic variables account for about 20 percent of within-sample variation in the ZIP code–level delinquency rate.

We next estimate the above specification when instead we explore the association between quarterly ZIP code–level delinquency rates and the four quarterly lags of the county-level variables to which a given ZIP code belongs. The county-level variables include the unemployment rate, house price growth, CLTVs, DTIs, and the county average VantageScore. Note that the county unemployment rate is only available on an annual basis, so we convert it to quarters, as we also do with annual income data. The second row of columns 1 and 2 in Table 5 shows the adjusted R^2 from these regressions. As we observe, county-level economic variables account for about 39 to 42 percent of within-sample variation in the ZIP code–level delinquency rate, a substantial improvement over national indicators.

Next, we move to an even more granular level and consider regressions with lagged ZIP code–level variables. Unfortunately, we do not have unemployment or income data at this level. The ZIP code–level variables only include house price growth, CLTVs, DTIs, and the average ZIP code–level credit scores of mortgage borrowers. The third row of columns 1 and 2 in Table 5 shows the adjusted R^2 from these regressions. As we observe, ZIP code–level variables account for about 67 to 87 percent of within-sample variation in the ZIP code–level mortgage delinquency rate, a very substantial improvement over both national- and county-level indicators. Moreover, the fourth, fifth, and sixth rows of columns 1 and 2 in Table 5 show that adding national- and county-level indicators to ZIP code–level ones leads to only minor increases in the adjusted R^2.

Columns 3 and 4 of Table 5 show the corresponding analysis for the ZIP code–level foreclosure rate. Again, as in the case of the delinquency rate and consistent with our analysis given above, we see that ZIP code–level indicators account for much more of the variation in ZIP code–level foreclosure rates than county- or national-level indicators.

We conclude this analysis by conducting additional robustness checks on our inferences. We do so by studying the predictability of local housing-related variables with corresponding lagged variables at different levels of geographic aggregation. Formally, we estimate a simple AR(12) process for each of the ZIP code–level local housing variables:

$$Y_{i,t} = \sum_{\tau=1}^{12} \beta_\tau Y_{j(i),t-\tau} + \varepsilon_{i,t}\ .$$

Again, $Y_{i,t}$ refers to the monthly ZIP code–level housing variables: house price, CLTV, DTI, delinquency rate, foreclosure rate, and real house price. All these variables are measured as demeaned growth rates. We regress this on 12 lagged terms of $Y_{j(i),t-\tau}$ with different degrees of aggregation. In particular, j is measured at the ZIP code and national levels.

We measure the performance of predictability R^2, that is, the percentage of variation explained by the 12 lagged terms of the same variable. For ease of illustration, we compare the predictability performance between ZIP code–ZIP code and ZIP code–national levels. We measure the predictability of housing variables in growth rate measures rather than level measures. The primary reason for this is that level variables exhibit strong autocorrelation. In this sense, a good fit of the level of housing market variables may not translate into a good prediction of the change of variables.

In unreported regressions, we find that real house price growth exhibits high predictability at the ZIP code level, using lagged ZIP code information. Lagged, ZIP code–level house price growth explains roughly 69 percent of the variation in house price growth, with a reduction in the root mean squared error of 0.56 percent. When using national house price growth, only 18 percent of the variation is explained. The other local housing variables are much more difficult to predict, even when using lagged local data. The adjusted R^2 for CLTVs is 2.64 percent; for DTIs 12.47 percent; for delinquency rates, 20 percent; and for foreclosures, 3.6 percent. At the national level, the explanatory power is approximately 0.

For ease of illustration, we present these inferences in online appendix Figure A6, where we illustrate the predictability at the local (ZIP code) and national levels by plotting the 12 autoregressive regression coefficients of local variable growth regressed on lagged local variable growth and lagged national variable growth, respectively. All these local variables exhibit mean reversion except the local house price, which exhibits short-term momentum.[19] For instance, an increase in the local delinquency growth rate of 1 percent predicts a subsequent decrease in the delinquency growth rate of 0.5 percent in the next month and 0.3 percent in the subsequent month. Similar patterns of mean reversion are also found for DTIs, CLTVs, and foreclosure rates. In contrast, national-level housing variables have little to no predictive power, with $R^2 \approx 0$.

[19] For a recent, comprehensive analysis of house price dynamics, see DeFusco, Nathanson, and Zwick 2017.

4.4 Summary

Overall, our evidence indicates that regional economic conditions display significant heterogeneity across U.S. states, counties, and ZIP codes, and over time. This heterogeneity, along with our analysis given above, suggest significant gains from using indexes that capture the local component of economic conditions in assessing the condition of the local housing markets relative to indexes at coarser levels of geographic aggregation. We also find evidence of significant instability over time in the strength of the relationships between key economic variables. In the next section, we discuss the implications of these findings for the design of mortgage contracts and debt relief policies.

5. Implications for Mortgage Contract Design and Debt Relief Policies

In the previous section, we presented evidence of significant heterogeneity in local economic conditions across space and time. Moreover, the strength of the relationships between key economic variables affecting households, such as housing and income risk, do not appear to be that stable over time, pointing to a time-varying distribution of these variables. Recall that our simple framework in subsection 3.2 illustrates that the successful implementation of indexed mortgage contracts or debt relief policies crucially depends on a correct understanding of the underlying structure of income and housing risk and its relation to the indexes on which such contracts or policies will be based. We now discuss the practical considerations in implementing temporary debt relief during economic downturns, drawing on insights from Section 3. We consider solutions that rely, in turn, on the design of mortgage contracts, ex post private renegotiation, and mortgage debt relief policies.

5.1 Mortgage Contract Design

We start by considering implementation of debt relief through changes in mortgage contract terms. We focus on different methods of mortgage contract indexation. As noted above, such indexation effectively implements an automatic temporary debt relief during economic downturns and can potentially circumvent the frictions discussed in Section 2.

5.1.1 ARMs

National Interest Rate Indexation

In practice, the most commonly used state-contingent mortgages are ARMs. Our discussion in Section 3 suggests that contracts indexed to national-level interest rate indexes may implement debt relief efficiently, as long as such indexes closely co-move with house prices and borrowers' incomes. Having said this, two points are worth emphasizing. First, as the top panel of Figure 16 shows, there is a large spatial variation in the ARM share of mortgages in a ZIP code. This implies that an automatic pass-through of low interest rates would be differentially passed through across regions, to the extent that FRMs would remain a popular contract choice.[20] Second, the empirical evidence we presented in Section 4 suggests that, due to significant regional heterogeneity, ARMs based on national indexes may not be as effective, especially in some regions.

One could argue that the most relevant aspect in successfully implementing debt relief through this mechanism is the close co-movement of interest rates on which ARMs are indexed with relevant local economic variables during recessions. Indeed, Figure 5 shows that during the Great Recession, all U.S. states experienced some decline in economic activity, though substantial heterogeneity remained in the strength of this effect. During this period, all main interest rate indexes also reached historically low levels. Consistent with this observation, the empirical evidence finds that ARMs resetting to a low rate after 2009 had a direct positive impact across borrowers and regions by reducing the default rate and increasing consumption, house prices, and local employment (Di Maggio et al. 2017; Fuster and Willen 2017).

One should be careful, however, not to overstate the benefits of ARM contracts. First, there were periods in the past, such as the stagflation episode, when interest rate indexes reached high levels during an economic downturn. In such an environment, a high share of ARMs in an economy could exacerbate the severity of the economic crisis. In particular, after a substantial increase in interest rate indexes along with the federal funds rate (see online appendix Figure A7), the ARMs reset to much higher rates during the period from late 2006 to early 2008. As a consequence, these ARM borrowers faced substantial rate increases, along with vanishing refinancing opportunities, due to the collapse of the subprime mortgage market by mid-2007. This aspect could have contributed to the mortgage default rate and the severity of the initial stage of the financial crisis.

To illustrate this more formally, we use monthly, loan-level panel data from BlackBox Logic on more than 1.8 million two-year, subprime ARMs that reset during this period. These two-year ARM contracts are loans that were mainly originated during the 2004–06 period. They faced a fixed initial rate for the first two years and subsequently a reset to the variable rate based on a short-term interest rate index (for example, the London Interbank Offered Rate). The top panel of Table 7 shows the summary statistics for these loans, including the mortgage interest rate before and after the first reset.

Columns 1 and 2 in the bottom panel of Table 7 show the regression results of the monthly mortgage interest rate and the default rate for two-year ARM borrowers on the time dummies for

[20] This share, however, needs to be interpreted with caution, because many subprime ARM contracts feature various caps and floors that may limit the extent of downward adjustment of their rates in response to a decline in interest rate indexes.

the three quarters before and four quarters after the change in the interest rate, with the fourth quarter before the reset period serving as the excluded category. This specification controls for a variety of borrower, loan, and regional characteristics, including the borrower's FICO credit score and the loan-to-value ratio. As we observe from the top panel, two-year ARM borrowers experience an increase of about 1.3 percentage points in their monthly interest rates after the reset, amounting to a relative increase of more than 17 percent. Column 1 in the bottom panel shows similar effects. The top panel of Figure 17 shows the corresponding monthly mean default rate (serious delinquency rate) for these loans around the first reset date that happens after month 24 of the loan's life. Note that it would take at least two months to see the effects of the reset, because for a loan to be considered seriously delinquent, it needs to be 60 days or more past due on payments. We observe a very substantial increase in the default rate just after the first reset date. Column 2 in the bottom panel of Table 7 confirm this inference by showing that these borrowers experience an estimated absolute increase in the monthly default rate of between 2.2 and 3 percent during the four quarters after the first reset, a very substantial effect (a relative increase of more than 100 percent).

One could worry that some of these effects reflect selection on unobservables, due to refinancings around the reset date. In particular, if better-quality borrowers refinance their loans before the reset, the increase in the default rate could reflect the change in the sample composition. To address this concern, we also consider a sample of more than 180,000 three-year ARMs that were originated between March 2004 and January 2005, when interest rates were at a historically low level (the top panel of Table 7 shows summary statistics for this sample). These loans experienced the first reset three years after their origination, which corresponds to the period April 2007 to January 2008. During this period, interest rates were at a relatively high level (see online appendix Figure A7), inducing substantial rate increases after resets and the private label refinancing market virtually collapsed, limiting refinancing opportunities. The two right-most columns show that there is also a strong association between the interest rate reset and the increase in the default rate in this sample (also see the bottom panel of Figure 17). In particular, three-year ARM borrowers in our sample experienced a rate increase of about 1.58 percentage points after the reset (a relative increase of about 25 percent) and an associated increase in the monthly default rate of between 1.2 and 2 percentage points (a relative increase in the default rate of more than 100 percent). Overall, this evidence suggests that an increase in interest rate indexes at the outset of the Great Recession contributed to the high default rate among ARM borrowers, possibly exacerbating the initial stage of the recent housing crisis. This effect could have been quite important, given that the majority of subprime loans originated before the crisis were ARMs.

The ability of mortgages indexed to national-level rate indexes to serve as effective debt relief also depends on the nature of monetary policy. In particular, policymakers could take into account

ARMs' potential role as an automatic stabilizer in setting national-level interest rates indexes. This, however, would require a careful and up-to-date assessment of economic conditions faced by borrowers, because ARM contracts could also accelerate the pass-through of policy mistakes to households and the real economy. Moreover, a system with a larger share of ARMs could also complicate the central bank's price stability objective, given that increases in interest rates can be highly unpopular with homeowners, creating political pressure to keep rates low for an extended period (Campbell 2013). Finally, in our analysis we have not considered inflation risk. In an environment with significant relative price instability, FRM contracts may provide additional benefits to borrowers by insuring them against fluctuations in nominal interest rates and the associated potential increase in their debt payments relative to their incomes (Campbell and Cocco 2003).

Regardless of such factors, the significant regional heterogeneity we document in the previous section indicates that one-size-fits-all contract indexation based on national-level variables may reduce the effectiveness of such solutions. Instead, there may be gains from pegging mortgage contracts to more granular regional conditions. Of course, such gains would need to be traded off with the potential costs of implementing such indexation, as noted above, including the costs of introducing new contracts with limited prior market experience.

House Price Indexation

An alternative indexation form consists of mortgages that depend on local house price indexes. Given significant regional heterogeneity, one key advantage of such contracts is that their terms can be more closely tied to local economic conditions, as opposed to ARMs that are tied to national-level interest rates. Shiller (2008) has long advocated for such "continuous workout mortgages." Piskorski and Tchistyi (2017) show that home equity insurance mortgages that are indexed to house prices, by alleviating incentives to default strategically, arise as an equilibrium contract in the private lending market, with empirically relevant frictions. Moreover, Greenwald, Landvoigt, and Van Nieuwerburgh (2018) show that indexation to local house prices can reduce financial fragility and improve risk-sharing if intermediaries retain a significant portion of loans on their balance sheets.

The widespread adoption of mortgages indexed to house prices would require timely and accurate regional house price indexes. Such indexes were unavailable in the past. ZIP code–level house price indexes have only recently been developed and started being offered by data providers such as Zillow (see online appendix Figure A8 for an example).

Labor Income Indexation

Our discussion in Section 3 suggests that conditioning mortgages on indexes capturing both local labor market conditions and house prices may provide additional efficiency. However, such arrangements would require timely and accurate regional local labor market indexes, which as of today are not commercially available. Having said this, current unemployment insurance programs implicitly provide a form of labor income indexation—partly insuring households against income shocks—which helps distressed households service their mortgage debt obligations (Hsu, Matsa, and Melzer 2018). Moreover, one could consider social transfer programs that directly provide temporary subsidies, reducing the mortgage payments of unemployed borrowers. Such programs could condition the terms of the transfers on the specific financial position of the borrower (for example, their debt burdens or DTIs). A potential downside of such approaches is that they could result in moral hazard risk for borrowers (Mayer et al. 2014), which we discussed in Section 2. For example, providing mortgage payment support to unemployed borrowers may erode their incentives to find a new job. The potential for such unintended consequences of social transfers for the unemployed has long been recognized in the unemployment insurance literature.

Discussion

It is important to note that we have not suggested designing contracts just based on individual CLTVs, house prices, or employment rates. Although potentially more efficient in capturing the risk dynamics of individuals, such contracts may create moral hazard in terms of incentives to pay or maintain a house. For this reason, we focus our discussion on indexes capturing these variables at the regional level (for example, at the ZIP code level).

It is worth discussing the fact that such indexation forms have not been widely implemented in the past, for several reasons. First, substantial government involvement in the mortgage market through a system of subsidies and regulations favors traditional contracts like FRMs and ARMs. This potentially suppresses the adoption of new mortgage designs.

