How Do Foreclosures Exacerbate Housing Downturns?

Adam M. Guren and Timothy J. McQuade*
Boston University and NBER; Stanford University

November 19, 2019

Abstract

This paper uses a structural model to show that foreclosures played a crucial role in exacerbating the recent housing bust and to analyze foreclosure mitigation policy. We consider a dynamic search model in which foreclosures freeze the market for non-foreclosures and reduce price and sales volume by eroding lender equity, destroying the credit of potential buyers, and making buyers more selective. These effects cause price-default spirals that amplify an initial shock and help the model fit both national and cross-sectional moments better than a model without foreclosure. When calibrated to the recent bust, the model reveals that the amplification generated by foreclosures is significant: Ruined credit and choosey buyers account for 25.4 percent of the total decline in non-distressed prices and lender losses account for an additional 22.6 percent. For policy, we find that principal reduction is less cost effective than lender equity injections or introducing a single seller that holds foreclosures off the market until demand rebounds. We also show that policies that slow down the pace of foreclosures can be counterproductive.

Journal of Economic Literature Classification: E30, R31.

Keywords: Housing Prices & Dynamics, Foreclosures, Search, Great Recession.

---

*guren@bu.edu, tmcquade@stanford.edu. We would like to thank John Campbell, Edward Glaeser, and Jeremy Stein for outstanding advice and Pat Bayer, Karl Case, Raj Chetty, Darrell Duffie, Emmanuel Farhi, Simon Gilchrist, Erik Hurst, Lawrence Katz, Arvind Krishnamurthy, Greg Mankiw, Chris Mayer, Monika Piazzesi, Andrei Shleifer, Alp Simsek, Johannes Stroebel, and seminar participants at Harvard University, Stanford University, the NBER Summer Institute, the Penn Search and Matching Workshop, and the Greater Boston Urban and Real Estate Economics Seminar for helpful comments. We would like to acknowledge CoreLogic and particularly Kathryn Dobbyn for providing data and answering numerous questions.

Foreclosures were one of the dominant features of the U.S. housing downturn. From 2006 through 2013, approximately eight percent of the owner-occupied housing stock experienced a foreclosure.[1] Although the wave of foreclosures has subsided, understanding the role of foreclosures in housing downturns remains an important part of reformulating housing policy.

The behavior of the housing market concurrent with the wave of foreclosures is shown in Figure 1. Real Estate Owned (REO) sales — sales of foreclosed homes by lenders — made up over 20 percent of existing home sales nationally for four years. Non-foreclosure sales volume fell 65 percent as time to sale rose. Prices dropped considerably, with aggregate price indices plunging a third and non-distressed prices falling by a quarter.

This paper uses a structural model to argue that quantitatively matching national and cross-sectional features of the bust requires that foreclosures play a significant role in exacerbating the bust. Indeed, in our calibrated model we find that the effect of having more homeowners with a foreclosure flag on their credit record and increased buyer choosiness due to the presence of foreclosures together account for 25.4 percent of the total decline in non-distressed prices and that the reduction in lending stemming from default-induced lender losses accounts for an additional 22.6 percent. Our model implies that foreclosures have far greater equilibrium effects than those found in micro-econometric studies, which use highly localized fixed effects and absorb much of the city-level variation reduction in supply and demand that is at the heart of our structural analysis. The role played by foreclosures opens the door to various foreclosure mitigation policies, and we use our model to compare several such policies. Our policy analysis reveals that both lender equity injections and a government facility to strategically hold foreclosures off the market are highly effective, while principal reductions similar to those pursued by the Home Affordable Modification Program (HAMP) are less cost-effective. Furthermore, if household liquidity shocks are persistent, slowing down foreclosures can prolong the bust to the point that doing so becomes counterproductive.

We analyze an equilibrium search model of the housing market with random moving shocks, undirected search, idiosyncratic house valuations, Nash bargaining over price, and endogenous conversion of owner-occupied homes to rental homes. The model builds on a literature on search frictions in the housing market, notably Wheaton (1990), Williams (1995), Krainer (2001), Novy-Marx (2009), Ngai and Tenreryo (2014), and Head et al. (2014). We add a mortgage market with a representative competitive lender facing a regulatory capital constraint and costly equity issuance, as well as mortgage default, whereby underwater homeowners default if hit by a liquidity shock, to this workhorse model. The mortgage sector in our model is related to recent research that models the interaction of the banking

[1] All figures are based on data from CoreLogic described in detail below.

Figure 1: The Role of Foreclosures in the Housing Downturn

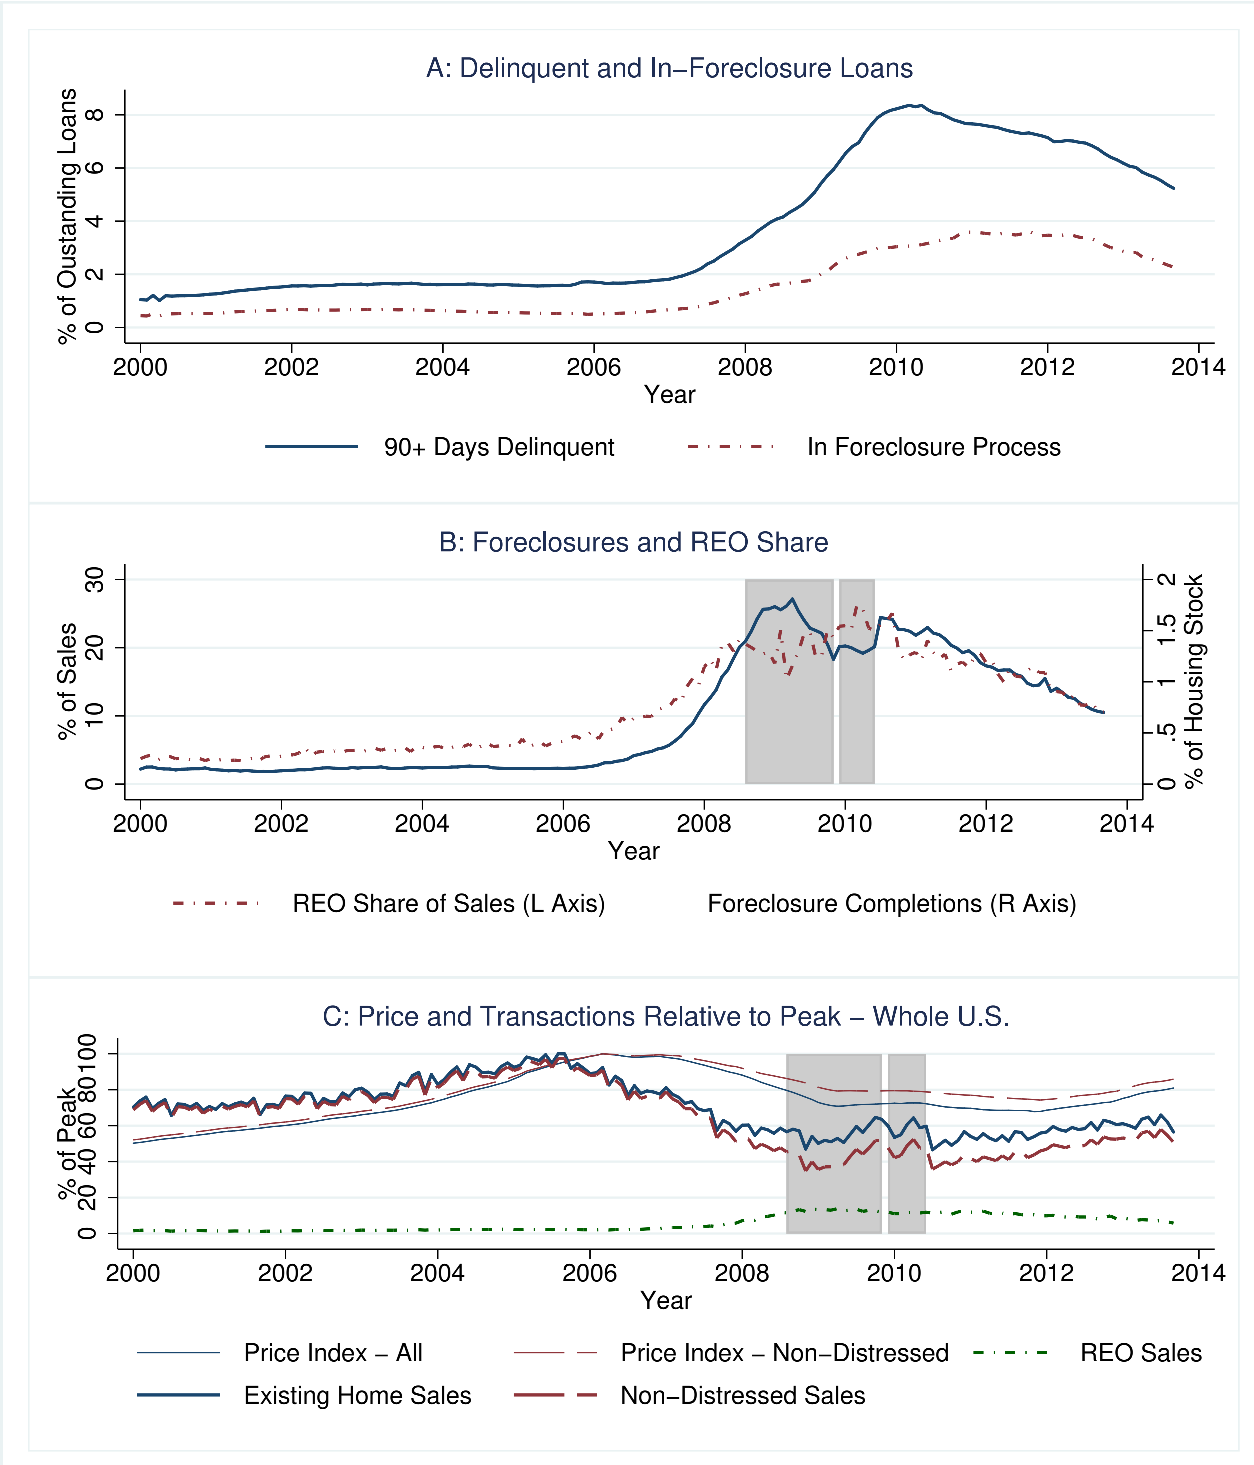

Notes: All data is seasonally adjusted national-level data from CoreLogic as described in the appendix. The grey bars in panels B and C show the periods in which the new homebuyer tax credit applied. In panel C, all sales counts are unsmoothed and normalized by the maximum monthly existing home sales while each price index is normalized by its separate maximum value.

sector and housing market, including Elenev et al. (2016), Greenwald et al. (2018), Hedlund (2016a, 2016b), Hendlund and Garriga (2019), and Paixao (2017).

Foreclosures have three distinguishing characteristics in our model: They cause losses for lenders and reduce lender equity, REO sellers have higher holding costs, and individuals who are foreclosed upon cannot immediately buy a new house. We show that foreclosures dry up the market for non-distressed sales and reduce volume and prices through three main equilibrium channels. First, reduced equity pushes lenders against their capital constraint and causes them to ration mortgage credit. This prevents some buyers from being pre-approved for a loan, which we call the lender rationing effect. Second, because foreclosed upon homeowners are for a time prevented from purchasing due to the foreclosure flag on their credit, foreclosures further reduce the number of buyers in the market relative to the number of sellers. We call this the foreclosure flag effect. Both the lender rationing effect and foreclosure flag effect result in an imbalance of buyers and sellers, reducing the probability that a seller contacts a buyer and lowering equilibrium prices. Typically, foreclosures are thought of as an expansion in supply; our model emphasizes that foreclosures also reduce demand. Third, the presence of distressed sellers increases the outside option to transacting for buyers, who have an elevated probability of being matched with a distressed seller next period and consequently become more selective. This choosey buyer effect endogenizes the degree of substitutability between REO and non-distressed sales.

In conjunction with the effect of foreclosures on prices, default amplifies the effects of a negative shock: An initial shock that reduces prices puts some homeowners under water and triggers foreclosures, which causes more price declines and in turn further default. Lock-in of underwater homeowners also impacts market equilibrium by keeping potential buyers and sellers out of the market, increasing the share of listings which are distressed. Endogenous conversion of owner-occupied units to renter-occupied in response to the increase in demand for rental units provides an important countervailing force.

We calibrate our model to match the nation-wide price decline, sales decline, REO share, and aggregate number of foreclosures from 2006 to 2013. The model fits a number of moments that are not direct calibration targets, including the decline in non-distressed prices. We use our model to quantitatively decompose the sources of the price decline. The choosey buyer and foreclosure flag effects together account for 32.3 percent of the decline in aggregate prices and 25.4 percent of the decline in non-distressed prices. We find that credit rationing associated with weakened lender balance sheets can explain an additional 27.1 percent of the decline in aggregate price indices and an additional 22.6 percent of the decline in non-distressed prices. Only 40.6 percent of the decline in aggregate prices and 52.0 percent of the decline in non-distressed prices is accounted for by an exogenous and permanent shock to aggregate prices that we interpret as a bursting bubble. These amplification effects are far more substantial than those found by micro-econometric studies of foreclosure externalities. Such studies compare the highly local effect of foreclosures on neighboring prices and find minor effects; we find more substantial effects because our structural approach allows us to analyze market-wide equilibrium effects that are absorbed into the constant in micro studies.

To further validate the extent to which foreclosures can account for the bust, we add a parsimonious amount of cross-sectional heterogeneity to our model and evaluate its ability to explain cross-sectional moments relative to a model without default. In particular, we allow cities, formally core-based statistical areas (CBSAs), to differ in their initial loan balance distribution, their unemployment rate, and the size of the preceding boom, assuming that the persistent component of the price decline in the bust is proportional to the boom. The calibrated model fits a number of cross-sectional moments despite the fact that most parameters are calibrated to pre-downturn moments and that these cross-sectional moments are not used in the calibration. In particular, the model does a good job of explaining quantitatively why cities with a larger boom had a more-than-proportionally-larger bust, which our model attributes to foreclosure. By contrast, a model without default does not explain these cross-sectional patterns and has a poorer fit because the size of the bust is proportional to the size of the boom.

Finally, we use the model to quantify the equilibrium impact of a number of government policies aimed at ameliorating the crisis, including principal reduction, equity injections, a facility to purchase foreclosures and hold them off the market until demand rebounds, and regulations to slow down the pace of foreclosures. To our knowledge, the only other papers to analyze foreclosure policy in a quantitative equilibrium model are Hedlund (2016b), who studies making non-recourse mortgages recourse, and Kaplan, Mitman, and Violante (2017), who consider debt forgiveness programs in a model with no price-default spiral. We first show that the effectiveness of slowing down the rate of foreclosure completions depends on the rate at which homeowners cure by regaining the ability to pay the mortgage or becoming above water. If the cure rate is fast enough, slowing down the rate of completions can be effective at limiting the amount of default. However, this policy also lengthens the crisis, which has a negative impact on prices. If the cure rate is sufficiently slow – which is likely the empirically-relevant case in a crisis – the effect of lengthening the crisis can dominate and slowing down the rate of completions can actually exacerbate the crisis. Second, we compare the equilibrium effects of three different government interventions to ameliorate the housing bust that operate on different margins: lender equity injections, homeowner principal reduction, and a government facility to purchase distressed homes, maintain them off the market, and re-introduce them once demand rebounds. We find the government facility to be highly cost-effective since it directly addresses the imbalance of supply and demand caused by foreclosures, which supports higher prices and reduces default. This in turn improves lender balance sheets, which further reduces the severity of the crisis. Government equity injections are also quite effective, while principal reductions have stronger aggregate effects than previous micro-econometric studies would suggest but are still the least-cost-effective of the policies we consider due to imperfect targeting.

Our analysis has a few limitations which one should be mindful of when interpreting the results. First, we consider only liquidity-driven default and do not consider a household's incentives to strategically default. Our model is a good approximation to the 2000s, as the literature has found that strategic default was limited in the crisis.[2] However, the absence of strategic default means that our model is not well suited to consider policies like the complete elimination of foreclosure, which would likely lead to widespread strategic default and cripple lender balance sheets. This pushes us to focus on *ex post* policies that would not dramatically change strategic default incentives if implemented carefully. Second, given our focus on a building a rich, structural framework of housing market dynamics and lender balance sheets, we do not fully model the household's dynamic budget constraint and non-housing consumption. Our approach yields useful insights about several prominent foreclosure policies, but it does limit our ability to consider some others like interest rate reductions (which we do our best to consider in an appendix) or alternate mortgage designs that work through the household budget constraint. Third, because we do not model the household's budget constraint, we abstract from general equilibrium effects such as the impact of tax or debt financing needed to pay for foreclosure policies or the effect of such policies on equilibrium interest rates. Despite these limitations, our framework delivers quantitative insights regarding the mechanisms driving housing crises and offers clean economic intuitions about the relative cost-effectiveness of various policies in stemming a price-default spiral given a particular funding scheme and level of government expenditures.

The remainder of the paper is structured as follows. Section 1 presents facts about the downturn across metropolitan areas. Section 2 introduces our basic model of the housing market, and Section 3 describes the calibration of the model. Section 4 explains and quantifies the forces at work in the national model, and Section 5 evaluates the model's ability to explain cross-sectional moments about the bust. Section 6 considers foreclosure policy, and

[2]Ganong and Noel (2019b) show that essentially all default by non-investors was liquidity triggered. Bhutta et al. (2017) estimate that the median non-prime borrower does not strategically default until their equity falls to negative 74 percent. Similarly, Gerardi et al. (2017) find that there were few strategic defaulters in the PSID as most defaulters do not have the assets to make a mortgage payment and maintain their consumption. The largest estimate of the share of defaults that are strategic is 15 to 20 percent (from Experian Oliver-Wyman). See also Elul et al. (2010) and Foote et al. (2008).

Figure 2: Price and Sales in Selected Cities With High Levels of Foreclosure

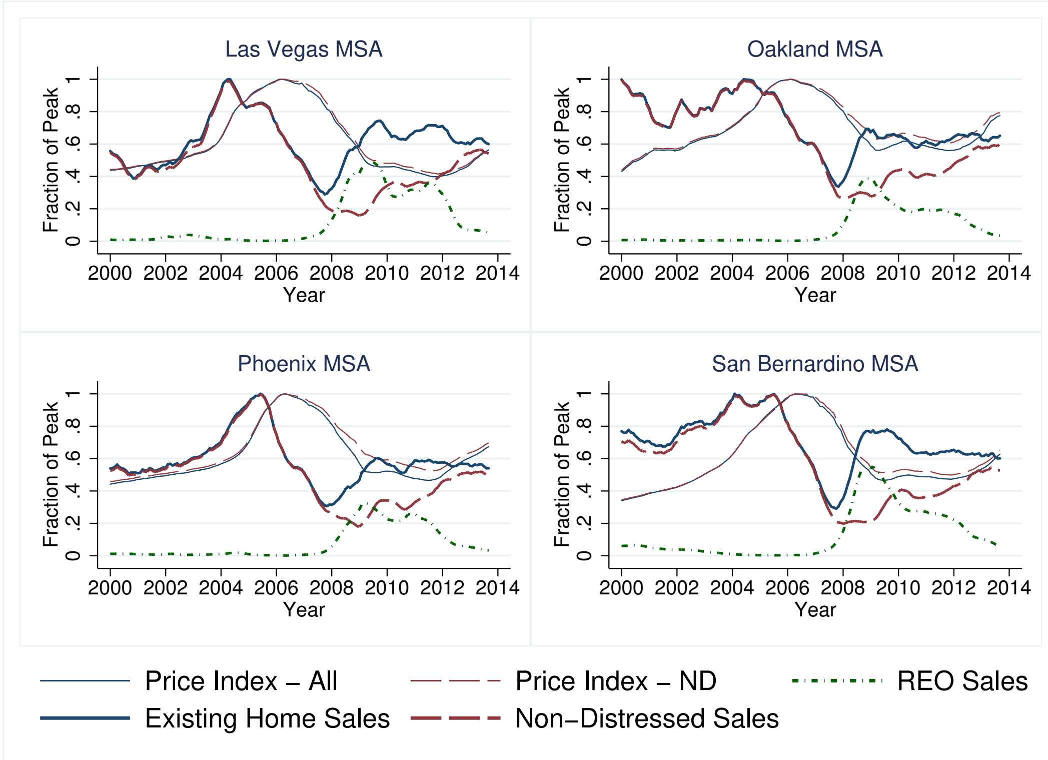

Notes: All data is seasonally adjusted CBSA-level data from CoreLogic as described in the appendix. Sales are smoothed using a moving average and normalized by the maximum monthly existing home sales, while each price index is normalized by its maximum value.

Section 7 concludes.

1 Empirical Facts

The national aggregate time series for price, volume, foreclosures, and REO share presented in Figure 1 mask substantial heterogeneity across metropolitan areas. To illustrate this, Figure 2 shows price and volume for four of the hardest-hit CBSAs. In Las Vegas, for instance, prices fell nearly 60 percent, and the REO share was as high as 75 percent.