Second, the implementation of a new mortgage design may require a significant amount of time due to private market inertia, learning, or a low perceived value of such innovations from the ex-ante perspective. For example, before the Great Depression, mortgage contracts were predominantly short-term loans. The inability to roll these loans over was a major factor in the collapse of the financial and housing markets. As a result, the government helped the private market develop and standardize the fully amortizing, long-term contracts, such as FRMs, that currently dominate the U.S. housing market.

Third, the widespread adoption of contracts indexed to local economic conditions would require timely and accurate regional indexes. Such indexes were unavailable in the past, and only recently have some been developed and started being offered by data providers (for example, house price indexes). As we discussed in subsection 4.3, given significant regional heterogeneity, the lack of such indexes at a sufficiently granular level may have significantly reduced the potential efficiency of such solutions, stifling incentives for their development.[21]

Fourth and finally, it is not clear that private market innovation would lead to the successful development of such contracts. In particular, Piskorski and Tchistyi (2017) show that in the competitive equilibrium setting with empirically relevant frictions, unrestricted competition in mortgage design may lead in some cases to market instability (that is, the nonexistence of equilibrium). Their findings highlight the potential, understudied role of the government-sponsored enterprises (GSEs): By subsidizing a restricted contract choice through their guarantee system, the GSEs may help facilitate the existence of a stable mortgage market by limiting private competition in mortgage design. In this regard, the government may also play a potentially important role in the practical implementation of new mortgage designs by promoting certain contracts through its system of subsidies. The downside of this approach is that it would require a continued operation of the GSEs and the Federal Housing Administration—institutions that are plagued by political economy concerns surrounding implicit and explicit government guarantees. Such an approach would also limit the ability to use market pricing for assessing the cost of insurance embedded in new mortgage designs.

5.2 Leverage Regulation and Downpayment Limits

Intuitively, as we discussed in Section 3, the benefits of indexed mortgage contracts or debt relief policies are much smaller for borrowers with significant housing equity. Hence, an alternative approach to decrease the likelihood and costs of future housing crises is preventing households from becoming highly leveraged in the first place.[22] This approach would also alleviate the impact of the frictions discussed in Section 2 by simply limiting the number of borrowers who require debt relief in the first place. One way to implement this in practice would be to impose contract restrictions, like stricter minimum down payment limits, in the current mortgage market setting. Of course, such policies, though potentially simpler to implement than other approaches, could result in additional welfare costs by delaying or preventing homeownership for some borrowers. There are also political economy considerations that might prevent such regulation from being imposed, especially when the housing market is booming. Finally, given significant evidence of misreporting the true extent of down payment and housing equity by financial intermediaries

[21] In addition, Hartman-Glaser and Hébert (2017) point out that if there are informational asymmetries between borrowers and lenders about the ability of such indexes to measure underlying states, the risk-sharing ability of state-contingent contracts based on such indexes can be limited.
[22] For a recent analysis of such policies, see DeFusco, Johnson, and Mandragon (2017).

before the recent crisis (Ben-David 2011; Piskorski, Seru, and Witkin 2015; Griffin and Maturana 2016), such policies may face additional implementation hurdles.

5.2 Private Mortgage Renegotiation

Another approach to implementing debt relief is to rely on private renegotiation efforts. Because foreclosure can induce significant deadweight costs, there should be instances where both borrowers and lenders would find it beneficial to temporarily reduce their household debt burden during economic downturns. As we discussed in Section 2, this approach faced a number of limitations during the Great Recession that were related to organizational frictions and capacity constraints in the intermediary sector, agency conflicts in securitization, ex post moral hazard concerns of intermediaries, and the inability of intermediaries to identify instances when renegotiation might be beneficial for both borrowers and lenders. More broadly, private ex post renegotiation may not take into account the positive externalities of debt relief and may also be of limited scope relative to solutions that rely on ex ante committed changes in contract terms.

5.3 Public Debt Relief Programs

An alternative approach to implementing debt relief policies is to leave the structure of the mortgage contracts intact and instead rely on large-scale government programs or monetary policy. Indeed, during the Great Recession the Federal Reserve reduced short-term interest rates and made large purchases of mortgage-backed securities in one attempt, among others, to support the prices of assets such as houses and lower the incidence of foreclosures.

Moreover, in response to the recent crisis, the administration passed two unprecedented and large-scale debt relief programs: HARP, aimed, again, at stimulating the mortgage refinancing activity of up to 8 million heavily indebted borrowers; and HAMP, aimed at stimulating a mortgage restructuring effort for up to 4 million borrowers at risk of foreclosure. Other notable programs during Great Recession included first-time buyer tax credits aimed at stimulating house purchases (Berger, Turner, and Zwick 2016) and programs aimed at stimulating consumer spending, such as economic stimulus payments (Parker et al. 2013) and subsidies for new car purchases (Mian and Sufi 2012).

As we discussed in Section 1, various implementation frictions—including the nature of mortgage contracts and the ability of financial intermediaries to quickly implement debt relief—can hamper the effectiveness of ex post solutions. For example, Di Maggio et al. (2017) show that a significant heterogeneity in the ARM share across regions (see the top panel of Figure 16) resulted in a significant differential pass-through of lower interest rates to households.

Unlike one-size-fits-all monetary policy, HAMP and HARP implemented a form of specific, individual targeting. This, given the evidence of significant regional and borrower heterogeneity, could in principle increase their efficiency. HAMP was mainly targeted at distressed borrowers with high DTIs. However, the need to verify such program criteria, coupled with the limited organizational ability of financial intermediaries, significantly hindered its effectiveness (Agarwal et al. 2017a). The middle panel of Figure 16 shows significant regional variation in the share of intermediaries with an organizational design that is conducive for renegotiation, suggesting that debt relief through a program like HAMP would be differentially passed through across space.

Wide-scale refinancing programs such as HARP may be easier to implement because they stimulate a more routine activity like refinancing rather than loan renegotiation. Moreover, a program like HARP, which was based on CLTVs, implicitly indexed local house prices on a granular level. However, because the implementation of this program was through intermediaries, its effectiveness was hampered both by intermediary friction, such as capacity constraints, and by market design, such as competition in the refinancing market (Agarwal et al. 2017b; Fuster et al. 2017). Moreover, the program targeted only loans issued with prior GSE guarantees (agency loans), which usually correspond to more creditworthy borrowers than subprime borrowers. The bottom panel of Figure 16 illustrates significant regional heterogeneity in the fraction of loans eligible for HARP (Agarwal and others 2017b). This suggests that debt relief through refinancing would be less likely in some regions than others.

To illustrate the importance of such factors, we conduct a simple analysis of the change in ZIP code–level quarterly delinquency and foreclosure growth after the period starting in 2009 relative to the prior period. We note that 2009 coincides with the introduction of various debt relief programs, including HARP and HAMP, and with a firm commitment to the prolonged policy of low interest rates. In particular, Table 8 shows the results of regressions of the difference in mean ZIP code–level delinquency or foreclosure growth on the ZIP code–level ARM share, high-intermediary capacity share, and HARP-eligible share.

The top panel of Table 8 shows the results for changes in ZIP code–level delinquency growth. Columns 1 through 6 regress the change in delinquency growth separately on each of the independent variables. We find that all coefficients in these columns are negative, indicating that ZIP codes with higher shares of ARMs, high-capacity servicers, and HARP-eligible loans experience faster declines in delinquency growth. Columns 7 and 8 include all these variables together and show that all the coefficients remain negative. The bottom panel of the table shows the results for ZIP code–level foreclosure growth. Consistent with results in the top panel, we find that ZIP codes with higher shares of ARMs, higher shares of loans serviced by high-capacity intermediaries, and higher shares of HARP-eligible loans experienced faster declines in foreclosures after 2009. This evidence is consistent with the findings of Agarwal et al. (2017a, 2017b) and Di Maggio et al. (2017), who show that these factors played an important role in the effectiveness of debt relief measures undertaken during the Great Recession and their differential impact across regions.

Overall, these findings provide guidance for designing large-scale debt relief programs in the future. First, in the case of programs aimed at stimulating mortgage renegotiation activity, such as HAMP, it may have been productive for the program to have allowed the easy transfer of distressed mortgages from inefficient servicers to those more capable of conducting many renegotiations. One way to address this issue in the future is to rely more heavily on special servicers, as is common in the commercial real estate market. Upon the occurrence of certain specified adverse events, the nonperforming loans would be automatically transferred to organizations better equipped to handle such assets. Moreover, there is a likely trade-off between screening more intensively—which limits the potential costs of such programs, including strategic defaults—and the reach and pace of the program.

Wide-scale refinancing program such as HARP may be easier to implement because they stimulate a more routine activity like refinancing rather than loan renegotiation. Moreover, a program like HARP, which was based on CLTVs, implicitly indexed on a granular level (local house prices) relative to national-level indexes. However, because the implementation of this program was through intermediaries, its effectiveness was hampered both by intermediary frictions, such as capacity constraints, and market design, such as competition in the refinancing market (Agarwal et al. 2017b; Fuster et al. 2017), and by its eligibility being restricted to loans issued with prior GSE guarantees. Such programs critically rely on the ability of the government to guarantee mortgage debt during a crisis (for example, through the GSEs). Their success also particularly depends on the speed and extent to which interest rates reach sufficiently low levels during housing market downturns, which may impose additional constraints on the conduct of monetary policy. HARP-like policies could also become a part of permanent market arrangements by automatically relaxing housing equity refinancing constraints in regions that have experienced sufficient declines in house prices.

6. Conclusion

In this paper, we have focused on understanding *design* and *implementation* challenges of ex ante and ex post debt relief solutions aimed at a more efficient sharing of aggregate risk between borrowers and lenders. Our analysis and discussion highlight an important trade-off that warrants more research. The indexed mortgage contracts have the advantage of circumventing financial intermediary and other fictions by facilitating a quick ("automatic") implementation of debt relief during economic downturns. However, as illustrated in subsection 3.2, for such contracts to be effective, lenders, policymakers, and borrowers may need a good ex ante understanding of the underlying distribution of risk and its relation to indexes being used when designing and choosing such contracts. Errors in beliefs about the structure of risk can reduce the benefits of such solutions. Given the vast heterogeneity in the nature of risk across space and time, such errors are likely, especially because a major change in the nature of mortgage contracts or housing policy can on its own significantly alter relationships between market equilibrium outcomes in a way that is potentially hard to quantify. Thus, it seems prudent to also rely on ex post debt relief solutions.

Ex post debt relief solutions, on one hand, have the advantage of being more fine-tuned to the specific realization of economic risk. On the other hand, they have limitations. As our analysis and discussion illustrate, these solutions are subject to various implementation frictions that could significantly delay debt relief and hinder its effectiveness.

More broadly, our evidence suggests that an effective mortgage design approach to debt relief requires a more in-depth analysis of the nature of relevant income and housing risk and its evolution across regions and borrowers. This could include the development of new and sufficiently granular indexes on which such contracts could be based. The recent “big data” revolution is promising in this regard. Such an analysis could also identify mortgage designs that effectively implement debt relief across a range of possible environments. As well, the approach relying on ex post solutions should focus on alleviating frictions that may hinder the effective implementation of policies like the ones we discussed in Section 5. Big data might again be useful here, because they could enable the development of an effective and easy-to-verify set of eligibility criteria for debt relief policies.

Our understanding of the design and implementation challenges for ex ante and ex post debt relief solutions suggests that a more resilient mortgage market system will involve a combination of ex ante and ex post policies with state contingencies. At minimum, the state contingency should involve the national-level variables; but as our analysis suggests, it would benefit from more granular variables to better address the variation faced by the borrowers (for example, by incorporating local house price indexes). It is clear that a better designed ex ante state contingency would limit the need to rely on ex post solutions. This may be desirable because a large quantity of distressed loans and the presence of intermediary frictions may not allow for the necessary level of debt relief (such as refinancing) to combat the crisis. This would also alleviate the pressure on financial intermediaries to implement large-scale debt relief (such as loan modifications) during a national crisis. Nonetheless, despite the best laid plans ex ante, it is likely that severe housing market downturns would require interventions ex post. What our analysis and discussion have demonstrated is that, though such ex post policies are easier to design, the implementation challenges are immense and must be thought about carefully. Finally, because the GSEs are likely to dominate the residential lending market, at least in the short to medium terms (Buchak and others 2017), a more resilient, redesigned mortgage system would likely require their active participation, an aspect about which we have been silent. Whether their presence would alleviate the various coordination and implementation hurdles of moving to a new mortgage market architecture remains an open area for discussion and research.

Table 1. Base Parameter Values for the Simple Framework

This table describes a set of base parameter values used in numerical examples in Section 3.2.

Parameter	Definition	Value
y	Income average	200
σ_y	Income standard deviation	70
P_1	House price average	150
σ_P	House price standard deviation	50
δ	Recovery rate when in default	0.7
P_0	Initial house price	100
$\bar{v}$	Loss of utility when in default	50
θ	Utility of living in a house	200
c	Fixed cost of indexed mortgage contract	0 or $0.01P_0$
D	Down payment of house	0 or $0.20P_0$

Table 2. National- and State-Level Economic Variables

This table presents summary statistics of the economic series used in the principal components analysis to extract a proxy for national and local economic conditions. Both GDP and income data are deflated using CPI for all urban consumers. Growth rates are measured at the annual level. Using growth rates avoids the issue of merging two separate datasets because growth rates are computed from the same series and then appended. For national housing data, we use the S&P/Case-Shiller U.S. National Home Price Index. All values are expressed as percentages. *Sources:* U.S. Bureau of Economic Analysis; U.S. Bureau of Labor Statistics; Zillow; Freddie Mac; Standard & Poor's; Federal Reserve Economic Data.

Variable	Mean	Std. Dev.	Minimum	Maximum
National-level				
Real GDP growth	2.28	2.14	-4.78	7.50
Real income growth	2.46	2.22	-5.95	6.21
Unemployment rate	6.35	1.66	3.90	10.80
Real house price growth	0.97	5.46	-11.97	9.96
State-level				
Real GDP growth	2.25	3.69	-27.49	31.12
Real income growth	2.43	2.95	-14.77	20.25
Unemployment rate	6.02	2.13	2.20	18.30
Housing price growth	0.55	6.42	-43.72	47.99

Table 3. Principal Component Analysis

This table displays the summary statistics about the weight on the economic variables in the first factor of the principal component analysis. The weight represents the relative loading of the local economic conditions proxy on local economic variables. The weights are the relative loadings on the first principal component. The analysis is at the state level. All values are expressed as percentages. We present the mean, standard deviation, min and max of the weights for the state level PCA. *Sources:* U.S. Bureau of Economic Analysis; U.S. Bureau of Labor Statistics; Zillow; Freddie Mac; Standard & Poor's; Federal Reserve Economic Data.