To provide a more systematic analysis of the heterogeneity of the bust across cities and to motivate, calibrate, and test the model, we use a proprietary data set provided by CoreLogic supplemented by data from the United States Census. CoreLogic provides monthly data for 2000 to 2013 for the nation as a whole and 99 of the 100 largest CBSAs. The data set includes a repeat sales house price index, a house price index for non-distressed sales only, sales counts for REOs and non-distressed sales, and estimates of quantiles of the LTV distribution. We seasonally adjust the CoreLogic data and smooth the sales count series

Figure 3: Price Boom vs. Price Bust Across Cities

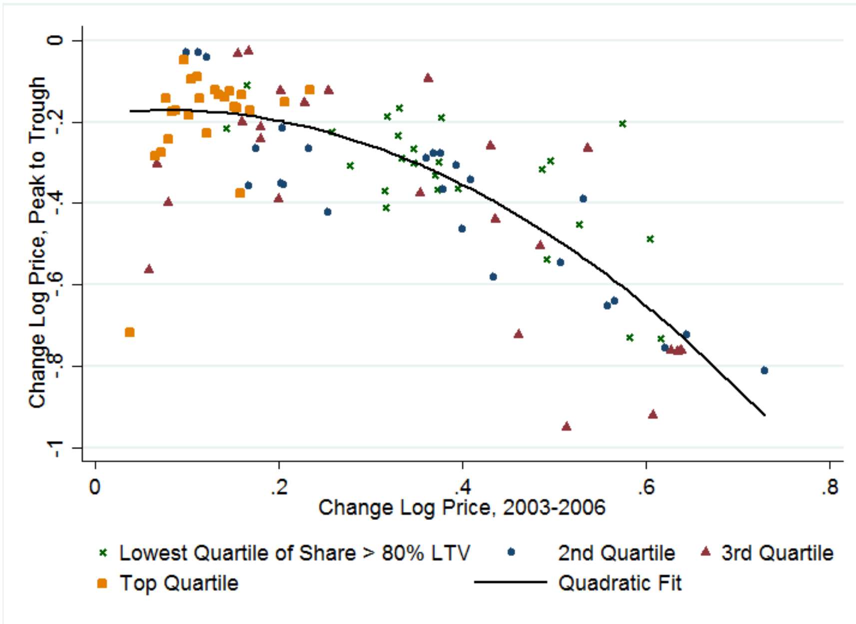

Note: Scatter plot of seasonally adjusted data from CoreLogic along with quadratic regression line that excludes CBSAs in greater Detroit which busted without a boom. The data are described in the appendix. Each data point is an CBSA and is color coded to indicate in which quartile the CBSA falls when CBSAs are sorted by the share of homes with over 80 percent LTV in 2006. Although the highest LTV CBSAs had almost no boom and no bust (e.g. Indianapolis), the CBSAs below the best fit line tend to be CBSAs with a large share of homeowners with high LTVs in 2006 (third quartile).

using a moving average. A complete description of the data and summary statistics are in the appendix.

The best predictor of the size of the bust is the size of the preceding boom. Figure 3 plots the change in log price from 2003 to 2006 against the change in log price from each market's peak to its trough. There is a strong downward relationship, which motivates a key feature of our model: The shock that causes the bust is a fall in home valuations that is assumed to be proportional to the size of the preceding boom.

Figure 3 also reveals a more subtle fact: Metropolitan areas that had a larger boom had a more-than-proportional larger bust. While a linear relationship between log boom size and log bust size has an r-squared of 0.62, adding a quadratic term that allows for larger busts in places with larger booms increases the r-squared to 0.68. The curvature can be seen in the best-fit line in Figure 3.[3] We argue that by exacerbating the downturn in the hardest-hit

[3]For the best fit line, the regressions that follow, and the model calibration, we exclude two outlier CBSAs in southeast Michigan which had a large bust without a preceding boom so that the non-linearity is not

areas, foreclosures explain why places with larger booms had disproportionately larger busts.

This explanation has an important corollary: Because default is predominantly caused by negative equity, a larger bust should occur not only in places with a larger initial negative shock to prices, but particularly in locations with a *combination* of a large shock and a large fraction of houses with high LTVs — and thus close to default — prior to the bust. To provide suggestive evidence that this prediction is borne out in the data, the points in Figure 3 have different shapes for each quartile of the share of homeowners in the CBSA with over 80 percent LTV in 2006. While the largest high LTV shares occurred in places that did not have a bust — home values were not inflated in 2006, so the denominator of LTV was lowest in these locations — one can see that the majority of CBSAs substantially below the quadratic trend line were in the upper end of the high LTV share distribution (triangles in the figure).

To formally investigate whether the interaction of many households with high LTVs and a large preceding boom is correlated with a deep downturn, we estimate:

$$\begin{aligned} Y_i &= \beta_0 + \beta_1 \Delta_{03-06} \log(P_i) + \beta_2 \left[\Delta_{03-06} \log(P_i)\right]^2 \\ &\quad + \beta_3 X_i + \beta_4 \left(\Delta_{03-06} \log(\mathrm{P}_i) \times X_i\right) + \varepsilon_i, \end{aligned} \tag{1}$$

where i indexes CBSAs, X_i is an interacted variable, and the outcome variable Y_i is either the maximum change in log price, the maximum peak-to-trough change in log non-distressed prices, the maximum peak-to-trough REO share, or the fraction of houses that experience a foreclosure. We use two X_is. First, to test whether the combination of a large bust and a large fraction of houses with high LTV creates a particularly large downturn, we use the z score of the share of mortgages with over 80 percent LTV in 2006. This regression is similar in spirit to Lamont and Stein (1999), who show that prices are more sensitive to income shocks in cities with a larger share of high-LTV households. Second, to more directly test the role of foreclosures, for price and non-distressed price we use the *ex-post* fraction of houses that experience a foreclosure. The regression results are shown in Table 1, with summary statistics and robustness checks in the appendix. With the share of mortgages with a high LTV in 2006 as the interacted variable, the fourth row of the first four columns shows the key result: The interaction term between the size of the run-up and the share of high-LTV homeowners is significantly negative for price, non-distressed price, and sales volume and significantly positive for the mean REO share of volume and the fraction of the housing stock that is foreclosed upon. This is consistent with a combination of a steep price run-up and high LTV homeowners triggering a price-default spiral.

---

overstated. All results are robust to including these two CBSAs.

Table 1: Cross-City Regressions on the Impact of the Size of the Boom and Its Interaction With High LTV Share

	1	2	3	4	5	6
Dependent Variable:	$\Delta \log(\mathrm{P})$	$\Delta \log(\mathrm{P}_{\mathrm{nd}})$	Mean REO Share	% Foreclosed	$\Delta \log(P)$	$\Delta \log(P_{nd})$
X (Interacted Variable)	Z LTV > 80%	Z LTV > 80%	Z LTV > 80%	Z LTV > 80%	% Foreclosed	% Foreclosed
$\Delta \log(\mathrm{Price})_{03-06}$	0.251	0.092	-0.531	-0.362	-0.662	-0.758
	(0.414)	(0.321)	(0.273)*	(0.139)**	(0.291)**	(0.225)***
$\Delta \log(\mathrm{Price})^2_{03-06}$	-1.886	-1.537	1.391	0.988	-0.031	0.355
	(0.498)***	(0.382)***	(0.331)***	(0.171)***	(0.504)	(0.392)
X	0.054	0.061	-0.002	-0.013	-1.01	-0.521
	(0.029)*	(0.024)**	(0.021)	(0.013)	(0.381)***	(0.297)*
$\Delta \log(\mathrm{Price})_{03-06} \times X$	-0.310	-0.336	0.205	0.235	-0.631	-1.519
	(0.123)**	(0.113)***	(0.075)***	(0.054)***	(0.905)	(0.711)**
r^2	0.690	0.716	0.467	0.651	0.783	0.815
N	97	97	96	96	98	98

Notes: * = 10% Significance, ** = 5% Significance *** = 1% significance. All standard errors are robust to heteroskedasticity. Each column shows estimates of (1) with the constant suppressed. For the first four columns, the interaction variable X is the z-score for the share of houses in 2006 with an LTV over 80%. In the last two columns, the interaction variable X is the share of the housing stock that experienced a foreclosure. P_{nd} is a non-distressed only price index. The mean REO share of sales volume is the average from 2008 to 2013, and the fraction foreclosed is the fraction of the housing stock foreclosed upon over the first 8 years of the downturn. All data is from CoreLogic and described in the appendix. Data are for 98 largest CBSAs excluding Detroit MI, Warren MI in all columns, Syracuse NY in columns 1-4, and, Birmingham AL for columns 3-4. All dependent variables are peak-to-trough maximums with the exception of percentage foreclosed upon. See the appendix for robustness and regressions including southeast Michigan.

Columns 5 and 6 show both a negative direct effect of foreclosures and a negative effect of the interaction between foreclosures and the size of the run up. While the interaction term is insignificant for the overall price index due to a large standard error, the effect is statistically significant for non-distressed prices. A negative direct effect and interaction is expected for the aggregate price index since foreclosures trade at a discount, but the negative effect on non-distressed prices provides evidence that foreclosures amplify the bust. Finally, adding the interaction with fraction foreclosed to the regression eliminates the strong and negative quadratic term on price, which suggests that the non-linearity in the size of the bust relative to the size of the boom can be accounted for by foreclosures, as will be the case in our model.

2 Model

We construct a a Diamond-Mortensen-Pissarides undirected equilibrium search model of the housing market to quantitatively analyze the effects of foreclosures and policies aimed at ameliorating their effects. Search frictions play an important role in housing markets: Houses are illiquid, most households own one house and move infrequently, buyers and sellers are largely atomistic, and search is costly, time consuming, and random.

We first describe the environment, agents, and shocks that agents receive. We then describe the housing and mortgage markets and define equilibrium.

2.1 Environment, Agents, and Shocks

Time is discrete and indexed by a t subscript, and the discount factor is β. All agents have linear utility. There are a unit mass of individuals who are assumed to be natural homeowners and a unit mass of houses which can either be owned and rented.[4] The model is thus a closed system with a fixed population and housing stock.[5]

Individuals can be in one of four states. A mass l_t of individuals are homeowners, v_t^b are buyers who are searching for a home and renting while they do so, v_t^r are renters who do not have a foreclosure on their record but have not qualified for a mortgage and are waiting to search until they do so, and v_t^f are renting and have a foreclosure on their record. The stock

[4]We focus on natural homeowners, but natural renters and transitions in and out of homeownership can be added without substantially changing the model.

[5]For tractability and to focus the paper, we abstract away from housing supply. In practice, there was little residential construction during the crisis. However, this assumption implies we may miss some features of the crisis, such as depreciation of the existing housing stock or foreclosed homes being demolished. It also means that we will not capture regional heterogeneity in housing supply elasticities when we perform our cross-city analysis.

of houses can have one of three statuses. l_t are owner-occupied, v_t^v are non-owner-occupied and not owned by a lender, and v_t^d are owned by lenders after a foreclosure. Of the v_t^v non-owner-occupied houses that are not owned by a lender, v_t^a are converted to rent temporarily, which precludes them from being listed for sale, and v_t^n are listed for sale. As we describe below, lenders have higher holding costs, and they consequently list all of the homes they own for sale.

Homeowners experience two different shocks. First, they experience moving shocks with probability γ that represent changes in tastes and life events that induce them to leave their house as in Krainer (2001) and Ngai and Tenreyro (2014). We assume that these shocks occur at a constant rate and that only individuals who are mismatched with their house attempt to find a new one. When a moving shock occurs, what happens depends on the homeowner's equity position. If a homeowner has negative equity, which we will define momentarily, she cannot pay off her mortgage balance with the proceeds from a sale, and consequently the homeowner is "locked in" to the current house. We assume the homeowner takes actions to accommodate the mismatch shock given her inability to move and remains a homeowner until she receives another shock.[6] If a homeowner has positive equity, she sells her house, pays off the lender, and attempts to buy another house.

To define a homeowner's equity position, we assume that there exists a competitive fringe of market-markets who pay homeowners a price $\widetilde{V}_t^n$ for the house and then market the home to buyers through a search and matching process described subsequently, similar to Hedlund (2016a).[7] Homeowners thus have negative equity and are locked in if their loan balance $L > \widetilde{V}_t^n$ and have positive equity and sell if $L \leq \widetilde{V}_t^n$. When positive equity homeowners sell their house, they pay off their mortgage and attempt to secure pre-approval for new financing.[8] Pre-approval occurs with equilibrium probability P_t. If the individual receives a pre-approval, she enters the housing market as a buyer. If the individual is unable to secure pre-approval, she becomes a renter and attempts to secure pre-approval again at exogenous rate γ_r. Pre-approval specifies $\Phi = (L, \mu)$, the loan amount and the interest rate at which financing can be secured at the time of purchase. We describe the endogenous pre-approval probability P_t and loan terms Φ when we introduce the mortgage market below.

[6] For parsimony we do not fully model the income and savings of households. Rather, we make the reduced form assumption that underwater homeowners experience lock-in. In practice, homeowners who are only slightly underwater may have sufficient savings to make up the difference.

[7] As an alternative, one can imagine that when a homeowner enters the market, they turn into both a seller and a buyer that are independent of one another as in Ngai and Tenreryo (2014). $\tilde{V}_t^n$ would then reflect the value of having a listing on the market, inclusive of marketing and maintenance costs.

[8] Our modeling of pre-approvals implicitly assumes that households require financing to purchase a home. In practice, few households purchase without a mortgage. Incorporating heterogeneous down payments would significantly increase model complexity while adding little economic insight.

The second type of shock a homeowner may receive is a liquidity shock, which occurs with time-varying probability ι_t. Again, what happens when a homeowner experiences a liquidity shock depends on her equity position. Homeowners with negative equity who experience a liquidity shock default because they are unable to afford their mortgage payments and the price they could receive from a market-maker is insufficient to pay off their mortgage. This "double trigger" default is the only source of default in our model. While "ruthless" or "strategic default" by borrowers – that is default by a homeowner with the liquidity to pay their mortgage – has occurred, there is a consensus in the literature that strategic default accounts for a very small fraction of mortgage defaults.[9] To keep the model tractable and maintain a focus on housing market dynamics, we thus do not model strategic default, nor do we model the strategic decision of the lender to foreclose, modify the loan, rent to the foreclosed-upon homeowner, or pursue a short sale, which are options that were not widely used until late in the crisis. Consequently, we assume that homeowners with $L > \widetilde{V}_t^n$ default if they experience a liquidity shock and enter the foreclosure process.

Homeowners who receive a liquidity shock who are above water, on the other hand, do not default. In our baseline model, we further assume that they can remain in their house despite the liquidity shock. In practice, homeowners with positive equity who receive a liquidity shock have various means to avoid having to sell. For example, borrowers with positive home equity could potentially borrow against it to cover temporary lost income. Homeowners with positive equity could also pursue a refinancing or a term extension. For parsimony, we do not fully model these options and instead make the reduced-form assumption that liquidity shocks do not force sales for households with positive equity. In the appendix, we consider an alternate model in which we assume that above-water homeowners who receive liquidity shocks are forced to sell as a robustness test. This model features the same economic forces as our baseline model and does reasonably well quantitatively. However, too many households are forced to rent during the crisis relative to the data, which is why we prefer our baseline assumption.

We further assume that foreclosure occurs immediately in our baseline model. In practice, foreclosure is not immediate, and some loans in the foreclosure process do cure before they are foreclosed upon. As another robustness check, in the appendix we consider a model in which there is a delay before foreclosure completion and find similar results. We also study the effects of slowing down foreclosure completions when we analyze policy in Section 6.

Homeowners who experience a foreclosure are prevented from buying for a period of time and must rent in the interim. Foreclosure dramatically reduces a borrower's credit score, and many lenders, the GSEs, and the FHA require buyers to wait several years after a foreclosure

[9] See references in footnote 2.

before they are eligible for a mortgage. Molloy and Shan (2013) use credit report data to show that households that experience a foreclosure start are 55-65 percentage points less likely to have a mortgage two years after a foreclosure start. Consequently, we assume that each period, individuals who defaulted become eligible to apply for a new mortgage with probability σ.

When a homeowner with positive equity sells, a market-maker takes possession of the home. As in Head (2014), each period market-makers have the option of listing a house for sale, which incurs a flow cost of m^n, or renting the house for r_t. When a homeowner with negative equity defaults and is foreclosed upon, the lender takes possession of the house and lists the house for sale, incurring a flow cost of m^d.[10] Lenders do not rent out distressed properties.[11] We anticipate $m^d < m^n < 0$, since vacant, distressed properties tend to depreciate faster and thus require higher ongoing costs to maintain.[12] This will manifest itself in equilibrium as distressed homes selling at lower prices than other listed homes.

Buyers who have been pre-approved, renters, and households with a foreclosure on their record all rent at the equilibrium rent r_t. We assume that a given unit of housing provides rental services for ζ renter households, with $\zeta < 1$ so that renters occupy less square footage than owner-occupants as in the data. The rental market is competitive and the supply consists of those owner-occupied homes being rented out plus a permanent stock of rental homes of mass v^{rs}. The endogenous conversion of owner-occupied homes to rentals to accommodate increased rental demand during the crisis is important in the bust.

2.2 Housing Market

Buyers and sellers in the housing market are matched randomly each period according to a standard fixed-search-intensity constant-returns-to-scale matching function. Defining market tightness θ_t as to the ratio of buyers to listed homes v_t^b/v_t^s where $v_t^s = v_t^n + v_t^d$, the probability a seller meets a buyer q_t^s, and the probability a buyer meets a seller q_t^b can be written as

[10] Alternatively, we could assume that lenders also sell to market-makers and that the costs to market-makers of selling foreclosures are higher than the costs of selling non-distressed properties. Lenders would then receive a price V_t^d for the property.

[11] We abstract away from the purchase of distressed home by institutional investors. In practice, these investors did not enter until late in the crisis (around 2012) and then only in a few markets, such as Atlanta, Tampa, and Phoenix. In our policy section we consider the impact of a government facility which purchases distressed homes and then re-introduces them into the housing market once demand rebounds, which is similar to the role played by investors.

[12] Lenders must make payments to security holders until a foreclosure liquidates, and they must also assume the costs of pursuing the foreclosure, securing, renovating, and maintaining the house, and selling the property. Even though they are paid additional fees to compensate for the costs of foreclosure and are repaid when the foreclosed property sells, the lender's effective return is far lower than its opportunity cost of capital.

functions of θ_t. Buyers meet each type of seller in proportion to their share of listed homes in the market.

When matched, the buyer draws a valuation for the house h from a distribution $F_t(h)$ which is time-varying. This valuation is a one-time utility benefit consumed at purchase and is common knowledge to both the buyer and seller. Prices are determined by generalized Nash bargaining with weight ψ for the seller. If the buyer and seller decide to transact, the seller leaves the market and the buyer obtains financing at the pre-approved terms and becomes a homeowner. If not, the buyer and seller each return to the market to be matched next period. The value of being a homeowner can be written as:

$$V_t^h(h,\Phi) = h + \Gamma_t(\Phi), \tag{2}$$

where $\Gamma_t(\Phi)$ is a continuation value function that is a function of today's loan terms Φ . This value function takes into account the possibility of moving, lock-in, and default. It is the only place where the value of being a renter both with and without a foreclosure on one's credit record enter. We relegate the formal definition of Γ_t and the value functions which are necessary to define it to the appendix.

Denote the total match surplus when a buyer pre-approved to receive a loan with terms Φ meets a seller of type $j \in \{n,d\}$ and draws a match quality h at time t by $S_t^{S,j}(h,\Phi)$, the buyer's surplus by $S_t^{B,j}(h,\Phi)$, and the seller's by $S_t^{S,j}(h,\Phi)$, with $S_t^j(h,\Phi) = S_t^{B,j}(h,\Phi) + S_t^{S,j}(h,\Phi)$. Let the price of the house if it is sold be $p_t^j(h,\Phi)$. The buyer's surplus is equal to the value of being in the house minus the down payment and the outside option of staying in the market:

$$S_t^{B,j}(h,\Phi) = V_t^h(h,\Phi) - \left(p_t^j(h,\Phi) - L\right) - \beta E_t B_{t+1}(\Phi) \tag{3}$$

where B is the value function for being a buyer. The seller's surplus is equal to the price minus the outside option of staying in the market:

$$S_t^{S,j}(h,\Phi) = p_t^j(h,\Phi) - \beta E_t \widetilde{V}_{t+1}^j, \tag{4}$$

where $\widetilde{V}_t^j$ is the value of having a vacant house that can be rented or put up for sale. Because utility is linear and house valuations are purely idiosyncratic, a match results in a transaction if h is above a zero-surplus threshold denoted by $h_t^j(\Phi)$:

$$V_t^h\left(h_t^j(\Phi),\Phi\right) = -L + \beta E_t\left[B_{t+1}(\Phi) + \widetilde{V}_{t+1}^j\right]. \tag{5}$$

We can then define the remaining value functions. The value of listing a house for sale as a market-marker n or as a lender d is equal to the flow payoff plus the discounted continuation value plus the expected surplus of a transaction times the probability a transaction occurs. Because sellers who list meet buyers with probability $q^s(\theta_t)$ and transactions occur with probability $1 - F_t\left(h_t^j(\Phi)\right)$, the value of putting a house up for sale as a $j = \{n, d\}$ type seller is:

$$V_t^j = m^j + \beta E_t \widetilde{V}_{t+1}^j + q^s(\theta_t) \int \left(1 - F_t\left(h_t^j(\Phi)\right)\right) E_t\left[S_t^{S,j}(h, \Phi) \,|h \geq h_t^j(\Phi)\right] dG_t^p(\Phi), \quad (6)$$

where $G_t^p(\Phi)$ is the distribution of pre-approval terms, and the integral is Lebesgue. The value of a vacant home to a market maker $\tilde{V}_t^n$ reflects the option to either list or rent the home, and in equilibrium market makers are indifferent so that:

$$\widetilde{V}_t^n = r_t + \beta E_t \widetilde{V}_{t+1}^n = V_t^n. \quad (7)$$

This condition pins down rents. In equilibrium, competitive market makers make zero profits in expectation, so $\tilde{V}_t^n$ is also the price they pay households selling their home. The value of being a buyer is defined similarly to that of a seller:

$$B_t(\Phi) = -r_t + \beta E_t B_{t+1}(\Phi) + \sum_{j=n,d} q^b(\theta_t) \frac{v_t^j}{v_t^n + v_t^d} \left(1 - F_t\left(h_t^j(\Phi)\right)\right) E_t\left[S_t^{B,j}(h, \Phi) \,|h \geq h_t^j(\Phi)\right]. \quad (8)$$

The buyer value function takes into account the possibility she can meet either normal or REO sellers.