Weight	Mean	Std. Dev.	Minimum	Maximum
Real GDP growth	56	4	46	69
Real income growth	53	5	45	73
Unemployment rate	-44	8	-54	-1
Real house price growth	43	11	-8	53
Explained variation	60	9	32	73

Table 4. Heterogeneity in State-Level Economic Factors

This table reports results from ordinary least squares regressions of national macroeconomic variables on state-level economic factors. The regressions are estimated separately for each state iteratively, adding one variable at a time. The national economic factor is the first factor of the principal component analysis at the national level. *Sources:* U.S. Bureau of Economic Analysis; U.S. Bureau of Labor Statistics; U.S. Department of the Treasury; Zillow; Freddie Mac; Standard & Poor's; Federal Reserve Economic Data; authors' calculations.

	Mean	Std. Dev.	Minimum	Maximum
Fraction of variation explained				
National economic factor	52.04	24.10	0.00	82.11
Real GDP growth	42.53	21.15	1.05	68.77
Unemployment	27.95	14.10	0.08	59.19
Unemployment change	18.17	12.03	0.15	53.40
Real income growth	35.37	15.40	3.71	65.08
Real house price growth	21.51	15.90	0.00	55.91
Federal funds rate	5.62	6.50	0.02	26.83
Treasury rate	4.74	4.92	0.06	21.43
Real Treasury rate	13.97	12.78	0.00	45.78
Change in real Treasury rate	2.95	3.78	0.00	14.89
Change in real Treasury rate (t – 1)	1.24	1.37	0.00	5.81
Mortgage interest rate	5.77	6.08	0.01	26.06
Coefficient estimate				
National economic factor	0.70	0.24	0.00	0.94
Real GDP growth	44.91	15.72	5.37	62.23
Unemployment	-46.96	17.55	-77.48	15.24
Unemployment change	-53.94	31.02	-109.69	31.77
Real income growth	40.15	11.29	12.38	55.71
Real house price growth	11.44	7.23	-5.29	22.27
Federal funds rate	-2.69	8.31	-18.12	12.93
Treasury rate	-0.46	9.13	-18.08	16.76
Real Treasury rate	18.49	17.81	-22.97	47.54
Change in real Treasury rate	11.96	10.29	-14.78	35.33
Change in real Treasury rate (t – 1)	6.55	8.24	-13.09	23.67
Mortgage interest rate	-4.33	10.14	-23.26	15.46

Table 5. The Importance of Local Economic Variables: The Upper Bound of R^2

This table shows the upper bound of ZIP code-level variables that can be explained by different levels of aggregation. Specifically, the ZIP code-level variable is regressed on contemporaneous geography-time fixed effects, and the unadjusted R^2, expressed as a percentage, is reported. All variables are demeaned at the ZIP code level and winsorized at 1 percent. Real house price data are from January 1997 to December 2017. The remaining variables are from July 2005 to December 2017. *Sources:* Zillow and Equifax.

	House price	Combined LTV ratio	Debt-to-income ratio	Delinquency rate	Foreclosure rate
Aggregation	(1)	(2)	(3)	(4)	(5)
Levels					
National	34.44	39.39	13.00	26.57	13.81
State	67.17	58.45	21.24	37.52	27.73
Metro	80.48	69.02	29.64	43.51	35.45
County	80.32	67.86	30.42	43.85	35.36
City	85.93	83.04	68.84	76.05	71.45
ZIP code	100	100	100	100	100
Growth rates					
National	24.91	4.45	1.13	1.06	0.58
State	36.88	6.89	1.83	1.92	1.50
Metro	52.23	14.47	8.75	8.36	9.25
County	55.67	16.60	10.75	10.49	12.01
City	77.05	62.17	60.41	60.59	56.93
ZIP code	100	100	100	100	100

Table 6. The Relative Importance of Local Economic Indicators in Accounting for Variation in Mortgage Delinquency and Foreclosure Rates

This table shows the association of various national-, county-, and ZIP code-level variables with ZIP code-level delinquency and foreclosure rates. We perform a series of simple linear regressions, with four lags for each independent variable, and with or without squared independent variables. The adjusted R^2, expressed as a percentage, is reported. The national variables include the unemployment rate, real house price growth, real income growth, the federal funds rate, the combined loan-to-value ratio, the debt-to-income ratio, and the national average VantageScore. The county variables include the unemployment rate, real house price growth, the combined loan-to-value ratio, the debt-to-income ratio, and the county average VantageScore. The ZIP code variables include real house price growth, the combined loan-to-value ratio, the debt-to-income ratio, and the ZIP code average VantageScore of mortgage borrowers. *Sources:* U.S. Bureau of Economic Analysis; U.S. Bureau of Labor Statistics; U.S. Department of the Treasury; Zillow; Freddie Mac; Standard & Poor's; Federal Reserve Economic Data; Equifax; authors' calculations.

	ZIP code-level delinquency rate		ZIP code-level foreclosure rate	
Independent variables	(1)	(2)	(3)	(4)
National	19.5	19.6	8.1	8.1
County	39.2	42.4	54.0	54.7
ZIP code	66.5	87.1	74.1	75.7
ZIP code and county	70.0	88.4	76.2	72.9
ZIP code, county, and national	71.3	89.4	78.3	79.3
Squared terms	No	Yes	No	Yes
No. of observations	262,258	262,258	262,258	262,258

Table 7: Impact of ARM Resets on Mortgage Default at the Outset of Great Recession

Panel A shows the summary statics for the sample of 2-year (2/28) ARMs and 3-year (3/27) ARMs used in our analysis (from BlackBox data). The sample of 2-year ARMs consists of loans that were mainly originated during the 2004-2006 period. They faced a fixed initial rate for the first two years and subsequently a reset to the variable rate based on a short-term interest rate index (e.g., LIBOR rate). The sample of 3-year ARMs consists of loans that were originated during the March 2004 – January 2005 period when interest rates were at a historically low level (Panel A of Table 7 shows summary statistics for this sample). These loans experienced the first reset three years after their origination which corresponds to the period of April 2007 – January 2008. During this period interest rates were at relatively high level (see Appendix A7) inducing substantial rate increases after resets and the private label refinancing market virtually collapsed limiting refinancing opportunities. Columns (1) and (2) of Panel B show the regression results of the monthly mortgage interest rate and default rate of 2-year ARM borrowers on the time dummies for the three quarters before and four quarters after the change in the interest rate due to the first reset with the fourth quarter before the reset period serving as the excluded category. The specifications controls for a variety of borrower, loan, and regional characteristics including the borrower's FICO credit score and the LTV ratio. Columns (3) and (4) of Panel B show the corresponding results for 3-year ARM borrowers. Standard errors are in the parentheses.

Panel A: Summary Statistics

	2-Year ARM Sample		3-Year ARM Sample	
	Mean	S.D	Mean	S.D
FICO	622.8	58.1	682.0897	64.35078
Loan Balance	199,435	130,569	221965.9	172645.7
Initial LTV	82.25	10.96	78.71714	14.78156
Initial Interest Rate	7.80	1.67	6.308084	1.425647
Interest Rate after Reset	9.13	2.10	7.897143	2.225515
Number of Observations	1,815,178		146, 078	

Panel B: Impact on Mortgage Interest Rates and Defaults

	2-Year ARM Sample		3-Year ARM Sample	
	Interest Rate (1)	Default Rate (2)	Interest Rate (3)	Default Rate (4)
Three Quarters Before	-0.0672	0.164	-0.0109	0.110
	(0.001)	(0.014)	(0.004)	(0.032)
Two Quarters Before	-0.183	0.296	-0.0151	0.236
	(0.002)	(0.015)	(0.004)	(0.033)
One Quarter Before	-0.268	0.679	0.0345	0.417
	(0.002)	(0.016)	(0.004)	(0.034)
One Quarter After	1.502	2.226	2.117	1.213
	(0.002)	(0.018)	(0.004)	(0.041)
Two Quarters After	1.426	3.081	1.901***	1.968
	(0.002)	(0.021)	(0.005)	(0.045)
Three Quarters After	1.600	2.500	1.455***	1.816
	(0.002)	(0.025)	(0.005)	(0.047)
Four Quarters After	1.546	2.314	1.323***	2.009
	(0.003)	(0.030)	(0.006)	(0.050)
Other Controls	Yes	Yes	Yes	Yes
Number of Observations	13,036,083	13,036,083	1,089,761	1,089,761
R-square	0.278	0.0130	0.564	0.0284

Table 8: The Association between Mortgage Contract Type and Financial Intermediary Factors and the Housing Recovery

This table reports the results from the OLS regressions of the difference between mean zip code delinquency growth and foreclosure growth after the major interventions in the housing market (2009-2016) and the preceding period (2006-2008) on various zip code level variables. We consider the following control variables at the zip code level: "ARM Share" that captures the share of loans that are of ARM type (from Di Maggio et al 2017), "High Capacity Share" that captures the share of loans serviced by high organizational capacity intermediaries (from Agarwal et al. 2017a), and "HARP Eligible Share" that captures the share of loans that are HARP eligible (from Agarwal et al. 2017b). Controls include DTI, CLTV, and vantage at the zip code level. Panel A shows the results for zip code delinquency growth and Panel B shows the results for zip code foreclosure growth. Standard errors are in the parentheses.

Panel A: Change in the Zip Code Delinquency Growth

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ARM Share	-0.194	-0.119					-0.164	-0.092
	(0.012)	(0.013)					(0.014)	(0.014)
High Experience Share			-0.193	-0.063			-0.083	-0.024
			(0.021)	(0.020)			(0.023)	(0.022)
HARP Eligible Share					-0.061	-0.091	-0.035	-0.073
					(0.011)	(0.011)	(0.012)	(0.011)
Control	No	Yes	No	Yes	No	Yes	No	Yes
Number of Observations	6739	6739	6739	6739	6739	6739	6739	6739
R-square	0.035	0.163	0.012	0.153	0.004	0.161	0.037	0.167

Panel B: Change in the Zip Code Foreclosure Growth

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ARM Share	-0.706	-0.587					-0.535	-0.395
	(0.034)	(0.037)					(0.039)	(0.041)
High Experience Share			-0.681	-0.549			-0.369	-0.39
			(0.058)	(0.059)			(0.064)	(0.063)
HARP Eligible Share					-0.401	-0.424	-0.323	-0.361
					(0.031)	(0.032)	(0.032)	(0.033)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Number of Observations	6739	6739	6739	6739	6739	6739	6739	6739
R-square	0.059	0.106	0.02	0.085	0.024	0.097	0.074	0.123

Figure 1. Utility Gains from Mortgage Indexation under the Simple Framework: Perfectly Correlated Income and Housing Risk

This figure shows the borrowers' utility gains (in percentage terms) under indexed-rate mortgage (IRM) compared to FRM for various degrees of correlation of rate index with income and assuming house price and income shocks are perfectly correlated. Panel (a) shows the results for the case of no indexation cost. Panel (b) shows the results for the case of mortgage indexation cost. Base parameters are as in Table 1.

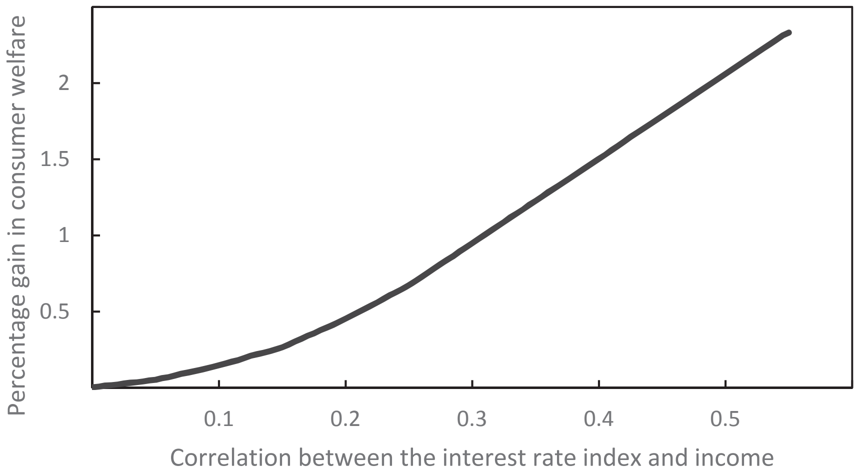

(a) No indexation cost

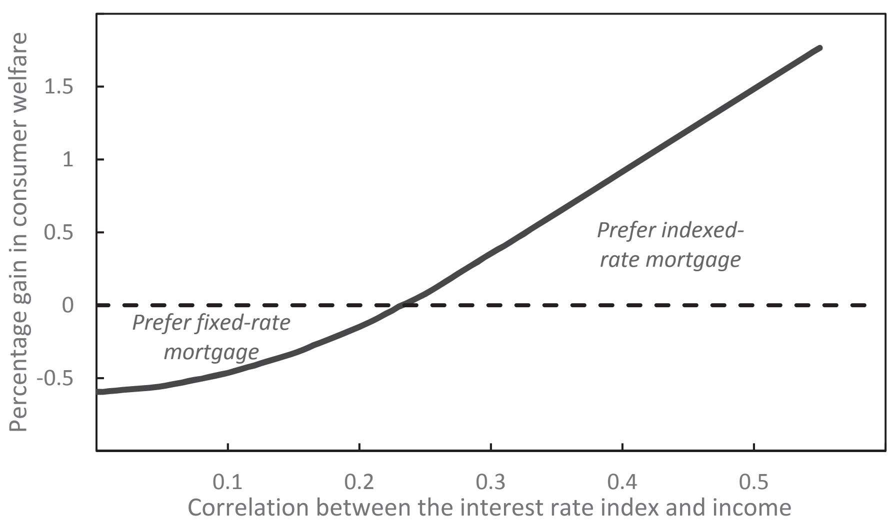

(b) With indexation cost

Figure 2. Utility Gains from Mortgage Indexation under the Simple Framework: No Indexation Cost and Imperfectly Correlated Income and Housing Risk

This figure shows the borrower's utility gain under an indexed-rate mortgage versus a fixed-rate mortgage. Panel (a) shows the results for correlation of income and house prices equal to 0.25 (low correlation) while Panel (b) shows the corresponding results for correlation of income and house prices equal to 0.75 (high correlation). Both cases assume no down payment. Panels (c) and (d) shows the corresponding figures for the case of 20% down payment. All panels assume no indexation cost. The base parameters are from Table 1.

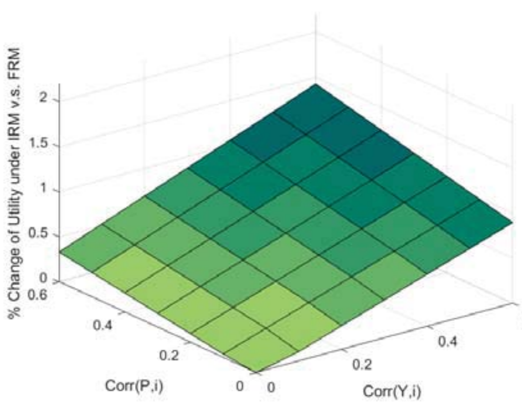

(a) Corr(y,P)=0.25, No Down-Payment

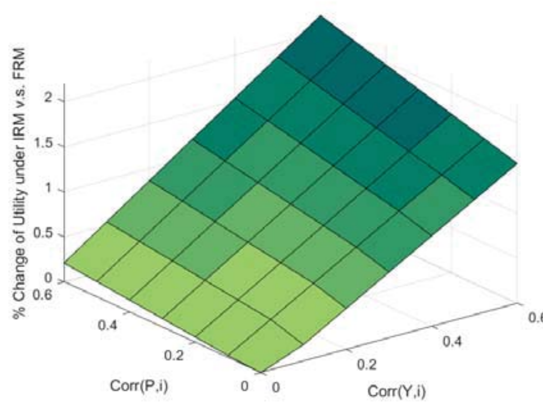

(b) Corr(y,P)=0.75, No Down-Payment

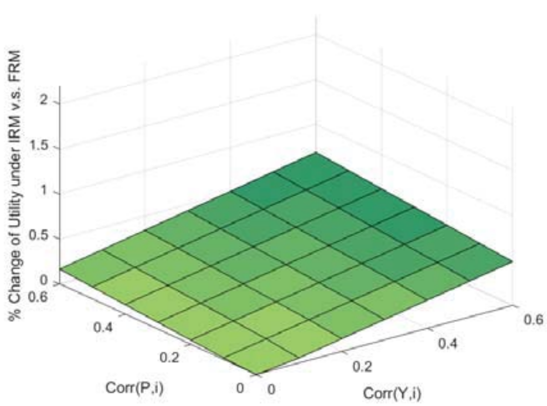

(c) Corr(y,P)=0.25, 20% Down-Payment

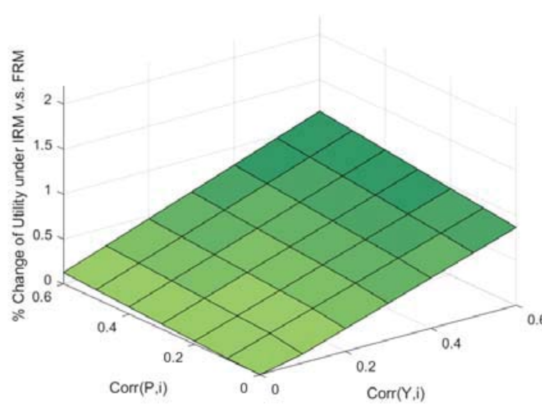

(d) Corr(y,P)=0.75, 20% Down-Payment

Figure 3. Utility Gains from Mortgage Indexation under the Simple Framework: Positive Indexation Cost and Imperfectly Correlated Income and Housing Risk

This figure shows the borrower's utility gain under an indexed-rate mortgage versus a fixed-rate mortgage. Panel (a) shows the results for correlation of income and house prices equal to 0.25 (low correlation) while Panel (b) shows the corresponding results for correlation of income and house prices equal to 0.75 (high correlation). Both cases assume no down payment. Panels (c) and (d) shows the corresponding figures for the case of 20% down payment. All panels assume positive indexation cost. The base parameters are from Table 1.

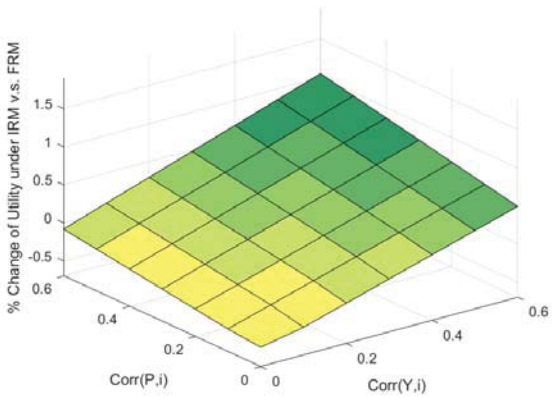

(a) Corr(y,P)=0.25, No Down-Payment

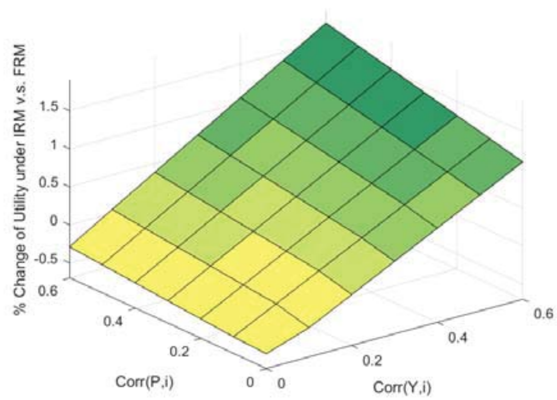

(b) Corr(y,P)=0.75, No Down-Payment

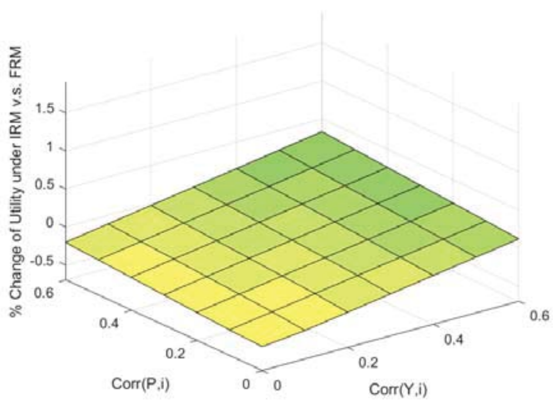

(c) Corr(y,P)=0.25, 20% Down-Payment

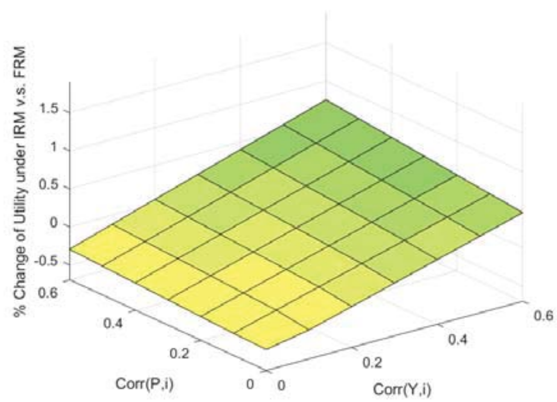

(d) Corr(y,P)=0.75, 20% Down-Payment

Figure 4. The Effect of Incorrect Beliefs about the Distribution of Relevant Economic Risk under the Simple Framework

This figure shows the borrower's utility (in percentage terms) under indexed-rate mortgage (IRM) designed for incorrectly high correlation of income and house prices (equal to 0.75) and high correlation of index with income and house prices (equal to 0.6) relative to the indexed-mortgage correctly designed knowing that the actual correlation of income and house prices is low (equal to 0.25) and the actual correlation of index with income and house prices is as shown in the figure. The computations assume no down-payment and no indexation cost. Other parameters are as in Table 1.

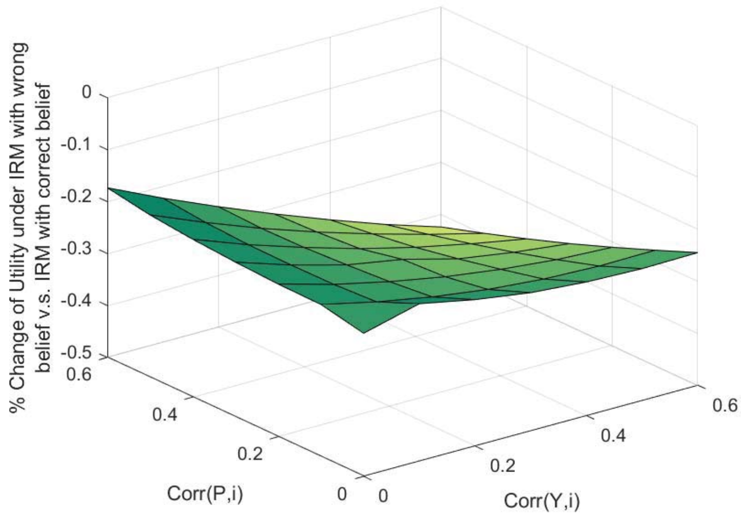

Figure 5. Regional Heterogeneity and the Evolution of the State Economic Factor

The solid black line is the mean across the 50 U.S. states and the District of Columbia of the state economic factor, a measure of local economic conditions. The shaded area denotes the 10th-90th percentile range of the distribution. *Sources:* U.S. Bureau of Economic Analysis; U.S. Bureau of Labor Statistics; Zillow; Freddie Mac; Standard & Poor's; Federal Reserve Economic Data; authors' calculations.

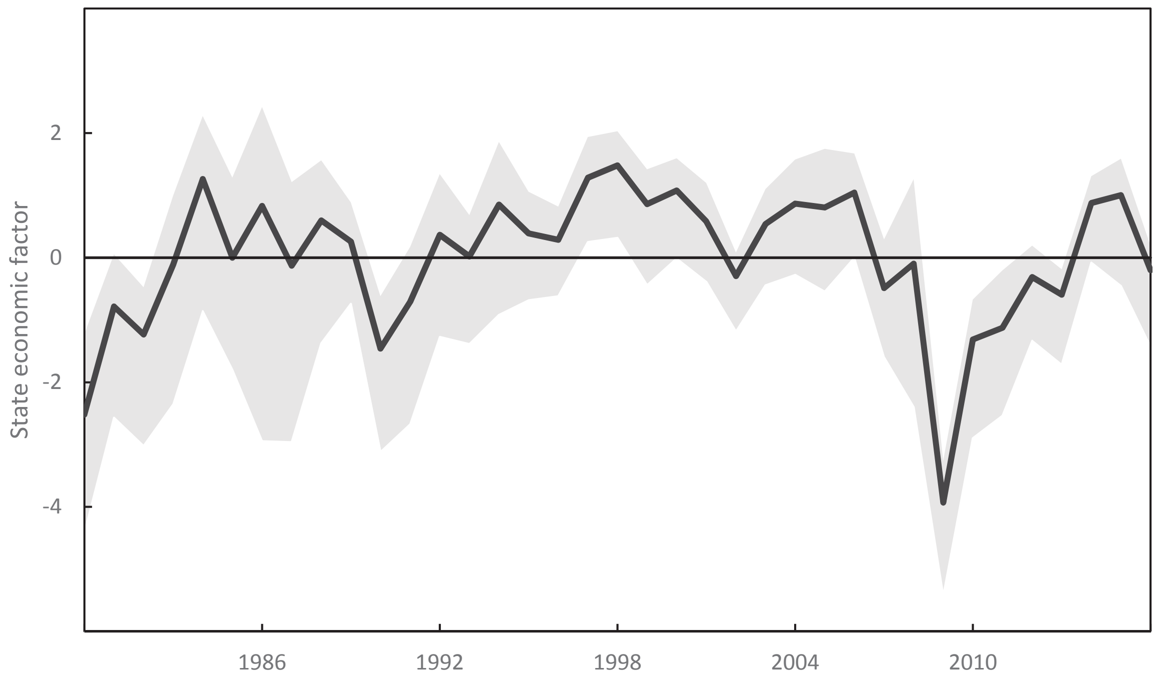

Figure 6. Correlation between State Economic Factors and Various State- and National-Level Variables

Panel (a) of this figure displays the kernel densities of the correlation between the state economic factors with various state economic variables. Panel (b) displays the corresponding kernel densities of the correlation between state economic factors with national economic conditions and various national-level macroeconomic variables. Panel (c) displays the corresponding kernel densities of the correlation between state economic factors with national-level interest rate indices. *Sources:* U.S. Bureau of Economic Analysis; U.S. Bureau of Labor Statistics; U.S. Department of the Treasury; Zillow; Freddie Mac; Standard & Poor's; Federal Reserve Economic Data.

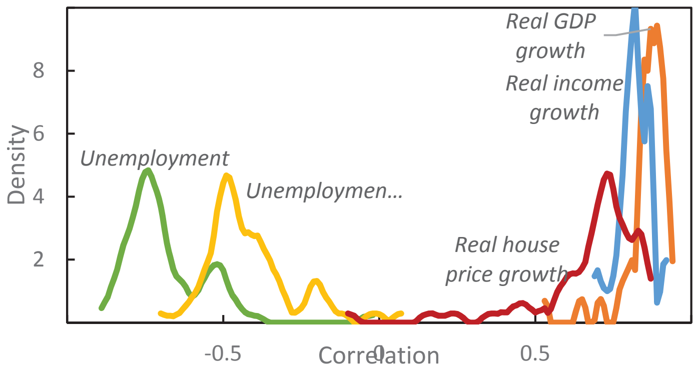

(a) Correlation with State-Level Variables

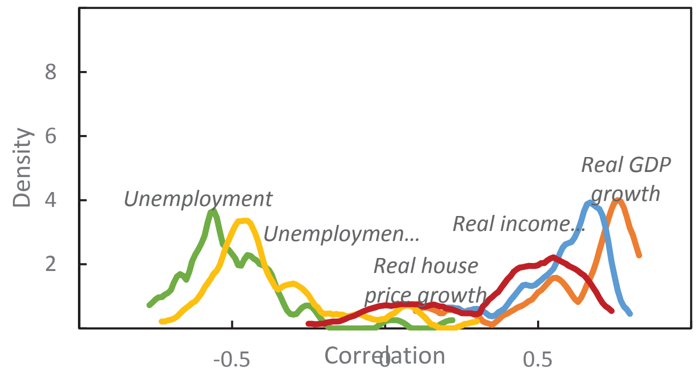

(b) Correlation with National-Level Variables

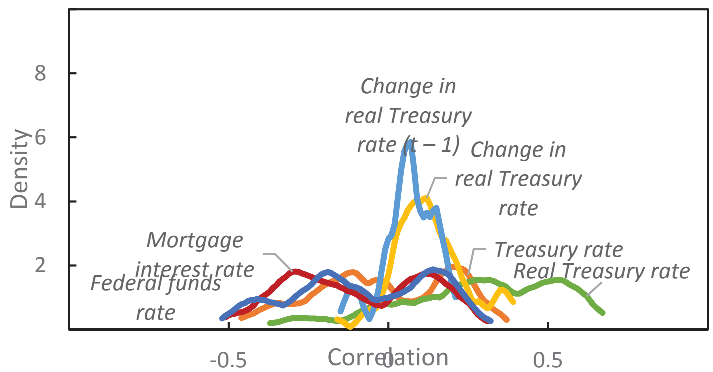

(c) Correlation with Interest Rate Indices

Figure 7. County-Level Income and Unemployment Rate

Panel (a) shows time-series of the means of county-level real income growth from 1998 to 2016. Panel (b) shows time-series of the means of county-level unemployment rate. Calculations are population weighted by county. The bottom panels show the standard deviation of these variables across counties in each year. *Sources:* U.S. Census Bureau, Small Area Income and Poverty Estimates; U.S. Bureau of Labor Statistics, Local Area Unemployment Statistics.