The conditional expectation of the total surplus given that a transaction occurs can be simplified as in Ngai and Tenreyro (2014) by using (2) together with (3) and (4):

$$S_t^j(h, \Phi) = V_t^h(h, \Phi) - V_t^h\left(h_t^j(\Phi), \Phi\right) = h - h_t^j(\Phi). \quad (9)$$

This implies $E_t\left[S_t^j|h \geq h_t^n(\Phi)\right] = E_t\left[h - h_t^n(\Phi)\,|h \geq h_t^n(\Phi)\right].$

Prices can be backed out by using Nash bargaining along with the definitions of the surpluses (3) and (4) and (9) to obtain:

$$p_t^j(h, \Phi) = \psi\left(h - h_t^j(\Phi)\right) + \beta E_t V_{t+1}^j. \quad (10)$$

This pricing equation is intuitive. The first term contains $h - h_t^j(\Phi)$, which is a sufficient statistic for the surplus generated by the match as shown by Shimer and Werning (2007). As the seller bargaining weight ψ increases, more of the total surplus is appropriated to the seller in the form of a higher price. The final two terms represent the seller's outside option of continuing as a seller next period. These terms form the minimum price at which a sale can occur, so that all heterogeneity in prices comes from the distribution of h above the cutoff $h_t^j(\Phi)$.

Define the joint distribution of current homeowner loan balances L and interest rates μ at time t by $G_t(L,\mu)$ and the marginal distribution of loans as $G_t(L)$. Given this specification for the housing market and the environment defined in the previous section, the laws of motion are:

$$\begin{aligned}
l_{t+1} &= (1-\gamma)\, l_t G_t\left(\tilde{V}_t^n\right) + (1-\iota_t)\, l_t\left(1 - G\left(\tilde{V}_t^n\right)\right) \\
&\quad + v_t^b q_t^b(\theta_t) \int \sum_{j=n,d} \frac{v_t^j}{v_t^n + v_t^d}\left(1 - F_t\left(h_t^j(\Phi)\right)\right) dG_t^p(\Phi) \qquad (11)\\
v_{t+1}^b &= \left[\gamma P_t l_t G_t\left(\tilde{V}_t^n\right) + \gamma_r P_t v_t^r + \sigma P_t v_t^f\right] \qquad (12)\\
&\quad + v_t^b\left[1 - q^b(\theta_t)\int \sum_{j=n,d} \frac{v_t^j}{v_t^n + v_t^d}\left(1 - F_t\left(h_t^j(\Phi)\right)\right) dG_t^p(\Phi)\right]
\end{aligned}$$

$$\begin{aligned}
v_{t+1}^r &= (1-\gamma_r P_t)\, v_t^r + \gamma(1-P_t)\, l_t G_t\left(\tilde{V}_t^n\right) + \sigma(1-P_t)\, v_t^f \qquad (13)\\
v_{t+1}^v &= \gamma l_t G\left(\tilde{V}_t^n\right) + v_t^n\left[1 - \int q^s(\theta_t)\left(1 - F_t\left(h_t^n(\Phi)\right)\right) dG_t^p(\Phi)\right] + v_t^a \qquad (14)\\
v_{t+1}^d &= \iota_t l_t\left(1 - G\left(\tilde{V}_t^n\right)\right) + v_t^d\left[1 - \int q^s(\theta_t)\left(1 - F_t\left(h_t^d(\Phi)\right)\right) dG_t^p(\Phi)\right] \qquad (15)\\
v_{t+1}^f &= (1-\sigma)v_t^f + \iota_t l_t\left(1 - G\left(\tilde{V}_t^n\right)\right). \qquad (16)
\end{aligned}$$

Equation (11) says that the stock of homeowners l_t increases due to buyers purchasing homes and decreases due to above-water homeowners receiving taste shocks and below-water homeowners receiving liquidity shocks. Equation (12) provides the law of motion for the stock of buyers v_t^b. New potential buyers come from above-water homeowners who receive a taste shock, individuals who previously defaulted losing their foreclosure flag, and individuals who were previously denied pre-approval. Entering the market as a buyer is conditional on receiving a pre-approval, which occurs with probability P_t. Equation (13) gives the law of motion for the stock of renters who have been denied a pre-approval. These renters re-apply for pre-approval at rate γ_r. Equation (14) gives the law of motion for houses not currently owner-occupied, which we denote as v_t^v. Because a vacant home can either be listed for sale or rented, the number of non-owner-occupied homes is equal to the number of sellers last period who did not sell plus the number of homes that were rented last period plus the flow of above-water homeowners who experience a moving shock. Equation (15) says that the mass of distressed sellers is equal to the inflow of foreclosures plus those distressed sellers who did not sell last period. Finally, equation (16) gives the law of motion for the stock of renters v_t^f locked out of mortgage market due to a foreclosure flag on their credit report. Equilibrium in the rental market requires:

$$v_t^a + v^{rs} = \zeta \left[v_t^b + v_t^f + v_t^r \right], \tag{17}$$

where v^{rs} is the stock of dedicated rental housing. We provide the laws of motion for $G_t(L, \mu)$ and $G_t^p(\Phi)$ in the appendix.

2.3 Mortgage Market

All homeowners purchase their house with a mortgage. Mortgages in our economy can be represented by (L, μ), a tuple that includes the loan balance L and interest rate μ. Mortgagees pay down a constant fraction of the house's principal γ_L plus interest each period as in Chatterjee and Eyigungor (2015). The required payment is thus $(\gamma_L + \mu) L$ and the principal evolves according to $L_{t+1} = (1 - \gamma_L) L_t$. To keep the model tractable, we assume that pre-payment can only occur in the event of a sale, so there is no endogenous refinancing.

Households obtain mortgages from a continuum of competitive, risk neutral lenders. We assume all lenders are identical and none go bankrupt unless the whole system becomes insolvent, so the balance sheet of the financial system can be represented by a representative lender. For tractability, we assume that at any point in time t, there is a single mortgage contract offered to borrowers. To do so, we assume that there is an institutional constraint that all mortgages are pre-approved at a fixed loan-to-value ratio ϕ relative to the average non-distressed price in the market at time t, $\bar{p}_t^n$.

The asset side of the representative lender's balance sheet is comprised of pre-approvals and originated mortgages. The book value Ω_t of these assets is equal to:

$$\Omega_t = l_t \int L dG_t(L) + v_t^b \int L dG_t^p(L), \tag{18}$$

where $G_t(L)$ is the marginal distribution of outstanding loan balances on originated mortgages and $G_t^p(L)$ is the marginal distribution of pre-approved loan balances. We assume that the lender raises the financing for the loan at the time of pre-approval and holds the proceeds in short-term marketable securities earning the risk-free rate μ_f until the loan is funded.

The representative lender funds itself by issuing short-term debt that is insured and thus risk-free from the perspective of its creditors. The lender is, however, subject to a regulatory capital requirement that mandates that its book equity E_t be greater than or equal to a certain percentage χ of its assets at all times:

$$E_t \geq \chi \varOmega_t. \tag{19}$$

The financial system can be in two states of the world. In the first, lenders can costlessly raise equity so that the capital constraint is satisfied. This is the default state that holds in steady state, and in this state $P_t = 1$ because the lender can serve all customers by raising equity. In the second, there is a breakdown in the equity issuance market, the ability to raise new private capital is limited, and the law of motion for lender equity is:

$$E_{t+1} = E_t + (1-\psi)\iota_t l_t \int 1\left[L_t > V_t^n\right]\left(V_t^d - L_t\right) dG\left(L_t\right) + \Delta_t^R, \tag{20}$$

where $G\left(L_t^i\right)$ is the distribution of loan balances at the beginning of period t. The second term represents equity losses due to default, equal to a fraction $1-\psi$ of total book losses, with ψ reflecting the impact of government bailouts and (limited) private equity issuance.[13] The outstanding mortgage balance is L_t, but the lender only receives V_t^d from selling a foreclosure. The final term, Δ_t^R, reflects increases in lender equity due to retained earnings. The expression for retained earnings Δ_t^R is provided in the appendix.[14]

The combination of capital requirements and the limited ability to issue new equity implies that lenders may be prevented from issuing enough debt to cover the full demand for new mortgage pre-approvals. In this case, for (19) to hold, Ω_t must fall, which occurs through P_t falling below one. The resulting equilibrium rationing occurs on quantity rather than price, and rationing is random. The interest rate is always set such that lenders break even in expectation. As a microfoundation, one can imagine that when households apply for financing at time t, there is random sorting into a queue. Only the first P_t fraction of

[13] This specification is natural, as internal and regulatory pressures to recapitalize are likely higher during periods of significant losses. This specification also puts government bailouts and private capital injections during the crisis on an equal footing, which makes economic sense and is useful when we consider policy. In Appendix D.2.4, we consider an alternative law of motion for lender equity where the private capital injection is additive rather than part of ψ. All of the economic insights of the paper continue to hold under this alternate law of motion for lender equity, although the model fit is worse than under our preferred specification.

[14] We assume that the lenders believe that the probability of a breakdown in the equity issuance is sufficiently small and the benefits of debt are sufficiently large such that lenders pay out all retained earning to their shareholders and that the capital requirement binds when equity can be freely issued.

applicants in the line can be serviced. Lenders and households bargain over the financing terms, with the Nash bargaining weight of the household equal to 1. This implies break-even pricing even in the rationing equilibrium.

Our mortgage and banking system is similar to approaches taken by a number of recent papers that study the interaction of housing markets and the banking sector. Elenev et al. (2016) and Greenwald et al. (2018) also include capital constrained lenders in a model with default, but they focus on bank failure in a model with heterogeneous lenders and tailor their model to study restructuring the GSEs and shared-appreciation mortgages, respectively. Hedlund (2016a, 2016b) and Hedlund and Garriga (2019) also include a mortgage sector in a search model of the housing market. However, their directed search model shuts down some of the main effects we highlight in our undirected search model, and their lending sector is competitive and not subject to a capital constraint. Finally, Paixao (2017) finds results of a similar magnitude in a model focused on household consumption.

2.4 Equilibrium

Given this setup we can define an equilibrium:

Definition 1. An equilibrium is defined by: masses l_t, v_t^b, v_t^n, v_t^d, v_t^v, v_t^a, v_t^r, v_t^f, value functions V_t^d, V_t^n, $\tilde{V}_t^n$, $V_t(h, \Phi)$, $\Gamma_t(\Phi)$, and $B_t(\Phi)$, purchase cutoffs $h_t^n(\Phi)$ and $h_t^d(\Phi)$, rents r_t and house prices $p_t^j(h, \Phi)$, a competitive interest rate μ_t, a book value of the representative lender's assets Ω_t, Equity E_t, a pre-approval probability P_t, and distributions $G_t(\Phi)$ and $G_t^p(\Phi)$ such that:

1. The laws of motion and adding up constraints (11), (12), (13), (14), (15), (16), (17), and the adding-up constraint $v_t^v = v_t^n + v_t^a$ are satisfied and the distributions $G_t(\Phi)$ and $G_t^p(\Phi)$, evolve according to the laws of motion in the appendix.

2. The value functions (2), (6) for $j = \{n, d\}$, (8), $\tilde{V}_t^n = V_t^n$, and the equation for $\Gamma_t(\Phi)$ in the appendix are satisfied.

3. The purchase cutoffs satisfy (5) for $j = \{n, d\}$.

4. Rents satisfy (7).

5. House prices satisfy (10).

6. The book value of the representative lender's assets satisfies (18). Equity satisfies (19) and, if equity issuance is limited, satisfies (20).

7. The pre-approval probability $P_t = 1$ if equity can be raised and $P_t \in [0,1]$ satisfies (19) if equity issuance is limited.

8. Lenders break even in expectation given the interest rate.

3 Calibration

We solve our model numerically. After choosing several functional forms, we take a four-step approach to calibrating the parameters of the model. First, we externally calibrate several parameters that correspond directly to parameters commonly used in the literature or which are directly observable in the data. Second, we select several parameters to match the model's steady state without default to pre-downturn empirical moments. Third, we calibrate parameters specific to the downturn to match the model's simulated nationwide downturn to moments from the recent housing crisis. Finally, we evaluate our calibrated model's ability to match cross-sectional heterogeneity in the severity of the downturn across different cities.

3.1 Parameterization of Functional Forms

We calibrate the model so that one period is a week. We use a constant-returns-to-scale Cobb-Douglas matching function so that with v_t^b buyers and v_t^s sellers there are $\Xi \left(v_t^b\right)^{\xi} \left(v_t^s\right)^{1-\xi}$ matches. The probability a seller meets a buyer is thus $q^s\left(\theta_t\right) = \Xi \theta_t^{\xi}$, and the probability a buyer meets a seller is $q^b\left(\theta_t\right) = \Xi \theta_t^{\xi-1}$. The elasticity of the matching function is set to $\xi = 0.84$ based on Genesove and Han (2012).[15]

We parameterize the distribution of idiosyncratic valuations $F(\cdot)$ as an exponential distribution with parameter λ shifted by $\bar{a}_t$, which represents the aggregate valuation of homes. Using an exponential is a neutral assumption because the memoryless property implies that $E_t\left[h - h_t^j\left(\Phi\right) | h \geq h_t^j\left(\Phi\right)\right] = \frac{1}{\lambda}$ for all Φ, which eliminates the effects of difficult-to-measure properties of the tail thickness of the distribution on the conditional expectation of the surplus. This assumption implies all movements in average prices $\bar{p}_t^j = \psi/\lambda + \beta E_t V_{t+1}^j$ work through V_t^j.

[15] Ξ is normalized to 0.5 so that the probability of matching falls on $[0,1]$, and the results are not sensitive to this normalization.

3.2 Externally-Calibrated Parameters

We set the annual household discount rate to five percent. Lenders and households discount at the same exogenous rate, so the riskless lender cost of capital is $\mu_f = 1/\beta - 1$. The probability of a moving shock is set to match a median tenure for owner occupants of approximately nine years from the American Housing Survey (AHS) from 1997 to 2005. We also use the AHS to set the fraction of an owner-occupied house's floor space occupied by a renter, ζ, to be 0.65. This reflects a conservative estimate of the average fraction of square footage per person and lot size occupied by renters who moved in the past year relative to owner occupants, as detailed in the appendix.

We set the non-distressed seller flow cost to reflect an annual maintenance cost of 3%. The Nash bargaining weight ψ is set to satisfy the Hosios condition.

We set the probability that a foreclosed-upon homeowner returns to being able to be pre-approved for a loan σ so that the average foreclosed-upon homeowner is out of the market for two-and-a-half years. Most lenders require one to seven years to pass after a foreclosure to be eligible for another mortgage. For instance, Veterans Administration loans require two years, Federal Housing Administration loans three years, and Fannie Mae and Freddie Mac required five years prior to 2011 and now require seven years, although reductions are allowed based on circumstances. We choose two-and-a-half years to fall in the middle of the range of waiting periods and alter this parameter in robustness checks.

The geometric rate of principal pay down is based on a thirty-year amortization. The LTV requirement is set to $\phi = 0.80$ to reflect the conforming loan limit, and our results are not sensitive to increasing this to $\phi = 0.90$. γ_r is set so that a household that is denied prepayment waits an average of eight weeks to seek another pre-approval. We set the capital requirement χ equal to 10% based on evidence from Begenau (2019).[16] Capital injections during the crisis $\psi = \psi_G + \psi_E$ reflect government bailouts ψ_G and limited private equity infusions. We set $\psi_G = 0.25$. based on evidence from Begenau et al. (2019) and let the data pin down ψ_E.[17]

3.3 Calibration to Steady-State Moments

We choose the initial housing preference parameter $\bar{a}_0$, the shape parameter for the exponential distribution of idiosyncratic valuations λ, the dedicated stock of rental housing v^{rs}, and

[16] Begenau (2019) estimates a Tier 1 capital requirement of 9.3%, taking into account banks' buffers over the regulatory threshold.

[17] Begenau et al. (2019) report total charge-offs by bank holding companies of approximately \$520 billion during the crisis. The "Paulson Plan" injected \$130 billion into these companies in the form of preferred equity (Veronesi and Zingales (2010)). While these numbers do not reflect the entire scope of losses or government interventions during the crisis, the implied bailout of 25% provides a reasonable baseline.

Table 2: Target Steady-State Moments

Moment	Target	Source
Mean House Price	\$300k	Adelino et al. (2012) mean price for 10 MSAs
REO Discount	12.5%	Clauretie and Deneshvary (2009), Campbell et al. (2011)
Non-Distressed Time on Market	26.00 Weeks	Piazzesi and Schneider (2009)
Buyer Time on Market	29.05 Weeks	1.117 times seller time from Genesove and Han (2012)

the REO seller flow cost m^d to match four pre-crisis moments in the data to the steady-state of the model, which we denote by dropping time subscripts. Because default was negligible pre-crisis, we consider a steady state in which $\iota = 0$, so defaults are measure zero. We also assume that in the steady-state equity issuance is costless, which implies that $P = 1$. These two assumptions simplify the steady state, as described in the appendix.

For these four parameters, we target the non-distressed seller time on the market, the average buyer time on the market, the mean house price, and the average REO discount. Intuitively, the mean price and seller time on the market are jointly determined by $\bar{a}_0$ and λ, the buyer time on the market relative to sellers is determined by v^{rs}, and the REO discount is determined by m^d. The REO discount in steady state is calculated as the sale price of an infinitesimal number of distressed sales in the housing market. The target values for the moments are summarized in Table 2 and detailed in the appendix, and these moments are matched exactly. Our resulting parameter values are listed in Table 3, with the parameters set exogenously in the top panel and the parameters set through moment matching in the bottom panel. Importantly, matching the REO discount implies $m^d < m^n < 0$.[18]

3.4 Simulating a National Housing Crisis

We simulate a housing crisis which matches the features of the national US housing market between 2006 and 2013. To do so, we start the model at its steady-state, with the exception that the initial LTV distribution matches the national loan balance distribution in 2006 from CoreLogic. We then compute the perfect foresight impulse response to a housing valuation shock and a concurrent increase in liquidity shocks. Specifically, at the time of the crisis, we hit the model with a permanent shock to home valuations, which we implement by

[18] Andersson and Mayock (2014) indicate that the total costs of foreclosing upon, maintaining, and selling an REO in the Great Recession were 8.5% of the house's value. In the model downturn when REO prices plunge and time on the market rises, the average REO seller's cumulative listing costs rise to 9% of the price. The m^d we use is thus of reasonable magnitude.

Table 3: Parameter Values Calibrated to Pre-Downturn Moments

Param	Value	Units	Param	Value	Units
β	$1-\frac{.05}{52}$	Weekly Rate	$\varXi$	0.5	
γ	$\frac{0.08}{52}$	Weekly Rate	ξ	0.84	
γ_L	$1/(30\times 52)$	Weekly Rate	ζ	0.65	
γ_r	$1/8$	Weekly Rate	ψ	0.16	
σ	$1/(2.5\times 52)$	Weekly Rate	m^n	-0.144	Thousands of $
χ	0.10		ψ_G	0.25	
$\bar{a}$	609.67	Thousands of $	m^d	-0.483	Thousands of $
λ	0.015	Thousands of $	v^{rs}	0.023	Mass (Housing stock = 1)

reducing the minimum idiosyncratic valuation $\bar{a}_0$ to $\bar{a}_t = a^{frac}\bar{a}_0$. We use a permanent shock since it reflects a bursting bubble and is consistent with the empirical facts in Section 1. We further assume that at the time of the crisis, there is an increase in the rate of liquidity shocks ι_t, which will create default among underwater homeowners. We assume $\iota_t = C_\iota Unemp_t$, where C_ι is a constant and $Unemp_t$ is a moving average of the time path of long-run unemployment in the Great Recession as detailed in the appendix. This assumption reflects the fact that large liquidity shocks come from persistent income shocks proxied by long-run unemployment. Finally, we assume that during the crisis lenders are limited in their ability to raise new equity capital for T_E periods. This reflects the temporary breakdown in financial markets that occurred in the Great Recession and creates a role for lender balance sheets in the crisis.

We choose four parameters that affect the downturn but not the steady state of the model to match four empirical moments of the national housing downturn. These parameters are the constant scaling the liquidity shocks C_ι, the size of the decline in housing valuations a^{frac}, the fraction of losses covered by raising private capital during the crisis ψ_E, and the date T_E that equity markets fully re-open. We target four moments: The peak-to-trough decline in the housing price index, the peak-to-trough decline in non-distressed transaction volume, the average REO share between 2006 and 2013, and the total number of foreclosures between 2006 and 2013.[19]

To solve the model, we discretize the possible mortgage balances using an equally spaced grid with 51 grid points. Further numerical details as well as the full system of equations that result from this discretization are provided in the appendix. We do a good job matching the target moments, with no moment more than 0.2 percent from its target value. The resulting parameter values of the calibration are reported in Table 4. The 24.75 percent fall in $\bar{a}$

[19]The moments are jointly determined, but each moment is principally controlled by a single parameter. C_ι determines the total number of foreclosures, T_E the average REO share given the number of foreclosures, a^{frac} the price decline, and ψ_E the volume decline.

Table 4: Parameter Values Calibrated to Downturn Moments

Param	Value
a^{frac}	0.7525
T_E	194 Weeks
ψ_E	.390
C_ι	0.328

implies that prices fall permanently by 11.91 percent nationally due to the valuation shock.

3.5 Simulating the Cross-Section of Downturns

To further evaluate the performance of our model, we consider its ability to match cross-sectional variation in the severity of the crisis across metropolitan areas. To do so, we simulate a separate housing downturn in each CBSA in our CoreLogic data. We assume that each CBSA is a closed system, with a housing market described by Sections 2.1 and 2.2.[20] We furthermore assume that there is a national representative lender. This is a good approximation to the fact that most loans are made by lenders with wide geographic coverage and most loans that are securitized are pooled geographically.[21] This implies that the path for P_t is the same for each CBSA and is calculated from the national downturn.

We allow for cities to differ in the size of the permanent price drop, the magnitude of liquidity shocks, and the initial loan balance distribution. We focus on these three dimensions of heterogeneity because the size of the preceding price run-up is the single best predictor of the size of the ensuing downturn as described in Section 1, because cities varied dramatically in their unemployment rate in the Great Recession, and because the loan balance distribution is critical to the strength of the price-foreclosure feedback in our model. The empirical loan balance distribution comes from proprietary estimates by CoreLogic, who report quantiles of the combined loan-to-value distribution for active mortgages in 2006 computed from public records and CoreLogic's valuation models.[22] We allocate mass to finer loan balance bins in the model within each quantile equally as described in the appendix. We incorporate heterogeneity in the magnitude of liquidity shocks by letting the path of liquidity shocks in city c be $\iota_t^c = C_\iota Unemp_t \frac{\max Unemp^c}{\max Unemp}$, where C_ι is from the national calibration, $Unemp_t$ is the national long-run unemployment time series, and $\frac{\max Unemp^c}{\max Unemp}$ is the ratio of the maximum

[20]This assumption of a closed system implies that we neglect migration between CBSAs. Modeling migration between CBSAs would greatly complicate the analysis and add little in terms of insights.

[21]While one could try to model several lenders with different market shares in different cities, this would introduce tremendous complexity into the model without substantial added value.

[22]Because our model concerns the entire owner-occupied housing stock and not just houses with an active mortgage, we supplement the CoreLogic data with the Census' estimates of the fraction of owner-occupied houses with a mortgage from the 2005-2007 American Community Surveys.

long-run unemployment rate in city c to the national long-run unemployment rate.[23] The permanent shock to home valuations a^{frac} is chosen to generate a price decline proportional to the log price gain from 2003 to 2006:

$$\Delta \log p_{permanent} = -\eta_1 \left(\Delta \log p_{2003-2006} - \eta_0 \right). \tag{21}$$

η_0 is an intercept term chosen so that $\Delta \log p_{permanent}$ matches the a^{frac} in the national calibration and η_1 is a slope term.[24] The limited amount of heterogeneity between cities creates a stringent test for the ability of our model to match the cross-sectional empirical patters of the housing crisis.

η_1 is chosen to minimize the sum of squared differences between the model and the data for the peak-to-trough log aggregate price decline for each metropolitan area. We start each CBSA in the initial steady state and calculate the perfect foresight impulse response given the initial loan balance distribution, the time path of the national lender's pre-approval probability P_t, the city-specific time path of liquidity shocks ι_t^c, and the city-specific a^{frac} calculated to satisfy (21). This yields a unique optimum of $\eta_0 = 0.113$ and $\eta_1 = 0.456$. This implies that if a city has a larger boom than the national average, roughly half of that boom relative to the national average is permanently lost when the bust hits. We return to the cross-sectional calibration in Section 5.

4 Decomposing the Effects of Foreclosures

4.1 Downturn Dynamics

Figure 4 illustrates the model dynamics for the national downturn. As shown in panel A, at the time of the shock, prices fall considerably for both REO and non-distressed sales and gradually return to steady state as the liquidity shocks dissipate. The aggregate price index dips more than the non-distressed price index since REO sales trade at a discount. This discount widens in the crisis, consistent with evidence from Campbell et al. (2011). Panel B shows that sales fall on the impact of the shock and continue to fall substantially as REOs become prevalent in the market. Panel C shows the decline in market tightness over the crisis, and Panel D shows the dynamics of foreclosures. The largest difference between the model and the data is that the model does not feature price momentum (Guren, 2018), so

[23]We take this approach this rather than using the time series of $Unemp^C$ directly because the BLS does not produce city-level long-run unemployment time series that are reliable at a high frequency.

[24]A handful of cities with small booms from 2003 to 2006 have a small permanent increase in their long-term price level under this formulation. Our results are not sensitive to capping $\Delta \log p_{permanent}$ at zero for these cities.

Figure 4: Dynamics of National Downturn

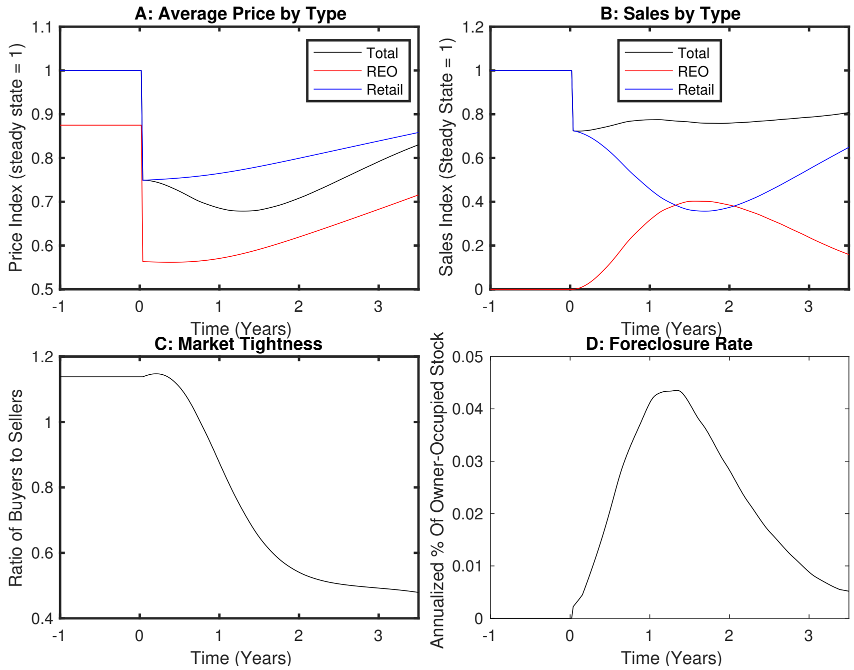

Note: This figure shows the results of the calibrated national model . Panels A and B show the average price and sales by type, with pre-downturn price and volume normalized to 1. Panel C shows market tightness. Panel D shows the annualized fraction of the owner occupied housing stock that is foreclosed upon at each point in time.

prices fall immediately on the impact of the shocks rather than gradually.

4.2 Forces Affecting Prices and Sales in a Downturn

There are three main forces through which default affects the non-distressed market in our model: a lender rationing effect, a foreclosure flag effect, and a choosey buyer effect. In this subsection, we describe the intuition and qualitative impact of each effect before quantifying their respective contributions.

First, there is a foreclosure flag effect. When a homeowner defaults, the lender sells the property as an REO, but the homeowner has a foreclosure flag on her credit record and is locked out of the housing market for a time until the flag is cleared. The foreclosure process thus creates listings immediately, but the corresponding demand only arrives with a delay.

Second, there is a lender rationing effect. Default leads to declines in lender equity. Lenders ration mortgage credit because they are temporarily unable to raise equity and would violate their capital constraint if they were to fully meet the demand for new mortgages by issuing debt. Mortgage rationing causes some homeowners who sell their house to delay their next purchase because they cannot secure financing. This too creates immediate listings, but the corresponding buyers only arrive with a delay as the lenders gradually recapitalize and households re-apply for financing.[25]

By creating an imbalance between the number of sellers and the number of buyers operating in the housing market, both the foreclosure flag and lender rationing effects reduce market tightness θ_t. The reduction in market tightness is partially offset by sellers who convert their owner-occupied house to rental space, which is required to meet the increased rental demand by foreclosed-upon households and households unable to secure financing. The fall in market tightness decreases the probability a seller meets a buyer, which in turn incentivizes sellers to transact faster, weakening their bargaining position and leading to lower prices. This incentive to sell is stronger for REO sellers who have a higher opportunity cost of not meeting a buyer, causing the REO discount to grow. Conversely, buyers are more willing to walk away from a deal, strengthening their bargaining position, which also leads to declines in price. Because the reduction in market tightness pushes down non-distressed prices, it pushes more homeowners underwater and leads to increased default, which leads to further rationing and foreclosure flag effects through a price-default spiral.

The third way that the non-distressed market is affected by default is a choosey buyer effect. Since REO sales trade at a discount, the presence of distressed sales increase the buyer's outside option to transacting, which is resampling from the distribution of sellers next period. This choosey buyer effect causes infra-marginal buyers to walk away from non-REO listings. This reduces non-distressed transaction volume and causes buyers to negotiate a lower price when they do transact. The choosey buyer effect is reinforced by the foreclosure flag effect and lender rationing effect. Both push down market tightness which has a bigger effect on the value functions of REO sellers and increases the REO discount. This in turn sweetens the prospect of being matched with an REO seller next period, amplifying the choosey buyer effect.

The choosey buyer effect is new to the literature and formalizes folk wisdom in housing markets that foreclosures empower buyers and cause them to wait for a particularly favorable transaction.[26] We expect that choosey buyer effects arise in other frictional asset markets

[25]See Paixao (2019) for empirical evidence that house price declines affected credit supply through bank balance sheets in the housing bust.

[26]For instance, The New York Times reported that "before the recession, people simply looked for a house to buy...now they are on a quest for perfection at the perfect price," with one real estate agent adding that "this is the fallout from all the foreclosures: buyers think that anyone who is selling must be desperate. They walk in with the bravado of, 'The world's coming to an end, and I want a perfect place'" ("Housing Market Slows as Buyers Get Picky" June 16, 2010). The Wall Street Journal provides similar anecdotal evidence, writing that price declines "have left many sellers unable or unwilling to lower their prices. Meanwhile, buyers remain gun shy about agreeing to any purchase without getting a deep discount. That dynamic has fueled buyers' appetites for bank-owned foreclosures" ("Buyer's Market? Stressed Sellers Say Not So Fast" April 25, 2011). Albrecht et al. (2007) introduce motivated sellers into a search model, but focus

Table 5: Decomposition of Effects in National Model

Statistic (Peak to Trough Unless Indicated)	No Liquidity Shocks	No Rationing, No Choosey Buyer	No Rationing	Full Model
Price Index	40.57%	71.49%	72.88%	100%
Non-Distressed Price Index	51.96%	74.19%	77.40%	100%
Non-Distressed Sales Volume	22.71%	51.86%	56.43%	100%
Average REO Share	0%	61.52%	62.75%	100%
Total Foreclosures	0%	76.84%	77.76%	100%

Note: Each cell indicates the fraction of the full model accounted for each column's model relative to the full model. No rationing turns off the lender rationing effect by setting $P_t = 1$. No rationing and no choosey buyer additionally turns off the choosey buyer effect by altering the buyer's value function so she does not take into account the possibility of meeting an REO seller, which leaves only the foreclosure flag effect. No default turns off all of the effects and entirely eliminates default.

with idiosyncratic valuations.

Finally, there are two other effects at work in our model that affect aggregates that are averages of the distressed and non-distressed markets. First, there is a lock-in effect. In the absence of liquidity shocks – and thus the absence of default – the permanent decline in housing valuations places some households underwater. These households are then locked into their current house until they pay down their mortgage enough to become above water. This lock-in effect has a minimal impact on prices but does lower sales volume. Second, there is a compositional effect. A greater share of REO sales makes the average sale look more like an REO, which sells faster and at a lower price both in and out of steady state. This affects sales-weighted averages such as total sales and the aggregate price index.

4.3 Quantitative Decomposition of Effects

To quantify the relative contributions of each force, we introduce them one by one. We first simulate a housing crisis in which we shut down default completely by eliminating liquidity shocks, so the only effects are the initial housing valuation shock and lock-in. We then simulate a crisis where we shut down both the lender rationing effect by assuming lenders can costlessly raise equity and the choosey buyer effect by assuming the buyer's value function does not take into account the possibility of meeting a REO seller. This leaves only the foreclosure flag effect. We then reintroduce the choosey buyer effect but continue to shut down the lender rationing effect.

on steady-state matching patterns (e.g. whether a high type buyer can match with a low type seller) and asymmetric information regarding seller type. Duffie et al. (2007) consider a liquidity shock similar to our foreclosure shock, but a transaction occurs whenever an illiquid owner meets a liquid buyer, so their model does not have a choosey buyer effect.

Table 5 shows the fraction of the decrease in various measures of prices, sales, and foreclosures in the full nation-wide model accounted for by each model. Our preferred measure of the extent to which foreclosures exacerbate housing downturns is the decline in non-distressed prices because it does not include compositional effects and is a direct measure of foreclosure spillovers. Table 5 reports that the model without any default account explains 51.96% of the peak-to-trough decline in non-distressed prices. The foreclosure flag effect explains an additional 22.23%, and the choosey buyer effect explains an additional 3.21%. The impact of the choosey buyer effect is significantly larger in the hardest hit CBSAs where the REO share was much higher. Finally, lender rationing explains the remaining 22.60%. Note that lender rationing plays a particularly large role in explaining the decline in sales volume. Indeed, to jointly match the substantial price and volume declines in the data, our model requires a significant lender rationing channel. Overall, our quantitative results suggest that accounting for default is crucial to understanding the dynamics of the housing bust.

These effects are substantially larger than the spillover effects of foreclosures estimated by microeconometric studies that cannot account for the general equilibrium effects of foreclosure at the search market level or the effects of foreclosures on lender balance sheets (e.g. Campbell et al., 2011; Gerardi et al., 2015; Anenberg and Kung, 2014). An exception is Mian et al. (2014), who find much larger effects using more macro variation arising from differences in foreclosure policies at state lines, which is consistent with our finding of a substantial search-market-level effect.

4.4 Robustness

In the appendix we present calibration robustness checks to show that our results are robust to some of our assumed parameter values, in particular γ_r, σ, and ζ. There is some dispersion in the strength of the various effects across calibrations, but generally the lender rationing effect accounts for between 20 and 25 percent of the overall decline in non-distressed prices, the choosey buyer effect accounts for approximately three percent, and the foreclosure flag effect accounts for 20 to 25 percent of the decline. In all cases, there are important spillovers from foreclosures to the non-distressed market, principally through the foreclosure flag and lender rationing channels.

The appendix also reports results for four alternate models. First, we show that a model without the lender rationing fails to account for the decline in sales in the data. Second, we show that a model in which there is some foreclosure delay can still fit the data quite well. Third, we show that a model in which positive-equity homeowners who experience an liquidity shock are forced to sell does reasonably well matching the cross-section but

Figure 5: Cross-CBSA Simulations vs. Data

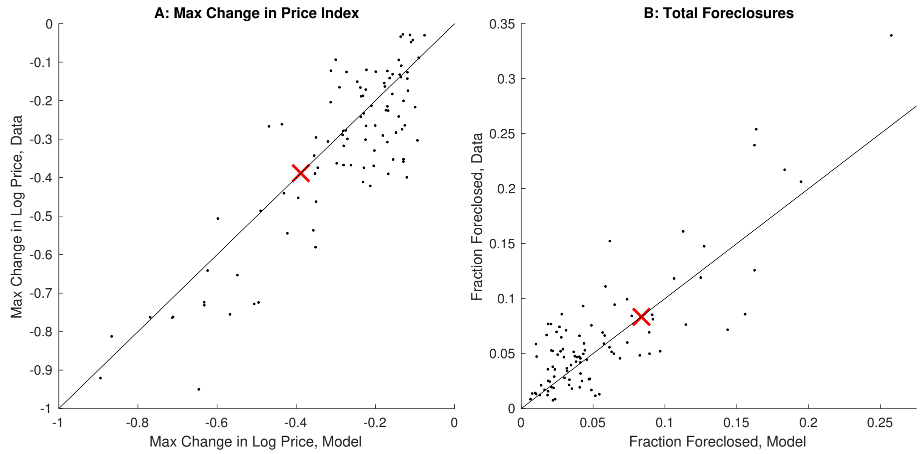

Note: Each panel shows scatter plots of data vs simulation results for 96 CBSAs in regression analysis. The red X represents the national simulation and each black dot is a CBSA. The 45-degree line illustrates a perfect match between the model and the data. The variable being plotted shown in each plot's title. The data are described in the appendix. The calibration methodology, which fits the cross-cities model only to the aggregate price decline in panel A, is described in text. The price decline is the maximum peak-to-trough change, while the fraction foreclosed which is the total from 2006 to 2013.

overstates the amount of conversion to rental relative to the data. Finally, we show that we obtain similar results with an alternate parameterization for private equity injections.

5 Cross-City Quantitative Analysis

To provide further support for our structural model and calibration, we now assess whether our model with a limited amount of heterogeneity can account for differences in the downturn across cities. We then illustrate the importance of default and foreclosure in explaining the downturn by showing that our model does a better job at matching the cross-sectional moments described in Section 1 than a model without default.

5.1 Cross-Sectional Model Fit

Figure 5 plots our simulated results from the baseline model against actual data for 96 CBSAs (black dots) and the national model (red X). Panel A shows the maximum log change in aggregate prices, which is used in the calibration of η_1. The model fits well, with the data points clustering around the 45-degree line across the spectrum of price declines. Indeed, when we regress the simulated data on the actual data we get a coefficient of 1.005, and we cannot statistically reject a coefficient of one and an intercept of zero. Panel B shows the fraction of the housing stock foreclosed upon over eight years. This is a moment used in the

Table 6: Cross-CBSA Simulations vs. Data

Variable	$\Delta \log P$	$\Delta \log (\mathrm{P}_{\mathrm{nd}})$	Mean REO Share	% Foreclose	% Convert
Reg Coef of	1.005	1.389	0.8283	0.992	0.434
Data on Model	(0.063)***	(0.087)***	(0.090)***	(0.069)***	(0.055)***
R^2	0.736	0.734	0.482	0.695	0.419

Note: Each column shows a comparison of the model and data for the given variable. The comparisons show the slope term of a regression of the actual data on the model simulated data. Standard errors are in parenthesis.

national calibration, so the national model is fits almost exactly, but it is not a target for the cross-section. Nonetheless, when we regress the simulated data on the real data, we get a coefficient of 0.992 and we cannot reject a coefficient of one and an intercept of zero. The limited heterogeneity in our model thus does a good job capturing the variation in default across cities.[27]

Table 6 summarizes the model fit for the national prices and foreclosures as well as three other metrics: the decline in non-distressed prices, the mean REO share, and a proxy for the share of the owner-occupied housing stock converted to rentals. For each outcome, we report the coefficient and r-squared we obtain when we regress the simulated data on the actual data. The coefficient is somewhat too high for non-distressed prices because the model under-predicts the decline in non-distressed prices in the hardest hit CBSAs. However, the non-distressed price index in these CBSAs indicates a declining foreclosure discount, which is inconsistent with the literature and suggests that these indices are biased downward by negative quality selection on the non-distressed houses that sell in the hardest-hit CSBAs. The model does well with REO share, although the coefficient is a bit too low because the model under-predicts the sales decline in the least-hit CBSAs. In the data, even cities with no price decline exhibited a significant volume decline. The volume decline from the decrease in national pre-approvals is not enough to fully match the volume decline in these cities.

The last column provides an additional out-of-sample test by comparing the maximum share of owner-occupied homes converted to rental homes in the model to approximate figures for 2006 to 2013 described in the appendix. This is important because if the model dramatically under-predicts the number of conversions, the change in market tightness in the bust will be too strong and the model will ascribe too much of the downturn to foreclosures.

[27]The appendix presents a calibration where we do not include heterogeneity in the liquidity shock series by CBSA. The model qualitatively has fits many of the features described in this section but does not quantitatively fit quite as well: if we regress the log change in the aggregate price index in the model on the corresponding data, we get an r-squared of 0.667 rather than 0.736. The model fits better with the unemployment heterogeneity because of a few hard-hit CBSAs like Las Vegas.