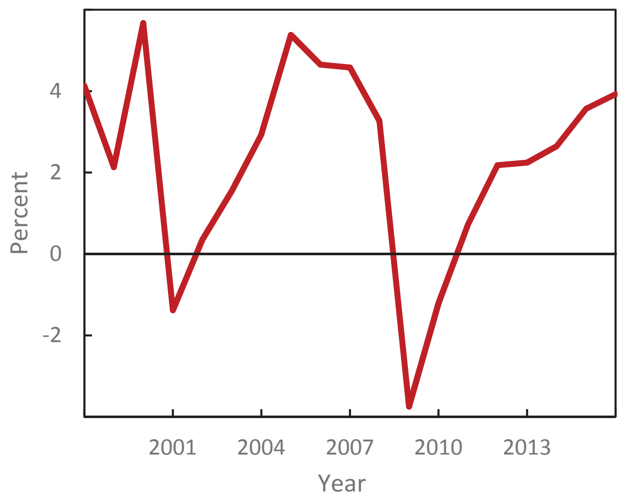

(a) Mean of Income Growth

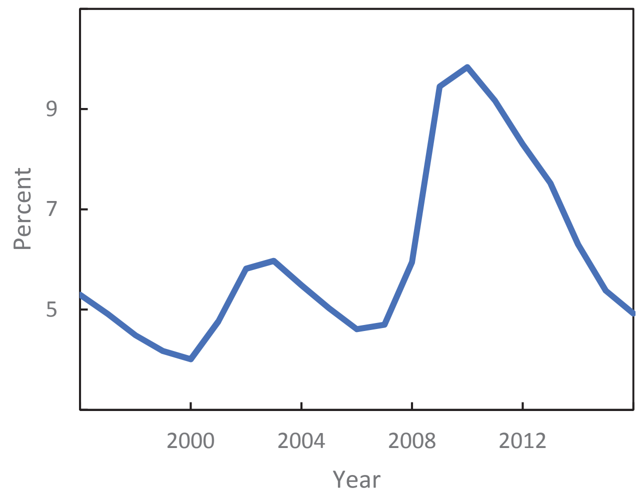

(b) Mean of Unemployment Rate

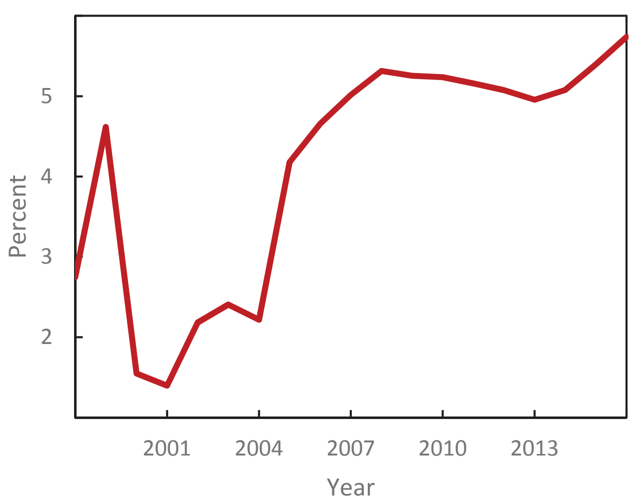

(c) SD of Income Growth

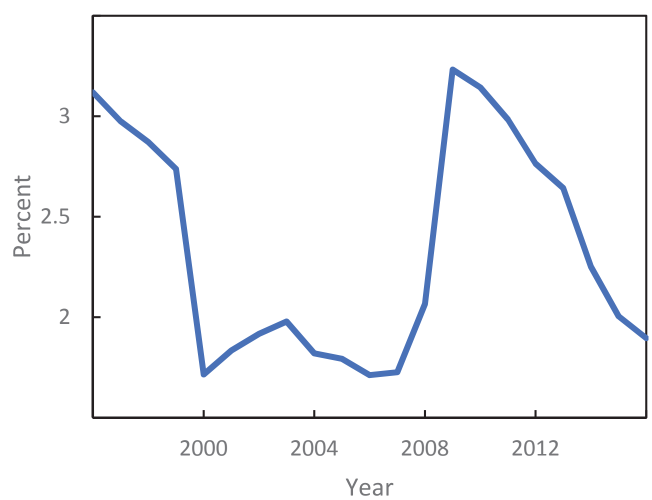

(d) SD of Unemployment Rate

Figure 8. County-Level Housing Variables

The left panels show time-series of the means of county-level house price growth, CLTV, and delinquency rate from 2005 to 2016. The right panels show the standard deviation of these variables across counties in each year. Calculations are population weighted by county. *Sources:* Zillow; Equifax.

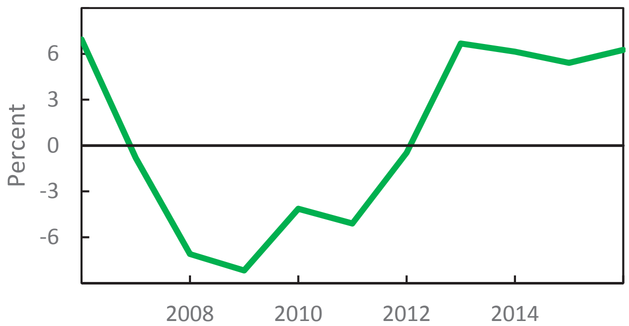

(a) Mean of House Price Growth

(b) SD of House Price Growth

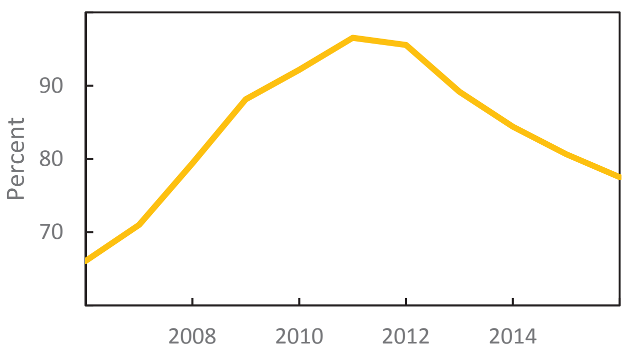

(c) Mean of CLTV

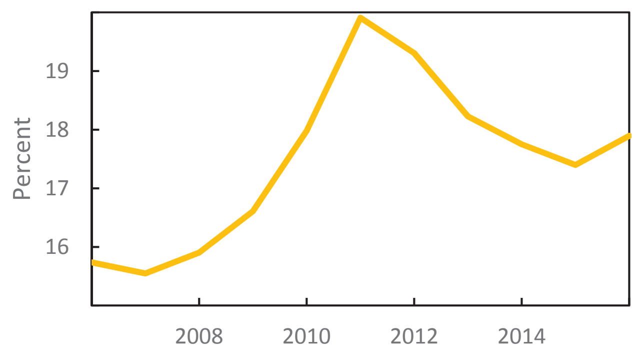

(d) SD of CLTV

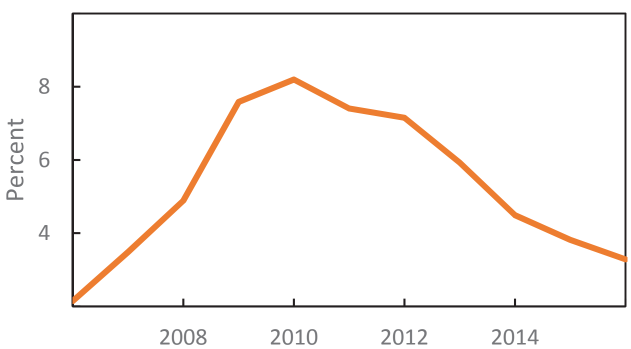

(e) Mean of Delinquency Rate

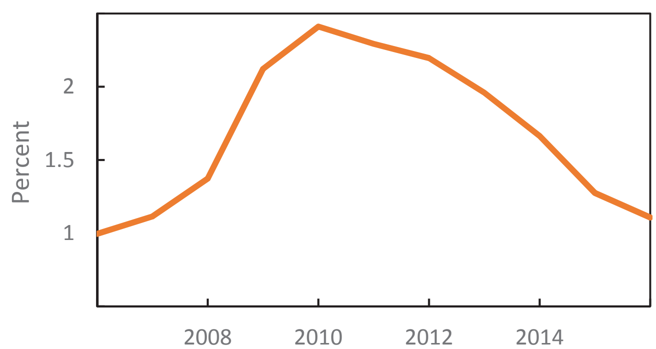

(f) SD of Delinquency Rate

Figure 9. County-Level Unemployment Rates

Panel shows county-level unemployment rate in 2005, the middle panel shows unemployment rate in 2010, and the bottom panel unemployment rate in 2016. *Source:* U.S. Bureau of Labor Statistics, Local Area Unemployment Statistics.

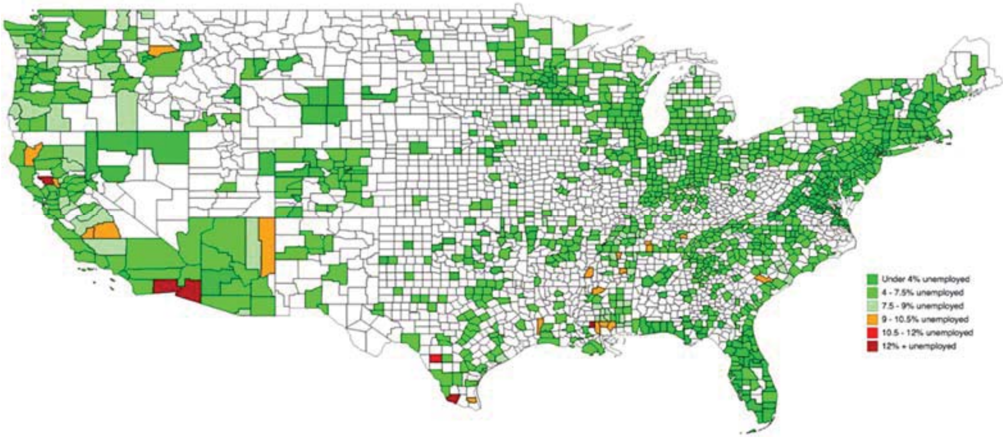

(a) 2005

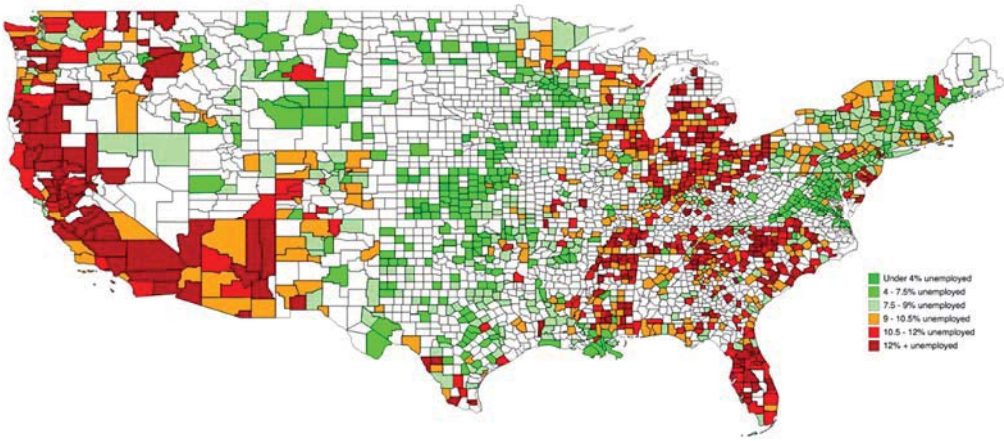

(b) 2010

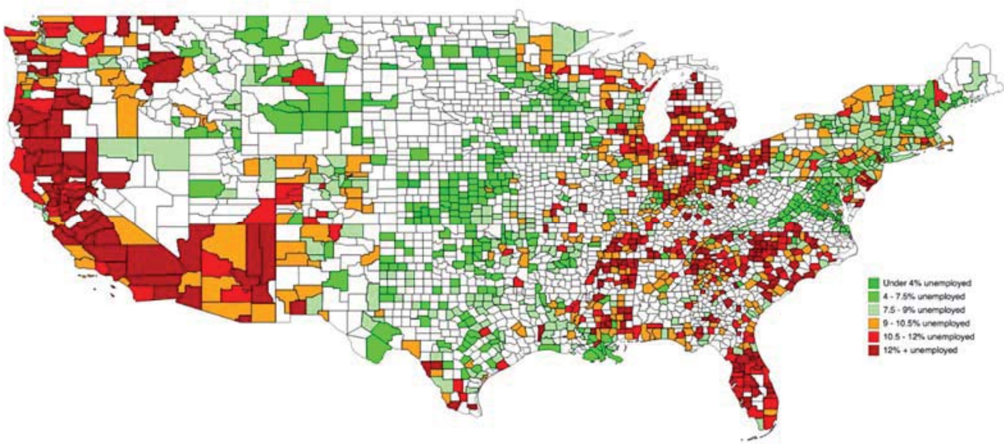

(c) 2016

Figure 10. County-Level House Price Growth

The top panel shows average annual county level house price growth from 2005 to 2006, the middle panel shows house price growth from 2007 to 2009 and the bottom panel shows house price growth from 2010 to 2016. *Source*: Zillow.

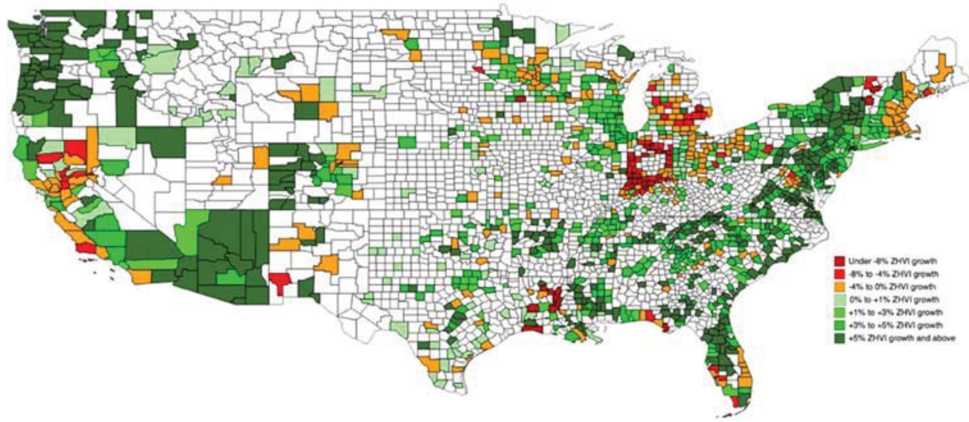

(a) 2005-2006

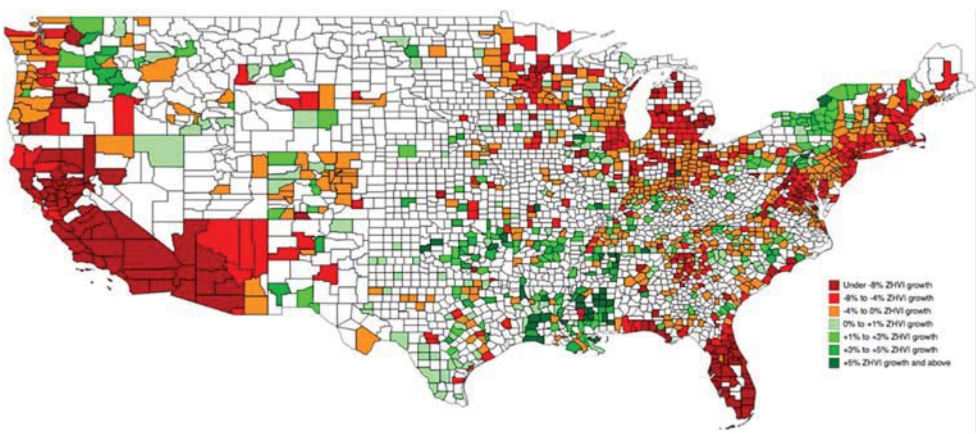

(b) 2007-2009

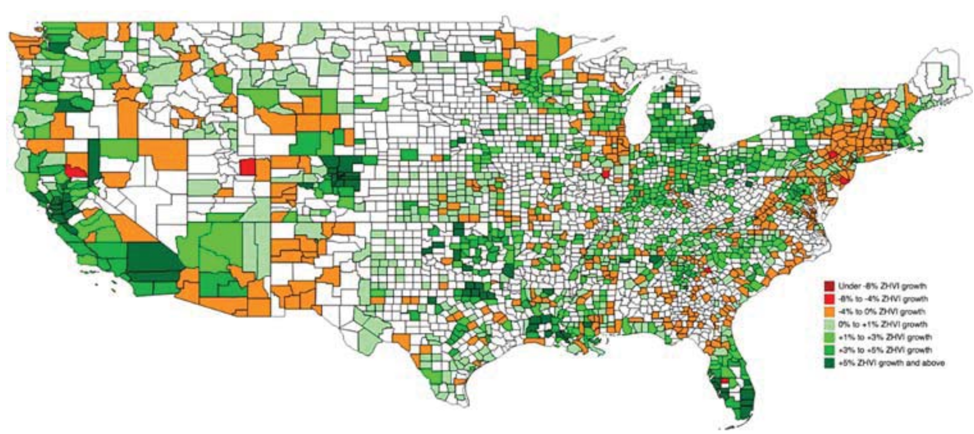

(c) 2010-2016

Figure 11. County-Level Housing Equity and Mortgage Default

The top panel shows average annual county CLTV in 2010. The bottom panel shows average serious mortgage delinquency rate in 2010 in a county. *Source:* Data comes from the Equifax representative panel of 10% of US population.

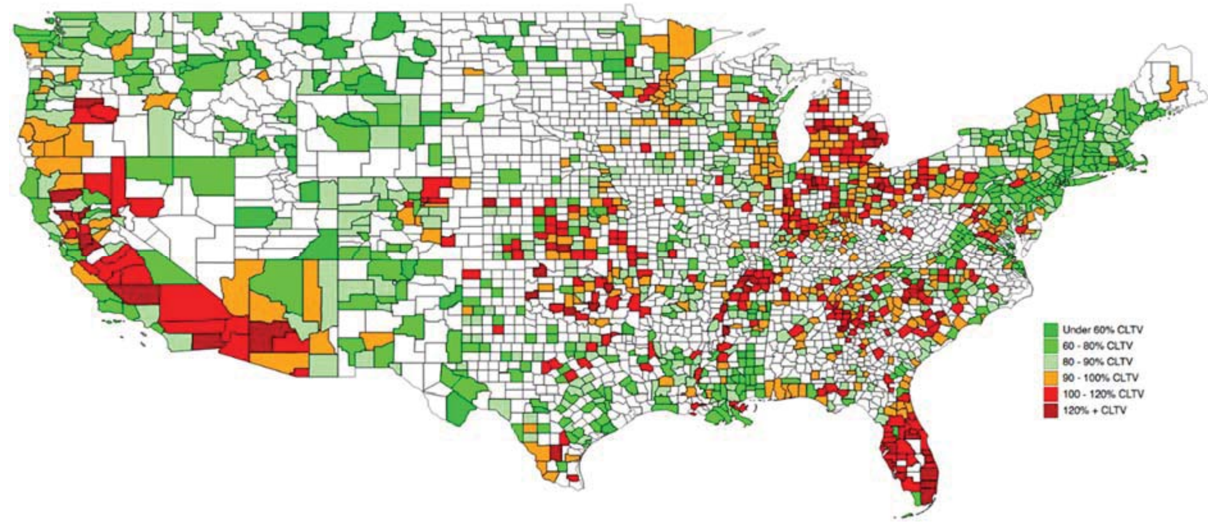

(a) CLTV

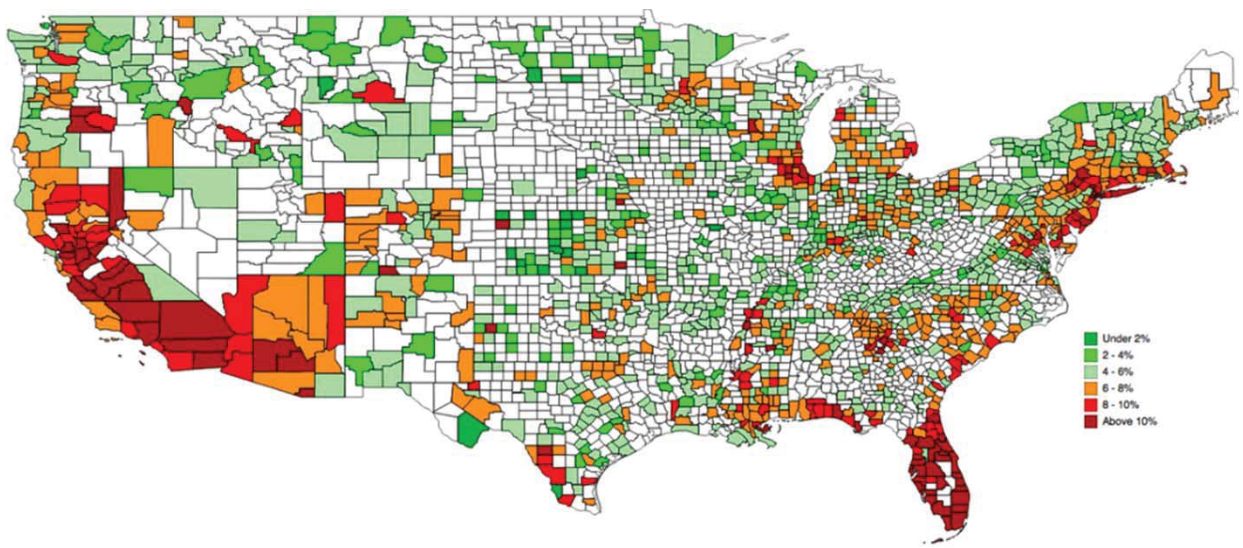

(b) Mortgage Default

Figure 12. ZIP Code-Level House Price Growth

The top panel shows average annual zip-code level house price growth from 2005 to 2006, the middle panel shows house price growth from 2007 to 2009 and the bottom panel shows house price growth from 2010 to 2016. *Source:* Zillow.

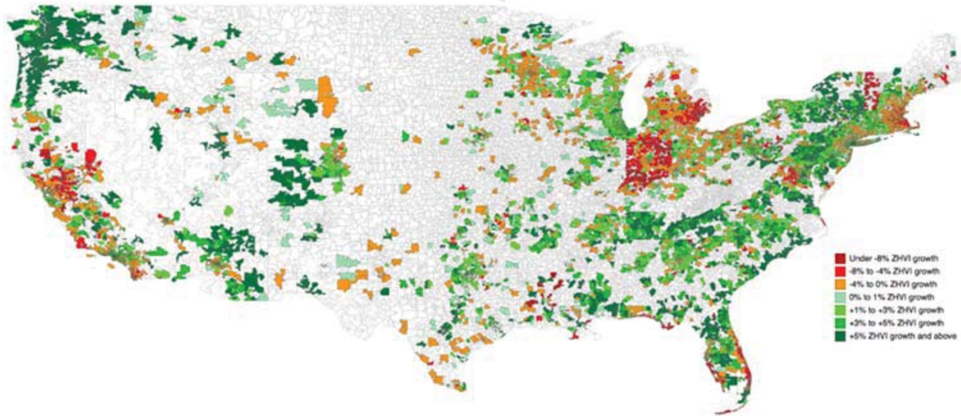

(a) 2005-2006

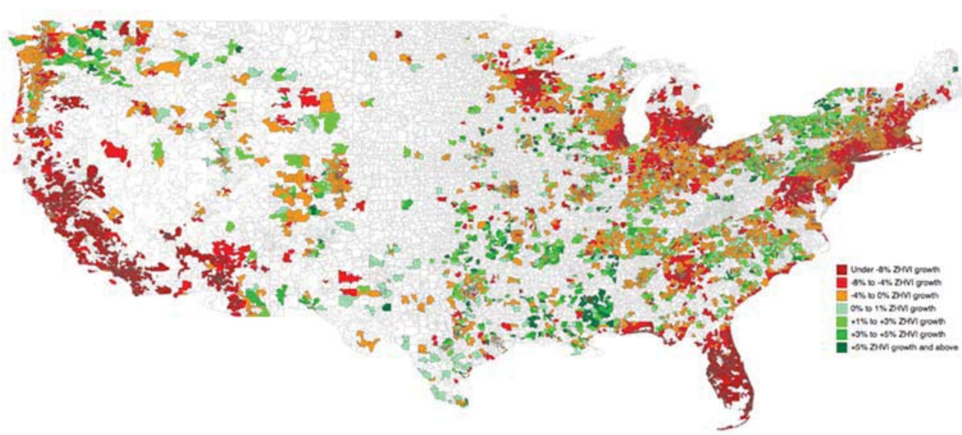

(b) 2007-2009

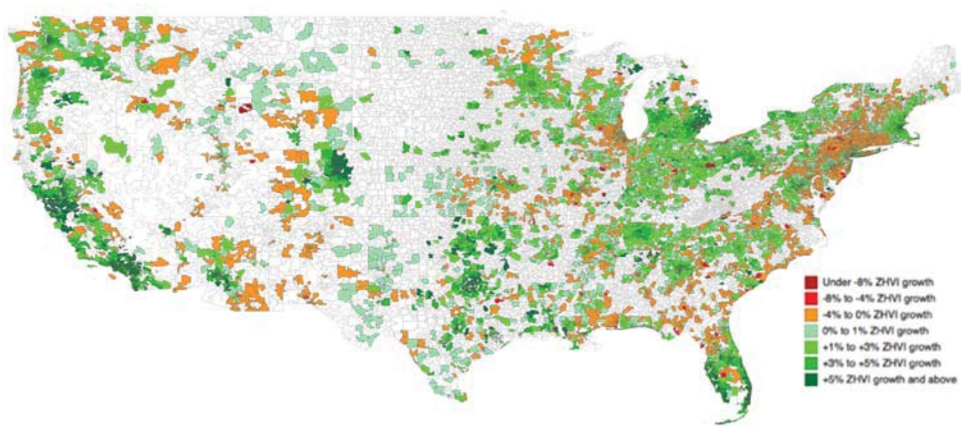

(c) 2010-2016

Figure 13. ZIP Code-Level Housing Equity and Mortgage Default

The top panel shows average annual zip code CLTV in 2010. The bottom panel shows average serious mortgage delinquency rate in 2010 in a zip code. *Source:* Equifax.

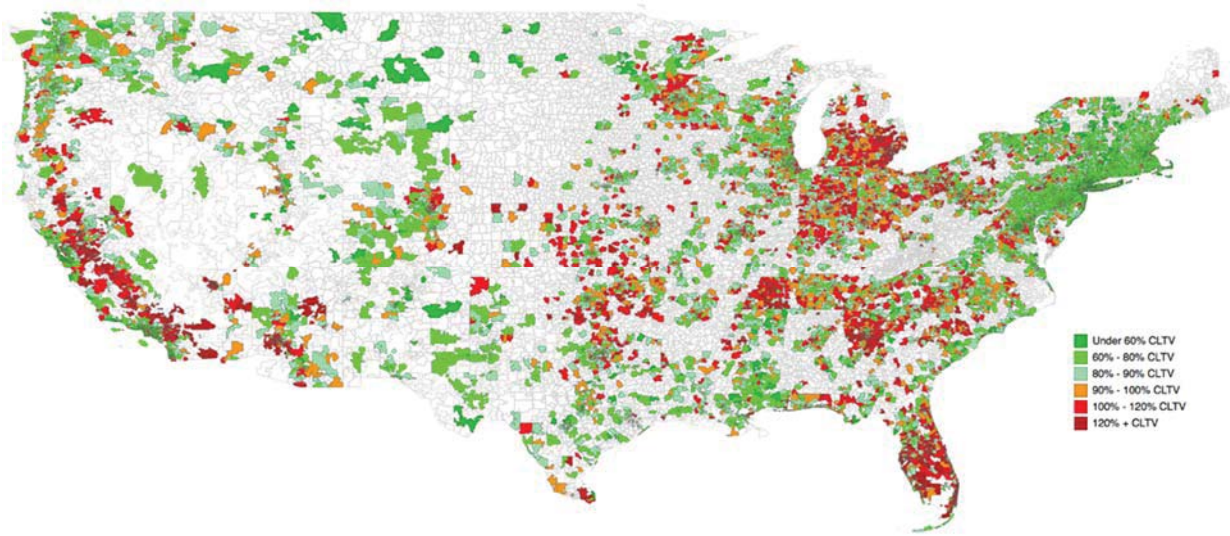

(a) CLTV

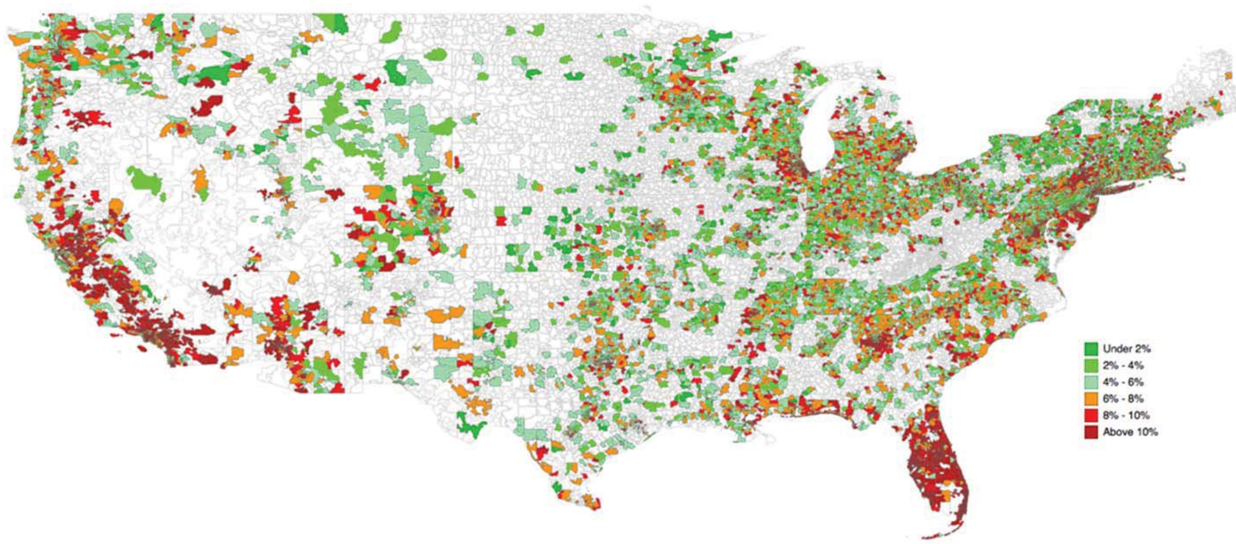

(b) Mortgage Default

Figure 14. The Relationship between County-Level Mortgage Default, House Prices, and Unemployment