Figure 6: Boom vs. Bust in Baseline and No Default Models

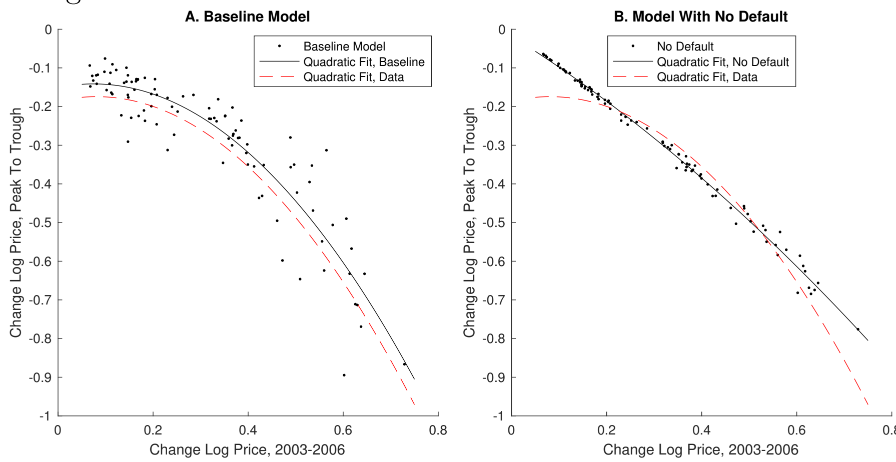

Note: The left panel shows the size of the boom vs. the model simulated data for the baseline calibrated model with default, while the right panel shows the same plot for the no default model. The black solid line is a best quadratic fit. The red dashed line shows the best quadratic fit to the actual data. For the no default model, $\eta_1 = 0.844$.

The model predicts that at the peak level of conversion, 7.38 percent of the owner-occupied housing stock is converted to rentals nationally relative to 4.35 percent in the data. Across cities, there is a positive correlation between the model and the data despite a considerable amount of noise due to the data we use for conversion being a crude proxy. The amount of endogenous conversion in our model is thus of the right order of magnitude.

5.2 Comparison With No Default Model

We now ask how well our model can account for cross-sectional variation in the data relative to a model with no default. To do so, we compare our model to a model with no liquidity shocks and thus no default or foreclosures. We calibrate the model to match the national price decline and optimally choose η_1 using equation (21) as before to give the no-default model the best possible opportunity to match the data.

Figure 6 replicates Figure 3 and plots the size of the bust against the size of the boom using model simulations for the baseline and no default models rather than the raw data. The solid black line shows the best quadratic fit to the model simulated data, while the red dashed line shows the best quadratic fit to the actual data as in Figure 3.

The baseline model quantitatively captures the non-log-linearity in the size of the bust relative to the size of the boom in the data: The solid black and dashed red lines are close to each other and have similar curvature. Furthermore, moving from a linear fit to a quadratic fit in the simulated data increases the r-squared from 0.78 to 0.85, relative to 0.62 to 0.68

Table 7: Model vs. Data: Interaction of Size of Boom With High LTV Share

$\Delta \log(\text{P}) \times$ Z LTV > 80%	$\Delta \log(\text{P})$	$\Delta \log(\text{P}_{\text{nd}})$	Mean REO Share	% Foreclosed
Baseline	-0.396	-0.140	0.307	0.167
	(.034)***	(0.013)***	(0.029)***	(.014)***
No Default	-.102	-.102		
	(.004)***	(.004)***		
Data	-0.310	-0.336	0.205	0.235
	(0.123)**	(0.113)***	(0.075)***	(0.054)***

Notes: * = 10% Significance, ** = 5% Significance *** = 1% significance. All standard errors are robust to heteroskedasticity. Each column shows estimates of (1) with the constant suppressed. P_{nd} is a non-distressed only price index. The mean REO share of sales volume is the average from 2008 to 2012, and the fraction foreclosed is the fraction of the housing stock foreclosed upon over the first 8 years of the downturn. All data is from CoreLogic and described in the appendix.

in the actual data. By contrast, the model without default is nearly linear in the size of the bust relative to the size of the boom and adding a quadratic term does little to improve the fit, with the r-squared rising from 0.990 to only 0.994. This is because essentially all of the price decline comes from the permanent decrease in prices that is proportional to the boom. Consequently, the model fit for price is much better in both the hardest and least hardest hit areas in the model with default relative to the model without default. A final measure of fit is the mean squared error for the aggregate price index. Without default, the mean squared error is 0.0166, while with default this falls to 0.0135.

To further explore the improvement in fit, Table 7 reports the interaction term from regression (1) from Table 1 using simulated outcomes from the baseline and no default models. The model does qualitatively well. The no default model has a negative interaction for price due to lock in, but the coefficient is too small relative to the data. The baseline model with default, by contrast, does a better job of matching the data, although the interaction is a too strong for regular prices and too weak for non-distressed prices.

Overall, we conclude that including default is crucial for models to match cross-sectional moments from the recent housing bust.

6 Foreclosure Policy

A number of foreclosure mitigation policies have been proposed to reduce the severity of a housing crisis. In this section, we use our calibrated model to perform a quantitative study of the positive effects of several different interventions, focusing on their effects on house prices declines and foreclosure rates. We begin by evaluating the effects of interventions at the local local level, that is in one city holding national credit conditions fixed and then move on to to evaluating the effects of national government policy interventions that allow credit conditions to be endogenous.

Before examining the policies, there are some important limitations to our policy analysis that the reader should keep in mind. First, since we do not fully model household consumption and the household costs of default, we are unable to make quantitative statements about welfare.[28] Second, while we compute the dollar costs to the government of various types of policy interventions, we do not model the means by which the government finances the expenditure, which may involve distortionary taxation or borrowing that affects equilibrium interest rates. Nonetheless, we are able to give clear intuitions on the positive effects of various government interventions and make clear, quantitative statements about the cost-effectiveness of various types of foreclosure policies in terms of mitigating a price-default spiral.

6.1 Limiting Foreclosure Completions

We begin our policy analysis by considering the impact of slowing down the rate of foreclosure completions, which was a policy intervention hotly debated during the housing crisis. Proponents argued that this policy could be implemented quickly and at a low cost relative to other government interventions. However, this policy also had its detractors. For example, during the 2012 Presidential campaign, Mitt Romney proposed removing legal barriers to foreclosure completion to get the economic damage over with rather than slowing them down to prolong things.

Since most of the variation in foreclosure timelines is at the the state or municipal level, we consider the local impact of slowing down the rate of foreclosure in a single city, taking national lender balance sheets as given. That is, we hold P_t fixed for this policy only. To incorporate the policy into our model, we continue to assume that homeowners becomes delinquent at the rate ι_t. However, we now further assume that a maximum Φ of the housing stock can be foreclosed upon each week due to institutional or legal constraints. In particular, we set Φ to half of the maximum rate of foreclosure completions in the baseline model, which corresponds to limiting the total number of foreclosures to 2.20% of the total local housing

[28]In practice, default and foreclosure can have significant welfare effects that are outside the scope of our model. For example, defaulting is costly for households. Foreclosures have significant non-pecuniary costs and lead to worse outcomes for children (Diamond et al., 2019), , adversely affect future employment outcomes (Brevoort and Cooper, 2013) and carry a significant cost of social stigma (Guiso et al., 2013). Furthermore, default and foreclosure amplify price declines in our model, and these declines in housing prices can have their own welfare impacts by impeding borrowing by households and firms (Iacoviello, 2005; Chaney et al., 2012; Adelino et al. 2015) and through aggregate demand effects (Mian and Sufi, 2011; Mian and Sufi, 2014).

Table 8: Effects of Limiting Foreclosure Completions

ω	$\Delta \log P_{agg}$	$\Delta \log P_{nd}$	Total Foreclosures
Baseline (No Policy)	-0.388	-0.289	8.37%
0	-0.363	-0.301	9.24%
1/(3×52)	-0.339	-0.281	7.38%
1/(2×52)	-0.335	-0.278	7.05%
1/(1×52)	-0.329	-0.273	6.53%

Notes: The table shows the peak-to-trough decline in log aggregate prices, log non-distressed prices, and the total foreclosures from 2006 to 2013 under each policy. All policies are fore a local intervention that does not affect the representative lender's balance sheet.

stock in a given year.[29]

Limiting foreclosure completions results in an equilibrium backlog of foreclosure starts waiting to be completed during the crisis. We assume that homeowners in the backlog are randomly processed. Finally, due to the lag, homeowners who are delinquent but who have not been foreclosed upon have the opportunity to "cure" out of foreclosure by becoming current on their loan, which we assume occurs to underwater homeowners each period with probability ω. This allows us to evaluate how the effectiveness of the policy depends on the rate at which homeowners recover from a liquidity shock such as a long-term unemployment spell. Homeowners also cure if prices rise to the point that they have positive equity, in which case they can sell their house.

Our results are reported in Table 8. We first consider the case in which homeowners never cure from their liquidity shock, so $\omega = 0$. In this case, slowing down foreclosures is actually mildly *counterproductive*. While there is a 6.44% smaller decline in the aggregate price index, this is purely a compositional effect, and the non-distressed price index actually falls by 4.15% *more*.

To understand why slowing down foreclosures makes non-distressed price declines larger, recall that the non-distressed average price $\bar{p}_t^n$ is a constant markup over the seller's value function V_t^n. By equation (6), the seller's value function is equivalent to holding a financial perpetuity which costs m^n each period but pays out θ/λ each period with time-varying probability equal to the probability that a seller transacts with a buyer. A policy of slowing down foreclosures leads to two competing forces on the value of this claim. First, in any given period, the policy reduces the imbalance between the number of buyers and sellers, leading to a smaller peak-to-trough decline in the equilibrium probability of a seller transacting. However, the policy also delays the recovery, so that the decline in the seller's probability of

[29]Note that capping the total number of foreclosure completions allowed in a given period is just one way to implement a slowing down policy. Alternatively, one could simply assume that foreclosure completions occur at some rate σ_f. We have explored this policy and find similar qualitative and quantitative results.

sale lasts longer. In our baseline calibration, the latter effect weakly dominates, leading to lower prices, and ultimately more default. Because of the lengthened crisis, few homeowners cure by coming above water. The relative strength of the these two forces is a numerical result and depends on our calibration.[30]

When homeowners exogenously cure ($\omega > 0$), slowing down the pace of foreclosures can reduce the the severity of the crisis if ω is high enough because since some homeowners recover before they are foreclosed upon, leading to a smaller price-default spiral. Quantitatively, if $\omega = 1/(2 * 52)$ so that homeowners cure on average after two years, halving the maximum flow of foreclosures reduces the non-distressed price decline by 3.81%. If $\omega = 1/52$, the same policy reduces the non-distressed price decline by 5.54%. Given the evidence of long-term scarring for displaced workers in the labor literature, ω is likely low in practice. However, our results suggest that if policy makers expect a quick recovery either in the labor market or in house prices, slowing down foreclosures may be modestly effective.

6.2 Government Cash or Equity Injections

We now turn our attention to three policy interventions implemented at the national level which can affect lender balance sheets and once again let the pre-approval probability P_t be endogenous.

The first policy we consider is an additional government equity injection into the financial sector. Our baseline calibration has such an intervention already built into it to reflect the 25% bailout of lenders losses during the crisis by the Federal government. We begin our analysis by considering the impact of larger government bailouts of financial institutions. Specifically, we consider the impact of cash injections equal to 30%, 40%, and 50% of bank losses

The results from this policy experiment are reported in the first three rows of Table 9. The baseline calibration with a cash injection of 25% of banking sector losses leads to a non-distressed price index log decline of -0.289 and a foreclosure rate of 8.34%. The present-value cost to the government is \$1,584, calculated on a per household basis.[31] [32] A 30% injection of cash as a percentage of losses results in a 10.67% smaller non-distressed price decline and a 8.39% reduction in the number of foreclosures and costs \$1,619 per household. A 40% injection of cash as a percentage of losses results in a 20.52% smaller non-distressed price decline and a 17.07% reduction in the number of foreclosures and costs \$1,819 per household.

[30] For example, if sellers discount the future by more, then the first effect can dominate.

[31] Formulae for the cost to the government of each policy can be found in the appendix.

[32] These figures reflect the cost of a cash bailout. A better policy is likely to purchase preferred stock. Since the lender does not default along the perfect foresight equilibrium path, this constitutes a riskless investment on the part of the government.

Table 9: National Policy Interventions

Policy	$\Delta \log P_{nd}$	Total Foreclosures	Per-Household Cost
Baseline (25% Bailout)	-0.289	8.34%	$1,584
30% Bailout	-0.258	7.64%	$1,619
40% Bailout	-0.230	6.92%	$1,819
50% Bailout	-0.220	6.73%	$2,175
25% Bailout, $5K Principal Reduction	-0.272	7.59%	$1,782
25% Bailout, $10K Principal Reduction	-0.259	6.98%	$1,917
25% Bailout, $\sigma_g = 1/52$	-0.159	4.98%	$2,150

Notes: The table shows the peak-to-trough decline in log non-distressed prices, the total foreclosures from 2006 to 2013, and per-household costs to the government. The intervention is a national intervention that affects the representative lender's balance sheet. The policy is as indicated, and formulae for computing the per-household cost are in the appendix.

A 50% injection of cash as a percentage of losses results in a 23.88% smaller price decline and a 19.30% reduction in foreclosures. It costs $2,175 per household.

Additional government bailouts of the financial sector ameliorate the impact of the housing crisis by increasing lender equity and thus reducing the amount by which the representative lender has to ration pre-approvals, a force which accounts for 22.6 percent of the decline in non-distressed prices. This reduced rationing has a direct effect on market tightness and drives up non-distressed prices, which brings some households above water, reduces default, and undoes some of the price-default spiral. The effect of additional equity injections is strong because the 10% capital requirement implies that for each 1% increase in bank equity, the bank can increase the size of the balance sheet by 10%.[33]

6.3 Mortgage Modification

The next policy we consider is principal reduction, which was implemented as part of the HAMP PRA program. This program financed principal forgiveness for homeowners, with the principal reduction capped at the point where the LTV ratio reached 115%. Using a regression discontinuity design, Ganong and Noel (2019a) document that, controlling for payment reduction, HAMP principal reductions for underwater homeowners had no effect on the short-run incidence of default. Scharlemann and Shore (2016) report similar results using a regression kink design.

To keep our analysis in line with the actual HAMP program, which only reduced principal up to the point where LTV ratios reached 115%, we consider a policy which offers principal forgiveness of Φ dollars or until 115% LTV is reached (whichever is less) to all homeowners

[33]Note that higher levels of private equity injections during the crisis would have similar effects.

with LTV ratios exceeding 115% at the onset of the crisis. We assume that the government pays the residual principal and interest payments to the bank throughout the life of the loan, so the bank does not incur any losses on its balance sheet. To isolate the impact of the principal reduction, we assume that incidence of the liquidity shocks remains constant. Finally, we again consider this policy as incremental, taking the baseline 25% lender bailout as given.

The quantitative results for principal reduction are in the fourth and fifth rows of Table 9. Relative to a baseline with no principal reduction, $5,000 principal reduction up to a 115% LTV cap lowers the price decline by 5.83% and the incidence of foreclosure by 9.04% and has a total cost of $1,782 per household. $10,000 principal reduction up to a 115% LTV cap lowers the price decline by 10.47% and the incidence of default by 16.35% and has a total cost of $1,917 per household.

The first result to note about principal reduction is that it can be effective at mitigating price declines and foreclosures. At first blush, this may seem to contradict to micro-econometric studies like Ganong and Noel (2019a) and Scharlemann and Shore (2016) who find that HAMP principal forgiveness had no effect on default at the margin. However, if we were to compare underwater households who receive principal reduction to those who do not in our model, we would also find no difference in short-run default rates since the policy does not bring households above water. Our model is thus entirely consistent with their micro findings.

The effectiveness of principal reductions is instead due to two equilibrium effects that are absorbed into a constant or fixed effect in micro-econometric studies. First, the policy affects default in the future because more households are above water than otherwise would be as prices rise.. As agents are forward-looking, the reduction in future default feeds back into a smaller price decline at the start of the crisis. Second, principal reductions act as an indirect bailout to lenders, as households that receive a liquidity shock and default have a smaller principal balance than they otherwise would have. Our equilibrium model suggests that micro studies on the impact of principal forgiveness like Ganong and Noel (2019a) and Scharlemann and Shore (2016) may understate the full impact of programs like HAMP because they cannot account for the equilibrium impacts of the policy.

The second result to note about principal reduction is that it is less cost-effective than the other policies we consider, as shown in Table 9. This is because principal forgiveness is imperfectly targeted. Intuitively, the government cannot tell who will receive a liquidity shock in the future and thus gives principal reduction to all households with sufficiently high LTV, which drives up the cost of the policy relative to a policy that targets households conditional on default. The other two policies are better-targeted because they condition on default: Equity injections cover a fraction of lender losses due to default, and government purchases of distressed homes only occur in the event of a foreclosure. These policies thus do not spend money to mitigate foreclosures that may not occur, while principal reduction does.

6.4 Government Purchase of Distressed Houses

The final policy intervention we consider is a government facility to purchase distressed homes and then slowly re-introduce them to the market. In our model, foreclosures create listings today but potential buyers are locked out of the market, which puts downward pressure on prices. Slowing down the rate at which foreclosed properties are listed on the market for sale could address this dynamic imbalance between supply and demand and lead to smaller price declines and less default. This does not occur in equilibrium because competitive lenders do not internalize the fact that by listing a property today they are creating more foreclosures and lowering the price they can command for a foreclosure. This externality creates scope for government intervention.

We implement such an intervention within our model by assuming that the government sets up a facility to purchase distressed homes at the price of V_t^d so that lenders are indifferent between listing the property for sale and selling it to the government. The government then re-introduces REOs into the market at the rate σ_g. The government pays per-period costs m^d while keeping the home off the market. Because the house is vacant, the government experiences increased depreciation and higher maintenance costs just as with lenders selling a foreclosure. When a house is re-introduced by the government back into the market at time T, the government sells the house to a market-maker at a price V_T^d.

Our results are reported in the final row of table Table 9. The policy is highly effective: A rate of $\sigma_g = 1/52$, corresponding to keeping the house off the market for an average of 1 year, leads to a 44.8% smaller national price decline and a 40.3% reduction in the number of foreclosures. The total cost to the government is \$2,150 per household. By contrast, Table 9 shows that a 50% lender bailout costs the government \$2,175 per household but delivers significantly smaller reductions in default and non-distressed prices. The policy is thus quite cost-effective, even accounting for the government's high maintenance costs when they hold homes off the market.

The reason this policy intervention is highly effective is that it mitigates both the foreclosure flag *and* lender rationing channels discussed in Section 4.[34] It clearly mitigates the foreclosure flag effect by directly addressing the dynamic imbalance of demand and supply

[34]Recall that, even in the absence of any lender rationing, the foreclosure flag effect itself accounts for 22.2% of the total price decline.

created by the foreclosure process, as described above. The lender rationing channel is also weakened because of the interaction between the foreclosure flag and lender rationing effects: By weakening the foreclosure flag effect, there is less default and smaller bank losses, which then mitigates the credit rationing by banks. The total result is significantly less amplification during the crisis and a far smaller price-default spiral.

Finally, it is useful to distinguish this policy from the policy of limiting foreclosure completions discussed in Section 6.1. At first glance, these policies may appear similar, but they are actually quite different. Recall that the foreclosure flag effect occurs because when a foreclosure happens, the home is sold now but the foreclosed-upon household does not return to the market until the foreclosure flag is removed from its credit record. The government facility directly addresses the foreclosure flag effect by holding the house off the market until demand picks up. Slowing down foreclosures, by contrast, delays the foreclosure, but the dynamic imbalance is unchanged because upon foreclosure completion the house is sold immediately while the foreclosed-upon homeowner is locked out. The foreclosure flag effect is thus just as potent, so the effects of the slowing down foreclosures policy is far more limited when cure rates are low.

7 Conclusion

This paper uses a structural analysis to show that foreclosures have important equilibrium effects that exacerbate housing downturns and to analyze foreclosure mitigation policy. We develop a quantitative search model of the housing market in which default erodes lender balance sheets, lenders sell foreclosed homes at a discount due to high holding costs, and homeowners who are foreclosed upon cannot immediately purchase another home. Lender rationing and the impact of foreclosure flags on potential buyers' credit records reduce the number of buyers relative to sellers, worsening seller outside options. Sellers, and particularly REO sellers, become highly motivated to sell, while buyers become more choosey due to the presence of distressed REO sellers offering properties at a discount. These effects create downward pressure on non-distressed prices, which in turn leads to additional default and a price-default spiral.

In our quantitative analysis, these equilibrium effects prove crucial to match the empirically observed declines in house prices and transaction volumes during the housing crisis. Furthermore, the model with default is better able to match the non-linearity in boom size relative to bust size across cities relative to a model without default. The deterioration in lender balance sheets generates a decline in non-distressed prices of 22.6 percent, while the foreclosure flag and choosey buyer effects cause an additional 25.4 percent decline.

The presence of sizable equilibrium effects of foreclosures opens the door to for foreclosure mitigation policy to ameliorate a downturn, and we use our quantitative model to compare several policies. The most cost-effective policy we consider is a government intervention that holds foreclosures off the market until demand rebounds, thereby rectifying the dynamic imbalance of supply and demand caused by foreclosures. Lender equity injections are also quite effective. The least cost-effective policy is principal reduction as implemented by HAMP, although it is more effective than a partial-equilibrium analysis would suggest since it acts as a poorly targeted bailout of lenders. Finally, slowing down foreclosures can be counterproductive if households do not cure out of foreclosure quickly enough because it lengthens the crisis, which offsets the benefits of having fewer foreclosures on the market at any one time. However, this policy can be effective if the cure rate is sufficiently fast.