This figure examines the relationships between county level variables through simple linear regressions. Dependent variables are regressed on county level independent variables that are interacted with annual dummies for each year from 2005 to 2016. Regressions are population weighted by county. Panel (a) shows the estimated relationship between annual change in the serious mortgage delinquency rate and annual change in unemployment rate along with 95% confidence interval. Panel (b) shows the estimated relationship between annual change in the delinquency rate and house price growth rate along with 95% confidence interval. *Sources:* U.S. Bureau of Labor Statistics, Local Area Unemployment Statistics; Equifax; Zillow.

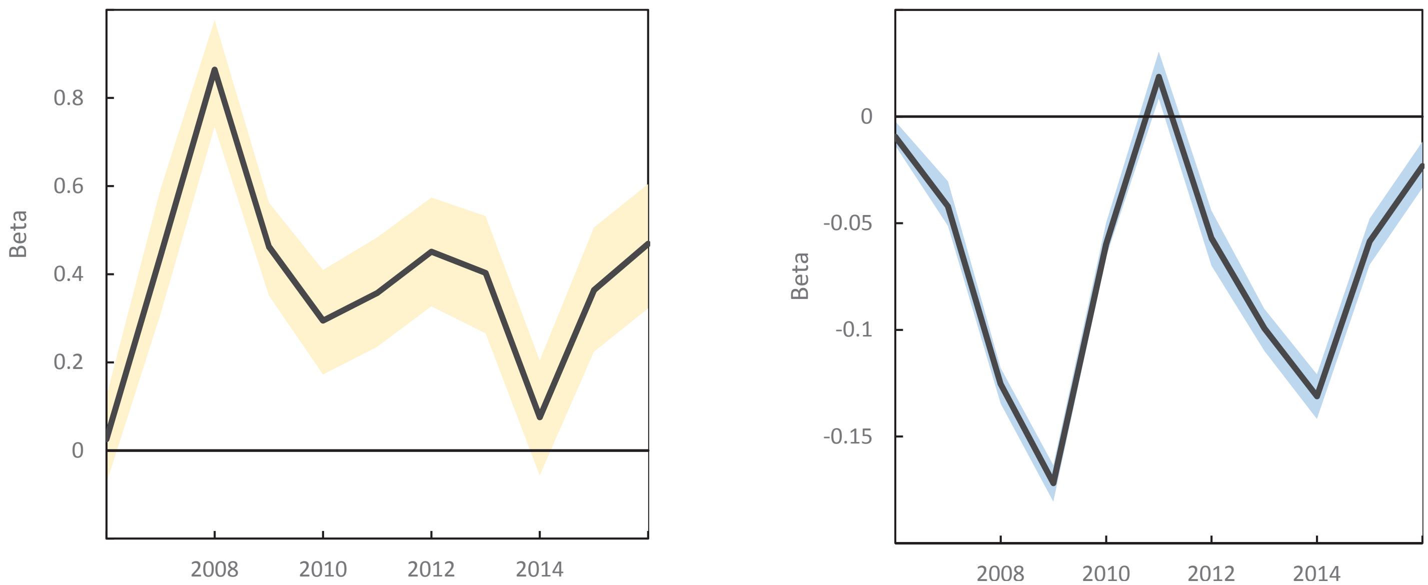

(a) Mortgage Default and Unemployment Relationship

(b) Mortgage Default and House Price Relationship

Figure 15. The Relationship between County-Level House Prices, Housing Equity, and Unemployment

This figure examines the relationships between county level variables through simple linear regressions. Dependent variables are regressed on county level independent variables that are interacted with annual dummies for each year. Regressions are population weighted by county. Panel (a) shows the estimated relationship between annual growth in house prices and unemployment rate change along with 95% confidence interval for years 1997 to 2016. Panel (b) shows the estimated relationship between annual change in the CLTV and annual change unemployment rate along with 95% confidence interval for years 2006 to 2017. *Sources:* U.S. Bureau of Labor Statistics, Local Area Unemployment Statistics; Equifax; Zillow.

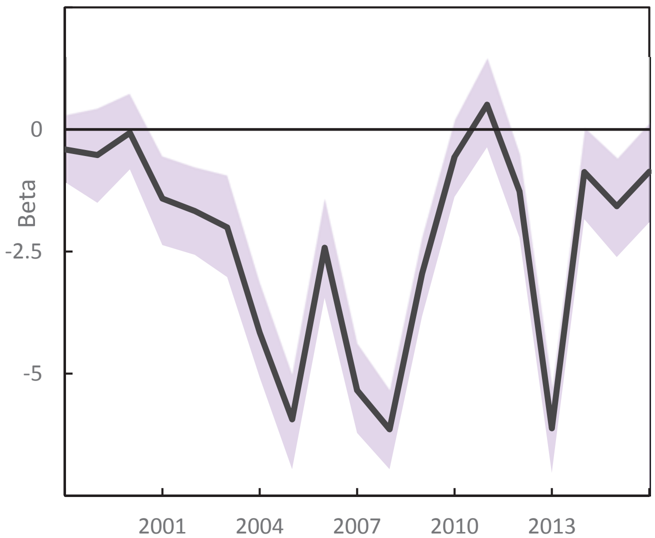

(a) House Prices and Unemployment Relationship

(b) Change in CLTV and Change in Unemployment Relationship

Figure 16. Spatial Variation in the Implementation of Debt Relief

The top panel of this figure shows the spatial variation of zip code ARM share in the US. We note that ARM loans can experience a quick "automatic" pass through of low interest rates. This share, however, needs to be interpreted with caution as many subprime ARM contracts feature various caps and floors that may limit the extent of adjustment of their rates. The middle panel shows the share of loans in a zip code serviced by intermediaries with low organizational capacity to service and modify loans. The bottom panel shows the share of loans in a zip code that were eligible for HARP based on their LTV level and the presence of GSE guarantee. Sources: DiMaggio et al. (2017); Agarwal et al. (2017a, 2017b).

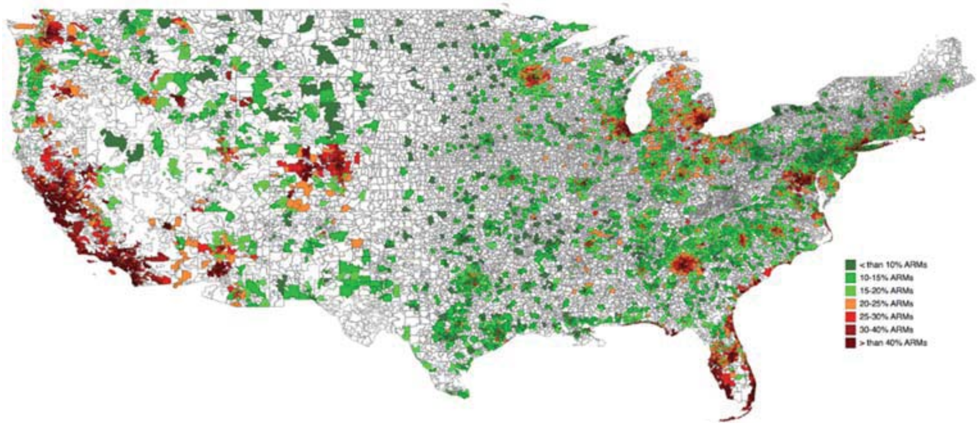

(a) ARM Share

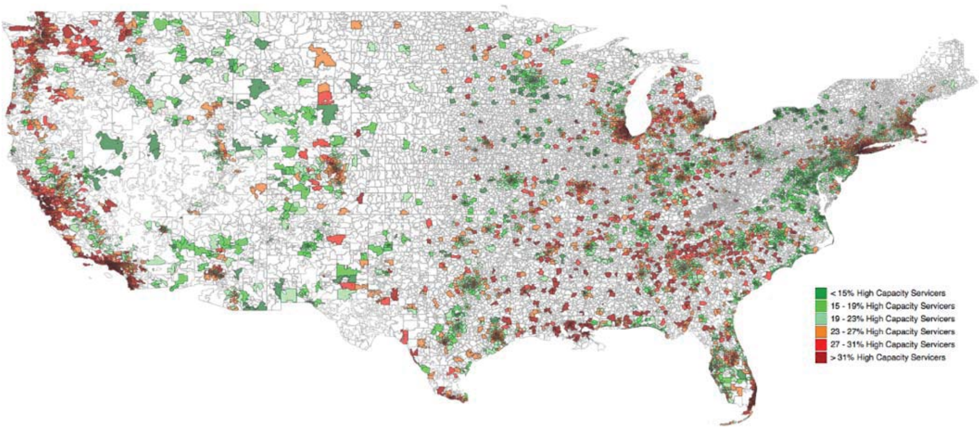

(b) Share of Loans Serviced by High Organizational Capacity Intermediaries

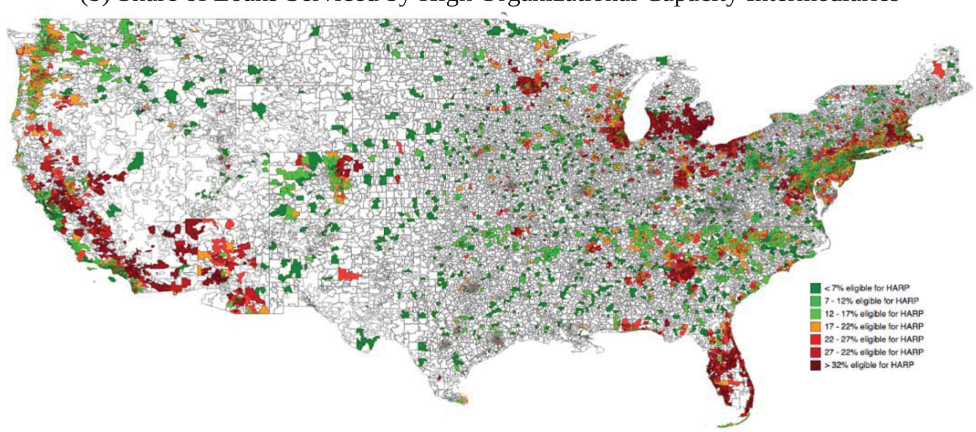

(c) HARP Eligible Share

Figure 17. The Impact of Adjustable-Rate Mortgage Resets on Defaults: Mid 2006-Early 2008

The top panel shows the mortgage default rate for two-year, subprime, adjustable-rate mortgages mainly originated during 2004-06. The loans faced a fixed initial rate for the first two years and subsequently were reset to a variable rate based on a short-term interest rate index. The vertical line marks two years. The bottom panel shows the mortgage default rate for three-year, subprime, adjustable-rate mortgages mainly originated between March 2004 and January 2005. The loans faced a fixed initial rate for the first three years and subsequently were reset to a variable rate based on a short-term interest rate index. The vertical line marks three years. Sources: BlackBox Logic.

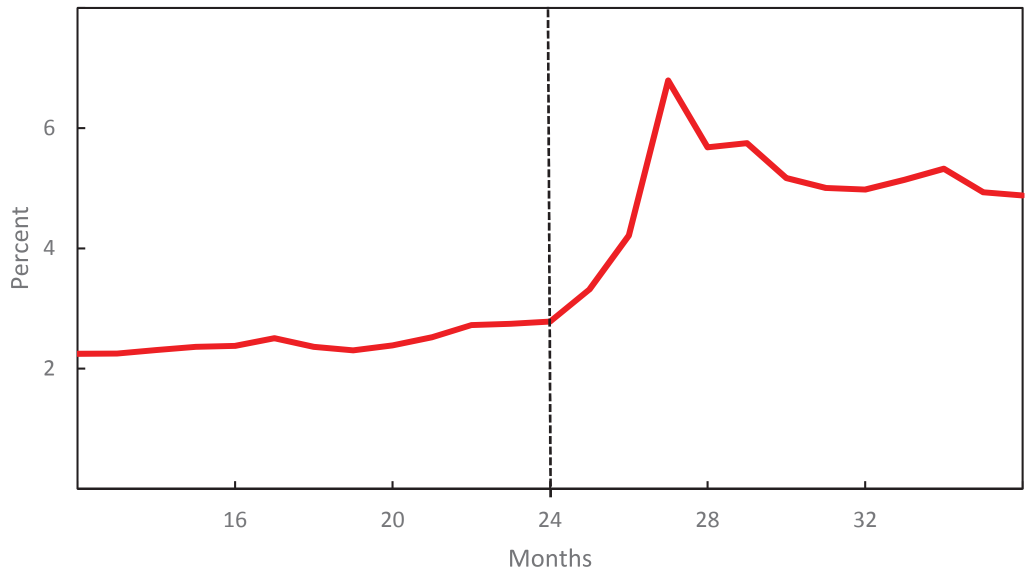

(a) Mortgage default rate (2-year ARMs)

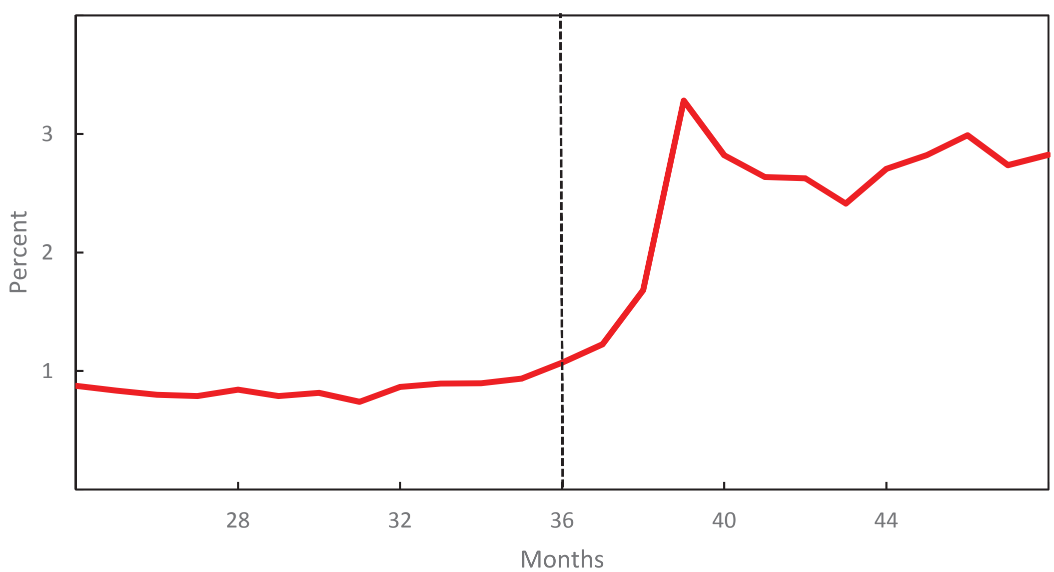

(b) Mortgage default rate (3-year ARMs)

References

Adelino, Manuel, Kristopher Gerardi, and Paul S. Willen. 2014. “Why Don’t Lenders Renegotiate More Home Mortgages? Redefaults, Self-Cures and Securitization.” *Journal of Monetary Economics* 60, no. 7: 835–53.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino. 2016. “Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class.” *Review of Financial Studies* 29, no. 7: 1635–70.

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, and Douglas D. Evanoff. 2011. “The Role of Securitization in Mortgage Renegotiation.” *Journal of Financial Economics* 102, no. 3: 559–78.

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru. 2017a. “Policy Intervention in Debt Renegotiation: Evidence from the Home Affordable Modification Program.” *Journal of Political Economy* 125, no. 3: 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 Refinance Program.” Working Paper no. 21512. Cambridge, Mass.: National Bureau of Economic Research.

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, no. 1: 129–90.

Andersen, Steffen, John Y. Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai. 2014. “Inattention and Inertia in Household Finance: Evidence from the Danish Mortgage Market.” Working Paper no. 21386. Cambridge, Mass.: National Bureau of Economic Research.

Auclert, Adrien. 2017. “Monetary Policy and the Redistribution Channel.” Working Paper no. 23451. Cambridge, Mass.: National Bureau of Economic Research.

Ben-David, Itzhak. 2011. “Financial Constraints and Inflated Home Prices during the Real-Estate Boom.” *American Economic Journal: Applied Economics* 3, no. 3: 55–87.

Benmelech, Efraim, Ralf R. Meisenzahl, and Rodney Ramcharan. 2017. “The Real Effects of Liquidity during the Financial Crisis: Evidence from Automobiles.” *Quarterly Journal of Economics* 132, no. 1: 317–65.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra. 2017. “Regional Heterogeneity and Monetary Policy.” Working Paper no. 23270. Cambridge, Mass.: National Bureau of Economic Research.

Berger, David, Nicholas Turner, and Eric Zwick. 2016. “Stimulating Housing Markets.” Working Paper no. 22903. Cambridge, Mass.: National Bureau of Economic Research.

Bhutta, Neil, and Benjamin J. Keys. 2016. “Interest Rates and Equity Extraction during the Housing Boom.” *American Economic Review* 106, no. 7: 1742–74.

Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru. 2017. "Fintech, Regulatory Arbitrage, and the Rise of Shadow Banks." Forthcoming in the *Journal of Financial Economics*.

Campbell, John Y. 2013. "Mortgage Market Design." *Review of Finance* 17, no. 1: 1–33.

Campbell, John Y., and João F. Cocco. 2003. "Household Risk Management and Optimal Mortgage Choice." *Quarterly Journal of Economics* 118, no. 4: 1449–94.

Campbell, John Y., and João F. Cocco. 2015. "A Model of Mortgage Default." *Journal of Finance* 70, no. 4: 1495–554.

Campbell, John Y., Stefano Giglio, and Parag Pathak. 2011. "Forced Sales and House Prices." *American Economic Review* 101, no. 5: 2108–31.

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

Caplin, Andrew, Noël Cunningham, Mitchell Engler, and Frederick Pollack. 2008. "Facilitating Shared Appreciation Mortgages to Prevent Crashes and Affordability Crises." Discussion Paper no. 2008-12. Brookings, Hamilton Project.

Chen, Hui, Michael Michaux, and Nikolai Roussanov. 2013. "Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty." Working Paper no. 19421. Cambridge, Mass.: National Bureau of Economic Research.

DeFusco, Anthony, Stephanie Johnson, and John Mondragon. 2017. "Regulating Household Leverage." Working paper.

DeFusco, Anthony A., Charles G. Nathanson, and Eric Zwick. 2017. "Speculative Dynamics of Prices and Volume." Working Paper no. 23449. Cambridge, Mass.: National Bureau of Economic Research.

Di Maggio, Marco, Amir Kermani, Benjamin J. 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, no. 11: 3550–88.

Di Maggio, Marco, Amir Kermani, and Christopher Palmer. 2016. "How Quantitative Easing Works: Evidence on the Refinancing Channel." Working Paper no. 22638. Cambridge, Mass.: National Bureau of Economic Research.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl. 2017. "The Deposits Channel of Monetary Policy." *Quarterly Journal of Economics* 132, no. 4: 1819–76.

Dunn, Kenneth B., and Chester S. Spatt. 1985. "An Analysis of Mortgage Contracting: Prepayment Penalties and the Due-on-Sale Clause." *Journal of Finance* 40, no. 1: 293–308.

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

Favilukis, Jack, Sydney C. Ludvingson, and Stijn Van Nieuwerburgh. 2017. "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium." *Journal of Political Economy* 125, no. 1: 140–223.

Foote, Christopher L., Kristopher Gerardi, and Paul S. Willen. 2008. “Negative Equity and Foreclosure: Theory and Evidence.” *Journal of Urban Economics* 64, no. 2: 234–45.

Fuster, Andreas, Laurie Goodman, David Lucca, Laurel Madar, Linsey Molloy, and Paul Willen. 2013. “The Rising Gap between Primary and Secondary Mortgage Rates.” *Economic Policy Review* 19, no. 2: 17–39.

Fuster, Andreas, Stephanie H. Lo, and Paul S. Willen. 2017. “The Time-Varying Price of Financial Intermediation in the Mortgage Market.” Staff Report no. 805. Federal Reserve Bank of New York.

Fuster, Andreas, and Paul S. Willen. 2017. “Payment Size, Negative Equity, and Mortgage Default.” *American Economic Journal: Economic Policy* 9, no. 4: 167–91.

Ganong, Peter, and Pascal Noel. 2017. “The Effect of Debt on Default and Consumption: Evidence from Housing Policy in the Great Recession.” Working paper.

Glaeser, Edward L., Joseph Gyourko, and Albert Saiz. 2008. “Housing Supply and Housing Bubbles.” *Journal of Urban Economics* 64, no. 2: 198–217.

Green, Richard K., and Susan M. Wachter. 2005. “The American Mortgage in Historical and International Context.” *Journal of Economic Perspectives* 19, no. 4: 93–114.

Greenwald, Daniel L. 2018. “The Mortgage Credit Channel of Macroeconomic Transmission.” Working paper.

Greenwald, Daniel L., Tim Landvoigt, and Stijn Van Nieuwerburgh. 2018. “Financial Fragility with SAM?” Working paper.

Griffin, John M., and Gonzalo Maturana. 2016. “Who Facilitated Misreporting in Securitized Loans?” *Review of Financial Studies* 29, no. 2: 384–419.

Gupta, Arpit. 2018. “Foreclosure Contagion and the Neighborhood Spillover Effects of Mortgage Defaults.” Forthcoming in the *Journal of Finance.*

Guren, Adam M., Arvind Krishnamurthy, and Timothy J. McQuade. 2017. “Mortgage Design in an Equilibrium Model of the Housing Market.” Working Paper no. 24446. Cambridge, Mass.: National Bureau of Economic Research.

Guerrieri, Veronica, and Harald Uhlig. 2016. “Housing and Credit Markets: Booms and Busts.” In *Handbook of Macroeconomics, Volume 2B*, edited by John B. Taylor and Harald Uhlig. Amsterdam: North-Holland.

Hartman-Glaser, Barney, and Benjamin Hébert. 2017. “The Insurance Is the Lemon: Failing to Index Contracts.” Working Paper no. 3569. Stanford, Calif.: Stanford University, Graduate School of Business.

Hsu, Joanne W., David A. Matsa, and Brian T. Melzer. 2018. “Unemployment Insurance as a Housing Market Stabilizer.” *American Economic Review* 108, no. 1: 49–81.

Hurst, Erik, Benjamin J. Keys, Amit Seru, and Joseph Vavra. 2016. “Regional Redistribution through the US Mortgage Market.” *American Economic Review* 106, no. 10: 2982–3028.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante. 2017. “The Housing Boom and Bust: Model Meets Evidence.” Working Paper no. 23694. Cambridge, Mass.: National Bureau of Economic Research.

Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig. 2010. “Did Securitization Lead to Lax Screening? Evidence from Subprime Loans.” *Quarterly Journal of Economics* 125, no. 1: 307–62.

Keys, Benjamin J., Tomasz Piskorski, Amit Seru, and Vikrant Vig. 2013. “Mortgage Financing in the Housing Boom and Bust.” In *Housing and the Financial Crisis*, edited by Edward L. Glaeser and Todd Sinai. University of Chicago Press.

Keys, Benjamin J., Devin G. Pope, and Jaren C. Pope. 2016. “Failure to Refinance.” *Journal of Financial Economics* 122, no. 3: 482–99.

Kruger, Samuel, 2017. “The Effect of Mortgage Securitization on Foreclosure and Modification.” Forthcoming in the *Journal of Financial Economics*.

Kung, Edward. 2015. “Mortgage Market Institutions and Housing Market Outcomes.” Working paper.

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider. 2015. “Housing Market(s) of San Diego.” *American Economic Review* 105, no. 4: 1371–407.

Maturana, Gonzalo. 2017. “When Are Modifications of Securitized Loans Beneficial to Investors?” *Review of Financial Studies* 30, no. 11: 3824–57.

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, no. 9: 2830–57.

Mayer, Christopher, Karen Pence, and Shane M. Sherlund. 2009. “The Rise in Mortgage Defaults.” *Journal of Economic Perspectives* 23, no. 1: 27–50.

Mayer, Christopher J., Tomasz Piskorski, and Alexei Tchistyi. 2013. “The Inefficiency of Refinancing: Why Prepayment Penalties Are Good for Risky Borrowers.” *Journal of Financial Economics* 107, no. 2: 694–714.

Melzer, Brian T. 2017. “Mortgage Debt Overhang: Reduced Investment by Homeowners at Risk of Default.” *Journal of Finance* 72, no. 2: 575–612.

Mian, Atif, and Amir Sufi. 2009. “The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis.” *Quarterly Journal of Economics* 124, no. 4: 1449–96.

Mian, Atif, and Amir Sufi. 2011. “House Prices, Home Equity–Based Borrowing, and the US Household Leverage Crisis.” *American Economic Review* 101, no. 5: 2132–56.

Mian, Atif, and Amir Sufi. 2012. “The Effects of Fiscal Stimulus: Evidence from the 2009 Cash for Clunkers Program.” *Quarterly Journal of Economics* 127, no. 3: 1107–42.

Mian, Atif, and Amir Sufi. 2014a. *House of Debt: How They (and You) Caused the Great Recession, and How We Can Prevent It from Happening Again*. University of Chicago Press.

Mian, Atif, and Amir Sufi. 2014b. “What Explains the 2007–2009 Drop in Employment?” *Econometrica* 82, no. 6: 2197–223.

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland. 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103, no. 6: 2530–53.

Piskorski, Tomasz, Amit Seru, and Vikrant Vig. 2010. "Securitization and Distressed Loan Renegotiation: Evidence from the Subprime Mortgage Crisis." *Journal of Financial Economics* 97, no. 3: 369–97.

Piskorski, Tomasz, Amit Seru, and James Witkin. 2015. "Asset Quality Misrepresentation by Financial Intermediaries: Evidence from the RMBS Market." *Journal of Finance* 70, no. 6: 2635–78.

Piskorski, Tomasz, and Alexei Tchistyi. 2010. "Optimal Mortgage Design." *Review of Financial Studies* 23, no. 8: 3098–140.

Piskorski, Tomasz, and Alexei Tchistyi. 2011. "Stochastic House Appreciation and Optimal Mortgage Lending." *Review of Financial Studies* 24, no. 5: 1407–46.

Piskorski, Tomasz, and Alexei Tchistyi. 2017. "An Equilibrium Model of Housing and Mortgage Markets with State-Contingent Lending Contracts." Working Paper no. 23452. Cambridge, Mass.: National Bureau of Economic Research.

Purnanandam, Amiyatosh. 2011. "Originate-to-Distribute Model and the Subprime Mortgage Crisis." *Review of Financial Studies* 24, no. 6: 1881–915.

Rajan, Uday, Amit Seru, and Vikrant Vig. 2015. "The Failure of Models that Predict Failure: Distance, Incentives, and Defaults." *Journal of Financial Economics* 115, no. 2: 237–60.

Scharfstein, David S., and Adi Sunderam. 2016. "Market Power in Mortgage Lending and the Transmission of Monetary Policy." Working paper.

Sinai, Todd. 2013. "House Price Moments in Boom–Bust Cycles." In *Housing and the Financial Crisis*, edited by Edward L. Glaeser and Todd Sinai. University of Chicago Press.

Shiller, Robert J. 2008. *The Subprime Solution: How Today's Global Financial Crisis Happened, and What to Do About It*. Princeton University Press.

Wong, Arlene. 2018. "Transmission of Monetary Policy to Consumption and Population Aging." Working paper.