Overall, our findings suggest that models of the housing market need to incorporate features that allow for default to cause overshooting in a bust. Our findings also imply that foreclosures and some foreclosure policies have far stronger effects than those implied by micro studies that absorb the equilibrium effects into a constant or fixed effects. We hope future work continues to refine our understanding of the equilibrium effects of foreclosures and their role in shaping foreclosure policy.

References

Adelino, M., A. Schoar, and F. Severino (2015). House Prices, Collateral, and Self-Employment. *Journal of Financial Economics 117*(2), 288–306.

Albrecht, J., A. Anderson, E. Smith, and S. Vroman (2007). Opportunistic Matching in the Housing Market. *International Economic Review 48*(2), 641–664.

Anenberg, E. (2016). Information Frictions and Housing Market Dynamics. *International Economic Review 57*(4), 1449–1479.

Anenberg, E. and E. Kung (2014). Estimates of the Size and Source of Price Declines Due to Nearby Foreclosures. *American Economic Review 104*(8), 2527–2551.

Begenau, J. (Forthcoming). Capital Requirements, Risk Choice, and Liquidity Provision in a Business Cycle Model. *Journal of Financial Economics*.

Begenau, J., S. Bigio, and J. Majerovitz (2019). Banks Adjust Slowly: Evidence and Lessons For Modeling. Working Paper.

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

Brevoort, K. P. and C. R. Cooper (2013). Foreclosure's Wake: The Credit Experiences of Individuals Following Foreclosure. *Real Estate Economics 41*(4), 747–792.

Campbell, J. Y., S. Giglio, and P. Pathak (2011). Forced Sales and House Prices. *American Economic Review 101*(5), 2108–31.

Chaney, T., D. Sraer, and D. Thesmar (2012). The Collateral Channel: How Real Estate Shocks Affect Corporate Investment. *American Economic Review 102*(6), 2381–2409.

Chatterjee, S. and B. Eyigungor (2015). A Quantitative Analysis of the U.S. Housing and Mortgage Markets and the Foreclosure Crisis. *Review of Economic Dynamics 18*(2), 165–184.

Clauretie, T. M. and N. Daneshvary (2009). Estimating the House Foreclosure Discount Corrected for Spatial Price Interdependence and Endogeneity of Marketing Time. *Real Estate Economics 37*(1), 43–67.

Diamond, R., A. M. Guren, and R. Tan (2019). The Effect of Foreclosures on Homeowners, Tenants, and Landlords. Working Paper.

Duffie, D., N. Garleanu, and L. H. Pedersen (2007). Valuation in Over-the-Counter Markets. *Review of Financial Studies 20*(5), 1865–1900.

Elenev, V., T. Landvoigt, and S. Van Nieuwerburgh (2106). Phasing Out the GSEs. *Journal of Monetary Economics 81*, 111–132.

Elul, R., N. S. Souleles, S. Chomsisengphet, D. Glennon, and R. Hunt (2010). What 'Triggers' Mortgage Default? *American Economic Review 100*(2), 490–4.

Foote, C. L., K. Gerardi, and P. S. Willen (2008). Negative Equity and Foreclosure: Theory and Evidence. *Journal of Urban Economics 64*(2), 234–245.

Fuster, A. and P. S. Willen (2017). Payment Size, Negative Equity, and Mortgage Default. *American Economic Journal: Economic Policy 9*(4), 167–191.

Ganong, P. and P. Noel (2019a). The Effect of Debt on Default and Consumption: Evidence from Housing Policy in the Great Recession. Working Paper, University of Chicago.

Ganong, P. and P. Noel (2019b). Why Do Borrowers Default on Mortgages? A New Test Using High-Frequency Data. Working Paper, University of Chicago.

Garriga, C. and A. Hedlund (2019). No Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession. Working Paper.

Genesove, D. and L. Han (2012). Search and Matching in the Housing Market. *Journal of Urban Economics 72*(1), 31–45.

Gerardi, K., K. Herkenhoff, L. Ohanian, and P. S. Willen (2017). Unemployment, Negative Equity, and Strategic Default. Working Paper.

Gerardi, K., P. S. Willen, E. Rosenblatt, and V. W. Yao (2015). Foreclosure Externalities: New Evidence. *Journal of Urban Economics 87*, 42–56.

Greenwald, D. L., T. Landvoigt, and S. Van Nieuwerburgh (2018). Finanical Fragility With SAM? Working Paper.

Guiso, L., P. Sapienza, and L. Zingales (2013). The Determinants of Attitudes Toward Strategic Default on Mortgages. *The Journal of Finance 68*(4), 1473–1515.

Guren, A. (2018). House Price Momentum and Strategic Complementarity. *Journal of Political Economy 126*(3), 1172–1218.

Head, A., H. Lloyd-Ellis, and H. Sun (2014). Search, Liquidity, and the Dynamics of House Prices and Construction. *American Economic Review 104*(4), 1172–1210.

Hedlund, A. (2016a). Illiquidity and Its Discontents: Trading Delays and Foreclsoures in the Housing Market. *Journal of Monetary Economics 83*, 1–13.

Hedlund, A. (2016b). The Cyclical Dynamics of Illiquid Housing , Debt, and Foreclosures. *Quantitative Economics 7*(1), 289–328.

Iacoviello, M. (2005). House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle. *American Economic Reivew 95*(3), 739–764.

Kaplan, G., K. Mitman, and G. L. Violante (2019). The Housing Boom and Bust: Model Meets Evidence. Working Paper.

Krainer, J. (2001). A Theory of Liquidity in Residential Real Estate Markets. *Journal of Urban Economics 49*(1), 32–53.

Lamont, O. and J. C. Stein (1999). Leverage and House-Price Dynamics in U.S. Cities. *RAND Journal of Economics 30*(3), 498–514.

Mian, A. and A. Sufi (2011). House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis. *American Economic Review 101*(5), 2132–2156.

Mian, A. and A. Sufi (2014). What Explains the 2007-2009 Drop in Employment? *Econometrica 82*(6), 2197–2223.

Mian, A., A. Sufi, and F. Trebbi (2015). Foreclosures, House Prices, and the Real Economy. *Journal of Finance 70*(6), 2587–2634.

Molloy, R. and H. Shan (2013). The Postforeclosure Experience of U.S. Households. *Real Estate Economics 41*(2), 225–254.

Ngai, L. R. and S. Tenreyro (2014). Hot and Cold Seasons in the Housing Market. *American Economic Review 104*(12), 3991–4026.

Novy-Marx, R. (2009). Hot and Cold Markets. *Real Estate Economics 37*(1), 1–22.

Paixao, N. (2017). Housing Prices and Consumer Spending: The Bank Balance Sheet Channel. Working Paper, Bank of Canada.

Paixao, N. (2019). Propogation of House Price Shocks Through the Banking System. Working Paper, Bank of Canada.

Piazzesi, M. and M. Schneider (2009). Momentum Traders in the Housing Market: Survey Evidence and a Search Model. *American Economic Review 99*(2), 406–411.

Scharlemann, T. C. and S. H. Shore (2016). The Effect of Negative Equity on Mortgage Default: Evidence From HAMP's Principal Reduction Alternative. *Review of Financial Studies 29*(10), 2850–2883.

Shimer, R. and I. Werning (2007). Reservation Wages and Unemployment Insurance. *Quarterly Journal of Economics 122*(3), 1145–1185.

Springer, T. (1996). Single-Family Housing Transactions: Seller Motivations, Price, and Marketing Time. *Journal of Real Estate Finance and Economics 13*, 237–254.

Veronesi, P. and L. Zingales (2010). Paulson's Gift. *Journal of Financial Economics 97*(3), 339–368.

Wheaton, W. C. (1990). Vacancy, Search, and Prices in a Housing Market Matching Model. *Journal of Political Economy 98*(6), 1270–1292.

Williams, J. T. (1995). Pricing Real Assets With Costly Search. *Review of Financial Studies 8*(1), 55–90.

A Model Details

In this appendix, we provide technical details on the model which were omitted in the main text.

We first provide the Bellman equation describing the value function of a homeowner. Recall that:

$$V_t^h(h,\Phi) = h + \beta E_t \Gamma_{t+1}(L(\Phi), \mu(\Phi)),$$

where h is the idiosyncratic valuation. The value function Γ_t is:

$$\begin{aligned}\Gamma_t(L,\mu) = {} & (1 - \gamma 1[L \leq V_t^n] - \iota_t 1[L > V_t^n])((\gamma_L + \mu)L + \Gamma_{t+1}((1-\gamma_L)L, \mu)) \\ & + \gamma 1[L \leq V_t^n](V_t^n - (1+\mu)L + P_t B_t(\Phi_t) + (1-P_t)R_t^v) \\ & + \gamma 1[L > V_t^n]\Upsilon_\gamma + \iota_t 1[L > V_t^n]\left(R_t^f + \Upsilon_\iota\right),\end{aligned}$$

where:

$$\begin{aligned}R_t^v &= -r_t + \beta\left[\gamma_r B_t(\Phi_t) + (1-\gamma_r) E_t R_{r+1}^v\right] \\ R_t^f &= -r_t + \beta\left[\sigma B_t(\Phi_t) + (1-\sigma) E_t R_{r+1}^v\right]\end{aligned}$$

are the value functions for being denied pre-approval and having a foreclosure flag, respectively. $\Upsilon_\gamma < 0$ and $\Upsilon_\iota < 0$ respectively denote the utility costs of lock-in and default. As an important special case, note that if there is no default, i.e. $\iota_t = 0$ and $\mu = 1/\beta - 1$, then the present-value of all future loan payments is equal to the loan balance. This implies that:

$$V_t^h(h,\Phi) = h - L + \beta E_t \tilde{\Gamma}_{t+1}(L(\Phi), \mu(\Phi)),$$

where:

$$\begin{aligned}\tilde{\Gamma}_t(L,\mu) = {} & (1 - \gamma 1[L \leq V_t^n])\tilde{\Gamma}_{t+1}((1-\gamma_L)L, \mu) \\ & + \gamma 1[L \leq V_t^n](V_t^n + P_t B_t(\Phi_t) + (1-P_t)R_t^v). \\ & + \gamma 1[L > V_t^n]\Upsilon_\gamma\end{aligned}$$

This simplifies the equilibrium cutoff condition to:

$$h_t^j(\Phi) + \beta E_t \tilde{\Gamma}_{t+1}(L(\Phi), \mu(\Phi)) = \beta E_t\left[B_{t+1}(\Phi) + V_{t+1}^j\right].$$

In steady-state or if there is no lock-in along the equilibrium path, this implies that equilibrium cutoff rules and buyer value functions are independent of the financing terms Φ.

We next discuss the dynamics of the loan balance and pre-approved loan balance distributions $G_t(\Phi)$, $G_t^p(\Phi)$. Since it corresponds closer to our numerical implementation, instead of providing laws of motion for these distributions, we provide the laws of motion for the mass of people with financing terms Φ, which essentially discretizes the $G(\cdot)$ and $G^p(\cdot)$

distributions. We denote these masses as l_t^{Φ} and $v_t^{b,\Phi}$. Note that:

$$
\begin{aligned}
l_t^{\Phi} &= l_t G_t\left(\Phi\right) \\
v_t^{b,\Phi} &= v_v^b G_t^p\left(\Phi\right).
\end{aligned}
$$

These masses follow the dynamics:

$$
\begin{aligned}
v_{t+1}^{b,\Phi} = &\left[\gamma P_t \sum_{\Phi} l_t^{\Phi} 1\left[L\left(\Phi\right) < V_t^n\right] + \gamma_r P_t v_t^r + \sigma P_t v_t^f\right] 1\left[\Phi_t = \Phi\right] \\
&+ \left(1 - q_b\left(\theta_t\right) \sum_{j=n,d} \frac{v_t^j}{v_t^n + v_r^d} E\left[1 - F_t\left(h_t^j(\Phi)\right)\right]\right) v_t^{b,\Phi}
\end{aligned}
$$

and

$$
\begin{aligned}
l_{t+1}^{\Phi} = &\, q_b\left(\theta_t\right) \sum_{j=n,d} r_t^j E_t\left[1 - F_t\left(h_t^j(\Phi)\right)\right] v_t^{b,\Phi} \\
&+ \left(1 - \gamma 1\left[\frac{L\left(\Phi\right)}{1-\gamma_L} \leq V_t^n\right] - \iota_t 1\left[\frac{L\left(\Phi\right)}{1-\gamma_L} > V_t^n\right]\right) l_t^{\left(\frac{L(\Phi)}{1-\gamma_L}, \mu(\Phi)\right)}.
\end{aligned}
$$

Finally, retained earnings are given by:

$$
\Delta_t^R = \sum_{\Phi} v_t^{b,\Phi} \mu_f L\left(\Phi\right) + \sum_{\Phi} l_t^{\Phi} \mu\left(\Phi\right) L\left(\Phi\right)\left(1 - \iota_t 1\left[L\left(\Phi\right) > V_t^n\right]\right) - \mu_f D_t,
$$

equal to the interest collected on all loans which do not default minus the interest the lender pays on its own debt D_t. Recall that the funds for pre-approved loans are invested in short-term marketable securities earning the risk-free rate. The interest rate on all loans is set such that lender's break even in expectation. We discuss this point more fully below.

As described in the main text, we consider a steady state in which $\iota = 0$, so defaults are measure zero. and equity issuance is costless so $P = 1$. This simplifies the steady state considerably. In particular, without default risk, the interest rate on all loans is the risk-free rate μ_f. Furthermore, equilibrium in the housing market is invariant to the current homeowner loan balance distribution as long as no homeowners are underwater, which is the case in steady state. Additionally, the equilibrium cutoffs $h^j\left(\Phi\right)$ are independent of the financing terms Φ.

B Computational Appendix

B.1 Steady-State System

In this section, we provide the full system of equations which solves for the steady-state of our model. Recall that $\iota = 0$ in the steady-state and lenders can costlessly raise equity so $P = 1$. This implies that $v^d = 0$, $v^f = 0$, and $v^r = 0$ in steady-state. Also, with no default, the buyer's value function and the equilibrium cutoffs are independent of Φ, and we suppress this notation. Dropping the t subscript to denote a steady state value, the full system is:

$$\begin{aligned}
l &= (1-\gamma)\, l + v^b q^b(\theta)\,(1-F(h^n)) \\
v^v &= \gamma l + v^n\left[1 - q^s(\theta)\,(1-F(h^n))\right] + v^a \\
v^b &= \gamma l + v^b\left[1 - q^s(\theta)\,(1-F(h^n))\right] \\
v^n + v^a &= v^v \\
v^{rs} + v^a &= \zeta v^b \\
V^n &= m^n + \beta \widetilde{V}^n + q^s(\theta)\,(1-F(h^n))\,\psi E\left[h - h^n | h \geq h^n\right] \\
V^d &= m^d + \beta V^d + q^s(\theta)\left(1 - F\left(h^d\right)\right)\psi E\left[h - h^d | h \geq h^d\right] \\
B &= -r + \beta B + q^s(\theta)\,(1-F(h^n))\,(1-\psi)\, E\left[h - h^n | h \geq h^n\right] \\
\widetilde{V}^n &= r + \beta \widetilde{V}^n \\
\tilde{V}^n &= V^n \\
\tilde{\Gamma} &= \gamma\,(V^n + B) + \beta\,(1-\gamma)\,\tilde{\Gamma} \\
h^n + \beta\tilde{\Gamma} &= \beta\left[B + V^n\right] \\
h^d + \beta\tilde{\Gamma} &= \beta\left[B + V^d\right]
\end{aligned}$$

The first three steady-state equations derive from the laws of motion for the homeowners, the stocks of non owner-occupied homes, and buyers respectively. The next equation says that the number of listed homes and homes being rented out must equal the total stock of non owner-occupied homes. The next is the market clearing condition for the rental market. The following four equations derive from the Bellman equations for the seller and buyer value functions. After those we have the equilibrium indifference condition between listing and renting. Then we have the equation pinning down the value of homeownership and finally the equations pinning down the equilibrium cutoff rules. Note that the expression for $\tilde{\Gamma}$ reflects the fact that there is no default, no lock-in, and $P = 1$ in the steady state. Under the exponential distribution, $E\left[h - h^j | h \geq h^j\right] = 1/\lambda$ for $j = n, d$ which further simplifies the numerical computation of the steady state.

B.2 Solving for the Downturn

We solve for the perfect foresight impulse response to the shocks we use to trigger a downturn under the assumption that households who buy during the downturn do not subsequently default or become locked-in. That is, all default arises from people who are already homeowners at the time of the crisis. We then verify ex-post that this is actually the case.[35] This makes solving the model considerably simpler. First, it implies that the interest rate on loans issued during the crisis is equal to the lender's discount rate μ_f. It further implies that the buyer value function and the equilibrium cutoffs do not depend on financing terms Φ. We therefore suppress this notation in what follows. Finally, if new homeowners are not expected to default, the value function for being a homeowner is considerably simpler.

We then discretize the loan balance distribution using and equally-spaced grid with 51

[35] Recall that all loans are pre-approved at an LTV of 80% relative to the average price. Because non-distressed prices rise on the equilibrium path, this implies that there is no re-default.

points. The equations describing equilibrium in the housing market are:

$$\begin{aligned}
V_t^n &= m^n + \beta E_t V_{t+1}^n + q^s(\theta_t)(1 - F_t(h_t^n))\psi E_t[h - h_t^n | h \geq h_t^n] \\
V_t^d &= m^d + \beta E_t V_{t+1}^d + q^s(\theta_t)\left(1 - F_t\left(h_t^d\right)\right)\psi E_t\left[h - h_t^d | h \geq h_t^d\right] \\
B_t &= -r_t + \beta E_t B_{t+1} + q^b(\theta(t)) \sum_{j=n,d} \frac{v_t^j}{v_t^n + v_r^d}\left(1 - F_t\left(h_t^j\right)\right)(1-\psi)E_t\left[h - h_t^j | h \geq h_t^j\right] \\
\tilde{\Gamma}_t &= \gamma(V_t^n + P_t B_t + (1 - P_t)R_t^v) + \beta(1-\gamma)E_t\Gamma_{t+1} \\
R_t^v &= -r_t + \beta\left[\gamma_r B_t(\Phi_t) + (1-\gamma_r)E_t R_{r+1}^v\right] \\
\tilde{V}_t^n &= V_t^n \\
\tilde{V}_t^n &= r_t + \beta E_t \tilde{V}_{t+1}^n \\
h_t^d + \beta E_t \tilde{\Gamma}_{t+1} &= \beta E_t\left[B_{t+1} + V_{t+1}^d\right] \\
h_t^n + \beta E_t \tilde{\Gamma}_{t+1} &= \beta E_t\left[B_{t+1} + V_{t+1}^n\right] \\
\bar{p}_t^n &= \psi E_t[h - h_t^n | h \geq h_t^n] + \beta E_t V_t^n \\
\bar{p}_t^d &= \psi E_t\left[h - h_t^d | h \geq h_t^d\right] + \beta E_t V_t^d \\
v_a(t) + v_{rs} &= \zeta(v_b(t) + v_f(t) + v_r(t)) \\
v_v(t) &= v_a(t) + v_n(t) \\
L_t^* &= \phi \bar{p}_t^n
\end{aligned}$$

Note that L_t^* is the size of pre-approved loans in period t. So the pre-approval financing terms in period t are $\Phi_t = (L_t^*, \mu_f)$.

We next describe the laws of motion. We first set up some notation. Let l_0^i denote the stock of homeowners at the time of the crisis with a loan balance equal to the value of the i^{th} grid point L_0^i. Then, l_t^i is the amount of these initial homeowners which still remain in the same house at time t. We let L_t^i denote their loan balance at time t. As people buy homes during the crisis, they flow into a stock which we denote as l_t^0. So $l_0^0 = 0$. According to our assumption, which we verify ex-post, homeowners in the l_t^0 bin do not default along the perfect foresight impulse response. Finally, let L_t^0 denote the average loan size among homeowners in l_t^0 and L_t^b the average pre-approval loan size at time t. Define:

$$\bar{L}_t^0 = l_t^0 L_t^0$$
$$\bar{L}_t^b = v_t^b L_t^b$$

The full set of laws of motion are:

$$
\begin{aligned}
l_{t+1}^{0} &= (1-\gamma)\, l_t^0 + v_t^b q_t^b(\theta_t) \sum_{j=n,d} \frac{v_t^j}{v_t^n + v_r^d}\left(1 - F_t\left(h_t^j\right)\right) \\
l_{t+1}^{i} &= \left(1 - \gamma 1\left[L_t^i \leq V_t^n\right] - \iota_t 1\left[L_t^i > V_t^n\right]\right) l_t^i \\
v_{t+1}^{b} &= \gamma P_t \sum_i l_t^i 1\left[L_t^i \leq V_t^n\right] + \gamma_r P_t v_t^r \\
&\quad + \sigma P_t v_t^f + v_t^b \left[1 - q_b(\theta_t) \sum_{j=n,d} \frac{v_t^j}{v_t^n + v_r^d}\left(1 - F_t\left(h_t^j\right)\right)\right] \\
v_{t+1}^{r} &= (1-\gamma_r P_t)\, v_t^r + \gamma(1-P_t) \sum_i l_t^i 1\left[L_t^i \leq V_t^n\right] + \sigma(1-P_t)\, v_t^f \\
v_{t+1}^{v} &= \gamma \sum_i l_t^i 1\left[L_t^i \leq V_t^n\right] + v_t^n\left[1 - q^s(\theta_t)\left(1 - F_t\left(h_t^n\right)\right)\right] + v_t^a \\
v_{t+1}^{d} &= \iota_t \sum_i l_t^i 1\left[L_t^i > V_t^n\right] + v_t^d\left[1 - q^s(\theta_t)\left(1 - F_t\left(h_t^d\right)\right)\right] \\
v_{t+1}^{f} &= (1-\sigma) v_t^f + \iota_t \sum_i l_t^i 1\left[L_t^i > V_t^n\right] \\
L_{t+1}^{i} &= (1-\gamma_L)\, L_t^i \\
\bar{L}_{t+1}^{b} &= \left(1 - q_t^b(\theta_t) \sum_{j=n,d} \frac{v_t^j}{v_t^n + v_r^d}\left(1 - F_t\left(h_t^j\right)\right)\right) \bar{L}_t^b \\
&\quad + \left[\gamma P_t \sum_i l_t^i 1\left[L_t^i \leq V_t^n\right] + \gamma_r P_t v_t^r + \sigma P_t v_t^f\right] L_t^* \\
\bar{L}_{t+1}^{0} &= (1-\gamma)(1-\gamma_L)\, \bar{L}_t^0 + q_t^b(\theta_t) \sum_{j=n,d} \frac{v_t^j}{v_t^n + v_r^d}\left(1 - F_t\left(h_t^j\right)\right) \bar{L}_t^b
\end{aligned}
$$

Note that we only need to keep track of the total stock of buyers since equilibrium cutoffs do not depend on Φ. We finally describe equilibrium in the mortgage market. Recall that in steady-state, the lender is at its capital requirement. The set of equations is:

$$\Delta_t^E = \iota_t \sum_i l_t^i 1\left[L_t^i > V_t^n\right]\left(V_t^d - L_t^i\right)$$

$$\Delta_t^D = \left[\gamma \sum_i l_t^i 1\left[L_t^i \leq V_t^n\right] + \gamma_r v_t^r + \sigma v_t^f\right] L_t^*$$

$$\Delta_t^Q = \left(\gamma + (1-\gamma)\gamma_L\right) L_t^0 + \gamma \sum_i l_t^i 1\left[L_t^i \leq V_t^n\right] L_t^i + \iota_t \sum_i l_t^i 1\left[L_t^i > V_t^n\right] L_t^i + \sum_i l_t^i \left(1 - \gamma \sum_i l_t^i 1\left[L_t^i \leq V_t^n\right] - \iota_t \sum_i l_t^i 1\left[L_t^i > V_t^n\right]\right) \gamma_L L_t^i$$

$$\Delta_t^D \tilde{P}_t - \Delta_t^Q = \chi \Delta_t^I + \Delta_t^E / \chi$$

$$\Delta_t^I = \mu_f \left[L_t^0 + L_t^b + \sum_i l_t^i \left(1 - \iota_t 1\left[L_t^i > V_t^n\right]\right) L_t^i\right]$$

$$P_t = \tilde{P}_t 1\left[t < T_e\right] + 1\left[t \geq T_e\right]$$

Here, Δ_t^E is the change in the lender's equity position due to losses, Δ_t^D is the total demand for new financing, Δ_t^Q reflects the decreases on the asset side of the balance sheet due to moving, default, and payment, and Δ_t^I is interest earnings. $\tilde{P}_t$ is set so that the lender continues to meet its regulatory capital requirement. At time T_E, the lender is able to recapitalize, setting $P_t = 1$.[36] The model is solved using the Dynare software package. For the cross-cities version of the model, the last equation is replaced by the time path of P_t from the national model.

C Data Sources and Calculations

C.1 Data Sources

The main data source is proprietary data from CoreLogic, which we supplement with data from the U.S. Census, American Housing Survey, Saiz (2010), the Wharton Land-Use Regulation Survey, and the Bureau of Labor Statistics.

CoreLogic provides us with a monthly data set for the nation and the 100 largest CBSAs[37]

[36] Along the equilibrium path, $\tilde{P}_t < 1$ for $t < T_E$.

[37] By CBSA code and name, they are: 10420 Akron, OH; 10580 Albany-Schenectady-Troy, NY; 10740 Albuquerque, NM; 10900 Allentown-Bethlehem-Easton, PA-NJ; 12060 Atlanta-Sandy Springs-Marietta, GA; 12420 Austin-Round Rock-San Marcos, TX; 12540 Bakersfield-Delano, CA; 12580 Baltimore-Towson, MD; 12940 Baton Rouge, LA; 13644 Bethesda-Rockville-Frederick, MD; 13820 Birmingham-Hoover, AL; 14484 Boston-Quincy, MA; 14860 Bridgeport-Stamford-Norwalk, CT; 15380 Buffalo-Niagara Falls, NY; 15764 Cambridge-Newton-Framingham, MA; 15804 Camden, NJ; 16700 Charleston-North Charleston-Summerville, SC; 16740 Charlotte-Gastonia-Rock Hill, NC-SC; 16974 Chicago-Joliet-Naperville, IL; 17140 Cincinnati-Middletown, OH-KY-IN; 17460 Cleveland-Elyria-Mentor, OH; 17820 Colorado Springs, CO; 17900 Columbia, SC; 18140 Columbus, OH; 19124 Dallas-Plano-Irving, TX; 19380 Dayton, OH; 19740 Denver-Aurora-Broomfield, CO; 19804 Detroit-Livonia-Dearborn, MI; 20764 Edison-New Brunswick, NJ; 21340 El Paso, TX; 22744 Fort Lauderdale-Pompano; Beach-Deerfield Beach, FL; 23104 Fort Worth-Arlington, TX; 23420

for 2000-2013 compiled from public records and mortgage data. CoreLogic estimates that its data covers 85 percent of the U.S. Our data set includes:

- The CoreLogic home price index and non-distressed home price index estimated from public records. In cases where public records do not include price (misreported observations or states where the price is not disclosed), this data is supplemented with data on individual mortgages that includes purchase prices. We refer to these as the aggregate and non-distressed price indices. The CoreLogic non-distressed price index drops REO sales and short sales from the database and re-estimates the price index using the same methodology.

- The number of pre-foreclosure filings and completed foreclosure auctions estimated from public records.

- Sales counts for REOs, new houses, non-REO and non-short sale resales, and short sales estimated from public records. Because short sales are not reported separately for much of the time frame covered by the data, we combine short sales and resales into a non-REO existing home sales measure which we call non-distressed sales. We calculate existing home sales by adding REO and non-distressed sales. We also use this data to construct the REO share of existing home volume, which we seasonally adjust.

- Estimates of 7 quantiles of the combined loan-to-value distribution for active mortgages: under 50%, 50%-60%, 60%-70%, 70%-80%, 80%-90%, 90%-100%, 100%-110%, and over 110%. These statistics are compiled by CoreLogic using public records and CoreLogic's valuation models.

Fresno, CA; 23844 Gary, IN; 24340 Grand Rapids-Wyoming, MI; 24660 Greensboro-High Point, NC; 24860 Greenville-Mauldin-Easley, SC; 25540 Hartford-West Hartford-East Hartford, CT; 26180 Honolulu, HI; 26420 Houston-Sugar Land-Baytown, TX; 26900 Indianapolis-Carmel, IN; 27260 Jacksonville, FL; 28140 Kansas City, MO-KS; 28940 Knoxville, TN; 29404 Lake County-Kenosha County, IL-WI; 29820 Las Vegas-Paradise, NV; 30780 Little Rock-North Little Rock-Conway, AR; 31084 Los Angeles-Long Beach-Glendale, CA; 31140 Louisville-Jefferson County, KY-IN; 32580 McAllen-Edinburg-Mission, TX; 32820 Memphis, TN-MS-AR; 33124 Miami-Miami Beach-Kendall, FL; 33340 Milwaukee-Waukesha-West Allis, WI; 33460 Minneapolis-St. Paul-Bloomington, MN-WI; 34980 Nashville-Davidson–Murfreesboro–Franklin, TN; 35004 Nassau-Suffolk, NY; 35084 Newark-Union, NJ-PA; 35300 New Haven-Milford, CT; 35380 New Orleans-Metairie-Kenner, LA; 35644 New York-White Plains-Wayne, NY-NJ; 35840 North Port-Bradenton-Sarasota, FL; 36084 Oakland-Fremont-Hayward, CA; 36420 Oklahoma City, OK; 36540 Omaha-Council Bluffs, NE-IA; 36740 Orlando-Kissimmee-Sanford, FL; 37100 Oxnard-Thousand Oaks-Ventura, CA; 37764 Peabody, MA; 37964 Philadelphia, PA; 38060 Phoenix-Mesa-Glendale, AZ; 38300 Pittsburgh, PA; 38900 Portland-Vancouver-Hillsboro, OR-WA; 39100 Poughkeepsie-Newburgh-Middletown, NY; 39300 Providence-New Bedford-Fall River, RI-MA; 39580 Raleigh-Cary, NC; 40060 Richmond, VA; 40140 Riverside-San Bernardino-Ontario, CA; 40380 Rochester, NY; 40900 Sacramento–Arden-Arcade–Roseville, CA; 41180 St. Louis, MO-IL; 41620 Salt Lake City, UT; 41700 San Antonio-New Braunfels, TX; 41740 San Diego-Carlsbad-San Marcos, CA; 41884 San Francisco-San Mateo-Redwood City, CA; 41940 San Jose-Sunnyvale-Santa Clara, CA; 42044 Santa Ana-Anaheim-Irvine, CA; 42644 Seattle-Bellevue-Everett, WA; 44140 Springfield, MA; 44700 Stockton, CA; 45060 Syracuse, NY; 45104 Tacoma, WA; 45300 Tampa-St. Petersburg-Clearwater, FL; 45780 Toledo, OH; 46060 Tucson, AZ; 46140 Tulsa, OK; 47260 Virginia Beach-Norfolk-Newport News, VA-NC; 47644 Warren-Troy-Farmington Hills, MI; 47894 Washington-Arlington-Alexandria, DC-VA-MD-WV; 48424 West Palm Beach-Boca Raton-Boynton Beach, FL; 48864 Wilmington, DE-MD-NJ; 49340 Worcester, MA.

We seasonally adjust the raw CoreLogic house price indices, foreclosure counts, sales counts, and delinquent and in-foreclosure loan shares using the Census Bureau's X-12 ARIMA software with an additive seasonal factor. For the CBSA sales counts, auctions counts, days on the market, and REO share, we smooth the data using a 5 month moving average (2 months prior, the current month, and 2 months post) to remove spikes in the data caused by irregular reporting at the county level.

For the calibration of the loan balance distribution and initial number of mortgages with high LTV ratios, we adjust the CoreLogic data using data from the American Community Survey as tabulated by the Census. The CoreLogic data only covers all active loans, while our model corresponds to the entire owner-occupied housing stock. Consequently, we use the ACS 3-year 2005-2007 estimates of the owner-occupied housing stock and fraction of houses with a mortgage at the national and county level, which we aggregate to the CBSA level using CBSA definitions.[38] From this data, we construct the fraction of owner-occupied housing units with a mortgage and the fraction of owner-occupied housing units with a second lien or home equity loan. We use these estimates to adjust the loan balance distribution so it represents the entire owner-occupied housing stock and in our regressions to construct the fraction of owner-occupied houses with over 80 percent LTV.

The LTV data is first available for March 2006, which roughly corresponds to the eve of the housing bust as the seasonally-adjusted national house price index reached is peak in March 2006. To approximate the size of the run-up, we average the seasonally-adjusted price index for March-May 2003 and March-May 2006 to calculate the change in log prices for 2003 to 2006.

We also calculate the maximum peak-to-trough log change in seasonally-adjusted aggregate and non-distressed prices, smoothed and seasonally-adjusted non-REO volume, and the average REO share for 2008 to 2013 for each geographical area. We estimate the minimum value between March 2006 and September 2013 and the maximum value between January 2002 and December 2007. We implement these restrictions so that the addition of counties to the CoreLogic data set prior to 2002 does not distort our results. We calculate the fraction of the owner-occupied housing stock that was foreclosed upon by adding up completed foreclosure auctions between March 2006 and September 2013 and dividing by the owner-occupied housing stock in 2006 as calculated from the ACS. Again, our results are not sensitive to the choice of dates.

From the 100 CBSAs we drop the Syracuse, New York CBSA which has incomplete data and the Indianapolis CBSA which has bad data on the 2006 loan balance distribution. We also exclude the Birmingham, Alabama CBSA from calculations on sales volume is dropped because a major county stopped reporting to CoreLogic in the middle of the downturn so the sales volume series is discontinuous. In the main text we exclude two CBSAs in the greater Detroit area—Detroit and Warren—because they experienced large price declines without a prior boom and thus create an exaggerated non-linear relationship between boom size and bust size. Below we show results including these two CBSAs, which are robust to including them. We thus have 97 CBSAs for price and 96 for sales volume.

[38]The 3-year ACS estimates include estimates of the housing stock and houses with a mortgage for all counties with over 20,000 residents. For a few MSAs, one or more small counties are not included in the ACS data. The bias on our constructed estimates of the fraction of owner-occupied homes with a mortgage and with a second lien or home equity loan due to these small missing counties is minimal.

The data on the share of the housing stock converted from owner-occupied to renter-occupied uses data from the American Community survey. Assuming that there are no purpose-built single family detached rental units, the share of the owner-occupied stock converted to renter-occupied from 2006 to 2013 is equal to the stock of single family detached rental homes from 2006 to 2013 divided by the mean stock of single family detached homes in the CBSA for 2005 to 2007. This is likely an upper bound because single-family detached homes were more likely to be converted in this time period. Note that the ACS data is not available at the CBSA level, so we link each CBSA to a MSA (made up of multiple CBSAs in some cases) and give each CBSA the same share converted as the overall MSA.

C.2 Robustness of Empirical Results in Section 1

For robustness tests, we merge data from Saiz (2010) into the CBSA data. The Saiz data includes his estimate of unusable land due to terrain, the housing supply elasticity, and the Wharton Land-Use Regulation Survey score for each CBSA. We are able to match every CBSA we have data on except for Sacramento CA and Honolulu HI. Summary statistics for the complete data set are in Table 10.

Table 10: MSA Summary Statistics

	Unweighted Mean	SD	Min	Max	N
Max $\Delta \log(\text{P})$	-0.326	0.219	-0.950	-0.028	97
Max $\Delta \log(\text{P}_{\text{non-distressed}})$	-0.278	0.195	-0.880	-0.035	97
Max $\Delta \log(\text{Sales}_{\text{non-distressed}})$	-1.124	0.295	-1.918	-0.390	96
$\frac{\text{Sales}_{\text{REO}}}{\text{Sales}_{\text{Existing}}}$	0.307	0.163	0.084	0.796	96
% Foreclosed	0.091	0.073	0.011	0.438	96
$\Delta \log(\text{Price})_{03-06}$	0.306	0.181	0.038	0.729	97
Share LTV > 80%	0.143	0.076	0.026	0.328	97
Frac Second Mort, 06	0.203	0.053	0.026	0.290	97
Saiz Land Unav	0.280	0.213	0.009	0.796	95
Wharton Land Reg	0.228	0.711	-1.239	1.892	95

Notes: Summary statistics for variables used in regression analysis. All data is from CoreLogic. Data is for the 100 largest CBSAs not including Syracuse NY and Detroit and Warren MI. For sales, the REO share and percent foreclosed, Birmingham AL is also omitted. For the Saiz land unavailability and the Wharton land-use regulation, Sacramento CA and Honolulu HI are omitted.

As a robustness test, Table 11 shows regression results when equation (1) is augmented to include the fraction of homes with a second mortgage in 2006, the land unavailability index, and the land use regulation index, and the interacted variable X is the z score of the share of mortgages with an LTV above 80% in 2006. We add the fraction of individuals with a second mortgage or home equity loan to the regression because these loans have received attention in analyses of the downturn (Mian and Sufi, 2011) and we want to make sure they are not driving the result. We use a land unavailability index and the Wharton land use regulation index from Saiz (2010) to proxy for the housing supply elasticity. We estimate:

$$
\begin{aligned}
Y_i \quad = \quad & \beta_0 + \beta_1 \Delta_{03-06} \log(P_i) + \beta_2 \left[\Delta_{03-06} \log(P_i)\right]^2 \\
& + \beta_3 \left(\text{Z Share LTV}_{2006,i} > 80\%\right) + \beta_4 \left(\Delta_{03-06} \log(\text{P}_i) \times \text{Z LTV}_{2006,i} > 80\%\right) \\
& + \beta_5 \left(\text{Z \% Second Mortgage}_{2006,i}\right) + \beta_6 \left(\Delta_{03-06} \log(\text{P}_i) \times \text{Z \% Second}_{2006,i}\right) \\
& + \beta_7 \left(\text{Z Saiz Land Unavailability}_i\right) + \beta_8 \left(\text{Z Wharton Land Use Regulation}_i\right) + \varepsilon_i .
\end{aligned} \tag{22}
$$

The key patterns in the main text continue to be present in this table. Table 12 shows the results are robust to including the two outlier CBSAs in the greater Detroit area that are dropped form the main analysis, although the results are less statistically significant.

C.3 Calibration Target Moments

This section provides details on the target moments used in the calibration in Section 3.

C.3.1 Steady State Moments

We target a mean house price for a non-distressed sales of \$300,000 as an approximation to Adelino et al.'s (2012) mean house price of \$298,000 for 10 MSAs. In reporting results, we normalize this initial house price to 1. The results are not sensitive to this figure, which is a normalization.

As discussed in the text, REO discounts are hard to estimate due to unobserved quality. Most estimates of REO discounts prior to the downturn were approximately 20 percent, but some estimates are closer to 10 percent (see Clauretie and Denshvary, 2009 and Campbell et al., 2011). In the main text we use 12.5 percent, which is a conservative figure that attributes a substantial amount of the discount to unobserved quality, and in the appendix we present results for 10 percent and 15 percent discounts.

We target a time on the market for non-distressed houses of 26 weeks as in Piazzesi and Schneider (2009). This number is a bit higher than some papers that use Multiple Listing Service Data such as Anenberg (2016) and Springer (1996), likely because of imperfect adjustment for withdrawn listings and re-listings.

We target a ratio of buyer to seller time on the market is 1.117 from Genesove and Han's (2012) analysis of National Association of Realtor surveys.

C.3.2 Externally-Calibrated Parameter Values

γ is baed on the median tenure for owner occupants of approximately 9 years comes from table 3-9 of the American Housing Survey National Summary Report and Tables for 1997-2005.

We use $\zeta = 0.65$ in our baseline calibration and $\zeta = 0.70$ in robustness tests. Recall that ζ is the fraction of floor space that a renter occupies relative to an owner. We calibrate this parameter using microdata from the American Housing Survey from 2001 to 2013. Table 13 reports the median (using survey weights) renter-to-owner ratio for a number of different statistics for all renters and for only renters who have moved in the past two years. Across years and measures, the estimates range between 0.5 and 0.7. Figure 7 panel A shows a

Table 11: Cross CBSA Regressions on the Impact of the Size of the Bubble and Its Interaction With High LTV Share

	$\Delta \log(\text{P})$	$\Delta \log(\text{P}_{\text{non-distressed}})$	Mean REO Share	% Foreclosed
$\Delta \log(\text{Price})_{03-06}$	0.586	0.374	-0.819	-0.430
	(0.403)	(0.305)	(0.258)***	(0.131)***
$\Delta \log(\text{Price})^2_{03-06}$	-2.234	-1.843	1.724	1.075
	(0.490)***	(0.370)***	(0.308)***	(0.162)***
Z Share LTV > 80%	0.062	0.068	-0.001	-0.011
	(0.029)**	(0.025)***	(0.022)	(0.014)
$\Delta \log(\text{P}) \times$ Z LTV > 80%	-0.330	-0.346	0.180	0.222
	(0.113)***	(0.112)***	(0.068)***	(0.060)***
Z Frac Second Mort, 06	-0.063	-0.044	-0.018	-0.016
	(0.020)***	(0.018)**	(0.018)	(0.009)*
$\Delta \log(\text{P}) \times$ Z Second	0.118	0.062	0.141	0.066
	(0.076)	(0.073)	(0.058)**	(0.032)**
Z Saiz Land Unav	-0.007	0.001	-0.008	-0.007
	(0.014)	(0.013)	(0.010)	(0.006)
Z Wharton Land Reg	-0.025	-0.022	0.015	0.006
	(0.011)**	(0.010)**	(0.009)*	(0.004)
Constant	-0.259	-0.196	0.256	0.110
	(0.066)***	(0.050)***	(0.044)***	(0.021)***
r^2	0.764	0.777	0.545	0.666
N	95	95	94	94

Notes: * = 10% Significance, ** = 5% Significance *** = 1% significance. All standard errors are robust to heteroskedasticity. Each column shows estimates of (1) as in Table 1 but with no coefficients suppressed. All data is from CoreLogic. Data is for 100 largest CBSAs excluding Syracuse NY, Detroit and Warren MI which are removed so that the outliers do not exaggerate the non-linear relationship, and for columns 3-5 Birmingham AL, which has an inconsistent sales series.

Table 12: Cross CBSA Regressions on the Impact of the Size of the Bubble and Its Interaction With High LTV Share: Including Outliers

	$\Delta \log(\mathrm{P})$	$\Delta \log(\mathrm{P}_{\text{non-distressed}})$	Mean REO Share	% Foreclosed
$\Delta \log(\text{Price})_{03-06}$	1.213	0.903	-0.967	-0.631
	(0.571)**	(0.468)*	(0.290)***	(0.194)***
$\Delta \log(\text{Price})^2_{03-06}$	-2.942	-2.441	1.894	1.303
	(0.683)***	(0.558)***	(0.344)***	(0.233)***
Z Share LTV > 80%	0.057	0.063	-0.005	-0.008
	(0.035)	(0.030)**	(0.023)	(0.015)
$\Delta \log(\mathrm{P}) \times$ Z LTV > 80%	-0.277	-0.299	0.181	0.202
	(0.139)**	(0.134)**	(0.070)**	(0.067)***
Z Frac Second Mort, 06	-0.054	-0.034	-0.011	-0.021
	(0.028)*	(0.024)	(0.019)	(0.011)*
$\Delta \log(\mathrm{P}) \times$ Z Second	0.079	0.025	0.125	0.084
	(0.097)	(0.089)	(0.059)**	(0.037)**
Z Saiz Land Unav	-0.017	-0.007	-0.007	-0.004
	(0.016)	(0.015)	(0.010)	(0.006)
Z Wharton Land Reg	-0.029	-0.025	0.014	0.008
	(0.012)**	(0.011)**	(0.009)	(0.005)
Constant	-0.363	-0.283	0.282	0.142
	(0.093)***	(0.076)***	(0.051)***	(0.031)***
R^2	0.689	0.704	0.524	0.614
N	97	97	95	96

Notes: * = 10% Significance, ** = 5% Significance *** = 1% significance. All standard errors are robust to heteroskedasticity. All data is from CoreLogic. Data is for 100 largest CBSAs excluding Syracuse NY, Sacramento CA, Honolulu HI, and for columns 3-5 Birmingham AL, which has an inconsistent sales series.

Table 13: AHS Microdata: Various Measures of ζ Across Years

Metric Sample	Sqft All	Sqft / Person All	Lot Size All	Lot Size / Person All	Sqft Movers	Sqft / Person Movers	Lot Size Movers	Lot Size / Person Movers
2001	0.529	0.640	0.593	0.705	0.529	0.600	0.511	0.600
2003	0.529	0.656	0.510	0.667	0.529	0.600	0.441	0.50
2005	0.506	0.639	0.786	0.786	0.506	0.587	0.786	0.667
2007	0.500	0.625	0.661	0.667	0.528	0.587	0.620	0.545
2009	0.528	0.625	0.686	0.667	0.544	0.594	0.643	0.583
2011	0.532	0.625	0.707	0.576	0.544	0.583	0.571	0.455
2013	0.542	0.625	0.590	0.554	0.542	0.600	0.556	0.462
Mean	0.524	0.634	0.648	0.660	0.532	0.592	0.590	0.545

Note: This table shows the median owner to renter ratio for the specified metric for each year. The first four columns use all renters, while the last four use only renters who moved in the previous two years. The last row shows the mean across years.

Figure 7: Histograms of Renter to Owner Square Footage Ratios

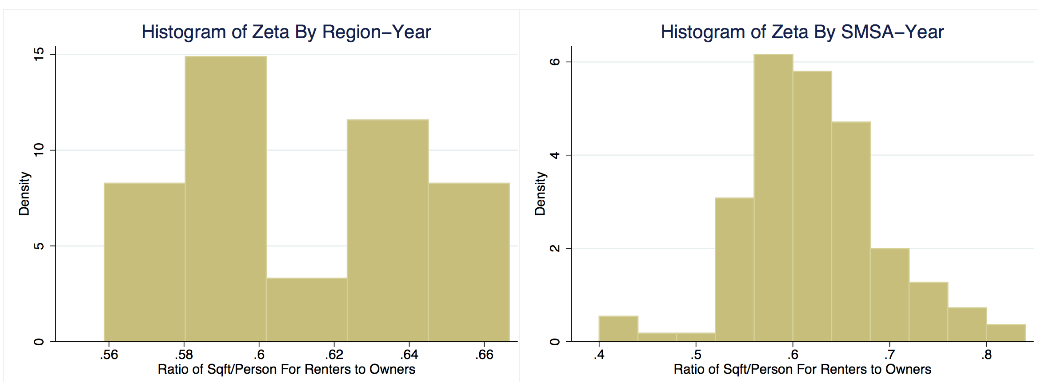

Note: The figure shows histograms of the median ratio of square feet per person for renters to owners. The left panel shows 28 census region × years and the right panel shows 138 SMSA × years with over 100 observations for both renters and owners in each cell.

histogram of the median renter-to-owner ratio of square feet per person by census region × year and panel shows the same histogram by SMSA × year (the AHS uses 1980 SMSA definitions for metropolitan areas) for SMSA × year cells with over 100 observations for both renters and owners. For region × years, the average is 0.61 with a standard deviation of .032. For SMSA × years, the mean is 0.62 with a standard deviation of .075. Because lower numbers imply less conversion and a stronger foreclosure effect, we conservatively choose $\zeta = 0.65$ for our baseline calibration and use $\zeta = 0.70$ for robustness tests.

σ is calibrated based on bank, FHFA, and GSE policies on the amount of time one cannot obtain a loan after a foreclosure as described in the main text. The policies for how long a buyer must wait after a foreclosure to be eligible for a mortgage can be found at http://www.nolo.com/legal-encyclopedia/when-can-i-get-mortgage-after-foreclosure.html and http://www.zillow.com/blog/boomerang-buyers-buying-again-after-a-foreclosure-or-short-sale-102457/.

As described in the text, the size of the government bailout is set to 25% based on a \$130 billion capital injection (Veronesi and Zingales, 2010) relative to losses of \$520 billion (Begenau et al., 2019).

C.3.3 Loan Balance Distribution

The loan balance distribution we use has 51 points distributed in 2% LTV intervals. Recalling that the steady state prices is \$300,000, this implies that the LTV bins have boundaries with intervals of \$6,000. We estimate the mass in each bin by matching the CDF in the model to the CDF implied by CoreLogic data on six quantiles of the combined loan-to-value distribution for active mortgages: under 50%, 50%-60%, 60%-70%, 70%-80%, 80%-90%, 90%-100%, and over 100%. We put all of the mass at or above 100% LTV at 100% LTV. We match the CDF exactly assuming that the mass is distributed equally to each bin within a quantile.

C.3.4 Time Path of Liquidity Shocks ι

We feed in a time path for liquidity shocks ι based on the national long-run unemployment rate. To calculate this rate, we divide the fraction of workers unemployed for 27 weeks or more (which we will call the long-run unemployment rate) by the total labor force as estimated by the Bureau of Labor Statistics. We then take a moving average of the unemployment rate that includes 6 months on either side of the month in question. We shift the series 6 months forward so that the shock corresponds to entering long-run unemployment rather than hitting the 6 month threshold. The time series we choose begins with date zero being August 2008, and we subtract off the long run unemployment rate from this date and normalize this to zero. We continue for 7.6 years until the long-run unemployment rate returns to its August 2008 level.

We adjust the ι series for each CBSA by the ratio of the maximum estimated long-run unemployment rate in the CBSA to the national maximum long-run unemployment rate. The BLS only reports aggregate unemployment rates by CBSA, not unemployment rates by duration. To estimate the local long-run unemployment rate, we use state-level data from the Geographic Profile of Employment and Unemployment Bulletin put out by the BLS and available at https://www.bls.gov/opub/gp/laugp.htm. For each year, we regress the state-level long-run unemployment rate on the state-level unemployment rate and use the coefficient to predict CBSA-level long-run unemployment rates from the CBSA-level unemployment rate. We find the maximum predicted CBSA-level long-run unemployment rate and multiply C_ι for each city by the ratio of this to the national long-run unemployment rate.

D Model Robustness

In this appendix we present robustness checks for the model calibration and present additional results from alternate models and parameterizations. In all cases, we alter a single moment or parameter and repeat the same calibration procedure described in the main text.

The parameters we alter are:

- A 10 percent REO discount rather than the 12.5 percent baseline.
- $\zeta = 0.70$ rather than $\zeta = 0.65$. As described in the appendix, $\zeta = 0.70$ is the upper limit of the ratio of space occupied by renters to the space occupied by owners in the AHS data.
- An average time out of market after a foreclosure of 2.0 years and 3.0 years rather than 2.5 years.
- $\gamma_R = 1/12$ assumes that renters re-apply for pre-approval after an average of 3 months rather than 2 months.

We also solve four alternate models:

- A model in which there is no lender rationing effect and we thus assume $P = 1$ throughout.
- A model in which we slow down foreclosures rather than assuming they occur immediately. We do so by assuming that that foreclosures take on average six months. We allow households to cure exogenously with the average household curing in two years.
- A model in which rather than not moving, households that receive a liquidity shock who are above water sell.
- A model with an alternate law of motion for lender equity that makes the private capital injection additive instead of multiplicative.

Finally, we provide a cross-cities calibration in which we adjust the liquidity shock series to account for heterogeneity across cities in the incidence of liquidity shock by multiplying the series by the ratio of the maximum long-run unemployment rate in each CBSA relative to the maximum-long run unemployment rate nationally.

D.1 Calibration Robustness

This subsection shows results for alternate calibrations for some of the parameters where our baseline parameters are more speculative. Table 14 evaluates the robustness of the role of each effect in reducing non-distressed prices. The top row reports the baseline, which is the same as the non-distressed price index row in Table 5. Each subsequent row reports results obtained from altering the indicated calibration target and recalibrating the model as described in Section 3.4.

The first column reports the model's peak to trough decline in non-distressed prices relative to the CoreLogic national non-distressed national price index. The baseline model almost exactly matches the CoreLogic index even though this is not a targeted moment. The alternate models continue to do well, never straying more than 10 percent from the data.

Table 15 provides results for the cross-city model fit for these various models The top shows the national fit, while the bottom shows the cross-cities fit, which does not change substantially across specifications.

Table 14: Robustness: Decomposition of Non-Distressed Price Index Effects in National Model

Changed Calibration Target	Model Rel to Data	No Default	No Rationing, No Choosey Buyer	No Rationing	Full Model
Baseline	96.92%	51.96%	74.19%	77.40%	100%
10% REO Discount	108.58%	54.99%	74.09%	77.09%	100%
$\gamma_r = 1/12$	97.40%	53.82%	76.64%	79.95%	100%
$\sigma = 1/(2*52)$	98.56%	52.83%	74.87%	78.16%	100%
$\sigma = 1/(3*52)$	96.70%	53.21%	75.94%	79.16%	100%
$\zeta = .7$	98.46%	52.28%	73.94%	77.24%	100%

Note: Each row is for a different robustness check in which one calibration target is changed as indicated in the first column. Each cell indicates the fraction of the full model's peak to trough decline in the non-distressed price index accounted for each column's model relative to the full model. No rationing turns off the lender rationing effect by setting $P_t = 1$. No rationing and no choosey buyer additionally turns off the choosey buyer effect by altering the buyer's value function so she does not take into account the possibility of meeting an REO seller, which leaves only the foreclosure flag effect. No default turns off all of the effects and entirely eliminates default.

D.2 Alternate Model Results

D.2.1 No Lender Rationing Effect

We now turn to a model with no lender rationing effect. We implement this by setting $P_t = 1$ and not imposing the equilibrium condition for P_t. This means that two of our parameters in fitting the model to the downturn, the capital requirement and the date that lenders regain the ability to raise equity, do nothing. Consequently, we use the other two parameters, the scaler for the liquidity shocks and the decline in the subjective valuation of houses, to match the decline in the aggregate price index and the fraction of the cumulative housing stock foreclosed upon.

The results are shown in Table 16 for the national model. One can see clearly how important the bank rationing effect is for matching non-distressed sales volume, as the model does very poorly for this metric.

D.2.2 Slowed Down Rate of Foreclosure

Table 17 shows the results of a model in which we relax the assumption that foreclosures occur immediately. Instead, we assume that there is a stochastic delay for foreclosures, which occur each period with probability $\sigma_f = 1/26$. This is a slightly different formulation than the one we use in Section 6.1, where we assume there is a maximum number of foreclosures that can be completed each period, but it is more numerically convenient and so we can introduce it into the cross-cities calibration. We are assuming that the average foreclosure takes six months to process. We also let $\omega = 1/(2*52)$ so that the average person cures out of their liquidity shock after two years. We then recalibrate everything as in the main text. One can see that the national model fit remains similar. For the cross-cities, for numerical reasons we do not include heterogeneity in unemployment by CBSA and instead assume $\iota_t = C_i Unemp_t$ for all CBSAs as in Appendix D.3. The cross-city results are very

Table 15: Robustness: Cross-CBSA Simulations vs. Data

Variable	$\Delta \log P$	$\Delta \log (\mathrm{P}_{\mathrm{nd}})$	Mean REO Share	% Foreclose	% Convert
National Data	-0.388	-0.298	0.190	8.34	4.35
Baseline Model	-0.388	-0.289	0.191	8.34	7.38
10% REO Discount	-0.388	-0.324	0.190	8.33	7.27
$\gamma_r = 1/12$	-0.388	-0.290	0.190	8.35	7.28
					7.34
$\sigma = 1/(3*52)$	-0.388	-0.288	0.190	8.35	7.32
$\zeta = .7$	-0.388	-0.293	0.191	8.34	7.87
Baseline Reg Coef of	1.005	1.389	0.823	0.992	0.435
Data on Model	(0.063)	(0.087)	(0.090)	(0.069)	(0.055)
R^2	0.736	0.734	0.482	0.695	0.419
10% REO Discount	1.037	1.244	0.809	0.971	0.446
Coef	(0.064)	(0.079)	(0.086)	(0.067)	(0.057)
R^2	0.742	0.732	0.497	0.699	0.414
$\gamma_r = 1/12$	1.005	1.385	0.828	0.991	0.443
Coef	(0.063)	(0.087)	(0.090)	(0.069)	(0.056)
R^2	0.736	0.734	0.483	0.695	0.422
$\sigma = 1/(2*52)$	1.037	1.369	0.860	1.017	0.397
Coef	(0.064)	(0.089)	(0.094)	(0.071)	(0.052)
R^2	0.742	0.72	0.482	0.689	0.399
$\sigma = 1/(3*52)$	0.941	1.357	0.771	0.939	0.492
Coef	(0.061)	(0.084)	(0.085)	(0.065)	(0.059)
R^2	0.721	0.741	0.474	0.694	0.443
$\zeta = .7$	1.013	1.373	0.828	1.007	0.464
Coef	(0.063)	(0.087)	(0.090)	(0.070)	(0.060)
R^2	0.740	0.729	0.485	0.694	0.409

Note: Each column shows a comparison of the model and data for the given variable. Each set of rows if for a different model as indicated. The top comparisons compare the data and model for the national simulation. The bottom comparisons show the slope term of a regression of the actual data on the model simulated data. Standard errors are in parenthesis. This table does not include stars for statistical significance as almost all coefficients are significant at the 1% level.

Table 16: No Lender Rationing Effect

Variable	$\Delta \log P$	$\Delta \log (\mathrm{P}_{\mathrm{nd}})$	$\Delta \log (\mathrm{Sales}_{nd})$	Mean REO Share	% Foreclose	% Convert
National Data	-0.388	-0.298	-1.029	0.190	8.34	4.35
No Lender Rationing	-0.389	-0.321	-0.581	0.159	8.35	3.93

Note: Each column shows a comparison of the model and data for the given variable for the national simulation.

Table 17: Slowed Down Rate of Foreclosure Calibration

Variable	$\Delta \log P$	$\Delta \log (\mathrm{P_{nd}})$	Mean REO Share	% Foreclose	% Convert
National Data	-0.388	-0.298	0.190	8.34	4.35
National Model	-0.388	-0.289	0.191	8.34	8.34
Reg Coef of	1.146	1.405	0.993	1.246	0.359
Data on Model	(0.071)	(0.089)	(0.107)	(0.086)	(0.046)
R^2	0.742	0.731	0.488	0.694	0.409

Note: Each column shows a comparison of the model and data for the given variable. The top comparisons compare the data and model for the national simulation. The bottom comparisons show the slope term of a regression of the actual data on the model simulated data. Standard errors are in parenthesis.

Table 18: Above-Water Households With Liquidity Shock Sell Calibration

Variable	$\Delta \log P$	$\Delta \log (\mathrm{P_{nd}})$	Mean REO Share	% Foreclose	% Convert
National Data	-0.388	-0.298	0.190	8.34	4.35
National Model	-0.388	-0.325	0.132	8.34	32.60
Reg Coef of	1.028***	1.204***	1.136***	0.922***	0.709***
Data on Model	(0.068)	(0.085)	(0.135)	(0.068)	(0.154)
R^2	0.716	0.687	0.440	0.668	0.198

Note: Each column shows a comparison of the model and data for the given variable. The top comparisons compare the data and model for the national simulation. The bottom comparisons show the slope term of a regression of the actual data on the model simulated data. Standard errors are in parenthesis.

similar to the results in Appendix D.3. We conclude that our model is robust to adding some foreclosure backlogs.

D.2.3 Above-Water Households With Liquidity Shock Sell

In the main text we assume that above-water households with an liquidity shock are able to stay in their house. In Table 18, we report results from a variant of the model in which we instead assume these households sell and rent until they cure out of their liquidity shock, which occurs with probability $\omega = 1/52$ and thus takes on average one year.

We follow the same calibration procedure except we do not calibrate to the mean REO share and instead fix T_E, the date at which lenders regain access to the equity market, at 3 years. We do so because if we try to fit the REO share, we end up with a calibration with a very long time until capital markets reopen and a *positive* a^{frac}, indicating that when the downturn starts prices permanently *rise*. This does not make sense, so we abandon matching the mean REO share. In the calibration where we ignore the mean REO share, the mean REO share is 13.2% instead of 19.0% in the data.

The model where above-water households with liquidity shocks sell does reasonably well matching the cross-sectional patterns. However, one can see that the model is off by an order of magnitude for the share of the housing stock converted to rental. This is because many

Table 19: Parameter Values Calibrated to Downturn Moments for Alternate Law of Motion

Param	Value
a^{frac}	0.728
T_E	204 Weeks
Δ_E	0.148
C_ι	0.293

households who are above water experience liquidity shocks, sell their house, and are forced to rent. The fact that the model overshoots this target by an order of magnitude – and by implication overstates the decline in the homeownership rate – suggests our assumption that households who are above water and receive an liquidity shock do not sell is more appropriate.

D.2.4 Alternate Law of Motion For Lender Equity

In the main text, we consider a law of motion for lender equity of the form:

$$E_{t+1} = E_t + (1-\psi)\iota_t l_t \int 1\left[L_t > V_t^n\right]\left(V_t^d - L_t\right) dG\left(L_t\right) + \Delta_t^R.$$

We let $\psi = \psi_G + \psi_E$, where ψ_G is a government bailout as a fraction of losses, which is calibrated to 25% based on Begenau et al. (2019), and ψ_E is private equity that can be raised as a fraction of losses, which we calibrate. We prefer this approach because it puts government and private equity injections on equal footing.

In this appendix, we consider an alternate approach as a robustness check which uses a law of motion:

$$E_{t+1} = E_t + \left(1-\psi_G\right)\iota_t l_t \int 1\left[L_t > V_t^n\right]\left(V_t^d - L_t\right) dG\left(L_t\right) + \Delta_t^R + \Delta^E.$$

Again ψ_G is a government bailout as a fraction of losses, which again we set to 25%. Now Δ^E are private equity injections. The difference here is that Δ^E is additive rather than multiplicative. The calibration procedure is the same with Δ^E replacing ψ_E.

The alternate law of motion for lender equity yields broadly similar results. The lender rationing effect is weakened and more of the bust is explained by an exogenous change in prices, but quantitatively these changes are minute.

Table 19 shows the parameters for the model. One can see that they are similar, although a^{frac} and C_ι are both somewhat smaller. The 27 percent fall in $\bar{a}$ implies that prices fall permanently by 13.26 percent nationally due to the valuation shock, up from 11.91 percent in the main text.

Table 20 repeats the decomposition of effects in Table 5. The model without any default account explains 56.56% of the peak-to-trough decline in non-distressed prices. The foreclosure flag effect explains an additional 22.68%, and the choosey buyer effect explains an additional 3.40%. Lender rationing explains the remaining 17.36%. Relative to the main text, the foreclosure flag and choosey buyer effects are similar, lender rationing explains

Table 20: Decomposition of Effects in National Model for Alternate Law of Motion

Statistic (Peak to Trough Unless Indicated)	No Liquidity Shocks	No Rationing, No Choosey Buyer	No Rationing	Full Model
Price Index	44.81%	76.67%	78.32%	100%
Non-Distressed Price Index	56.57%	79.24%	82.64%	100%
Non-Distressed Sales Volume	58.01%	62.02%	82.57%	100%
Average REO Share	0%	69.38%	70.92%	100%
Total Foreclosures	0%	82.01%	83.11%	100%

Note: Each cell indicates the fraction of the full model accounted for each column's model relative to the full model. No rationing turns off the lender rationing effect by setting $P_t = 1$. No rationing and no choosey buyer additionally turns off the choosey buyer effect by altering the buyer's value function so she does not take into account the possibility of meeting an REO seller, which leaves only the foreclosure flag effect. No default turns off all of the effects and entirely eliminates default. This table uses the alternate law of motion.

roughly 5% less of the decline in non-distressed prices, and the model without default explains roughly 5% more. Table 21 evaluates the robustness of the role of each effect in reducing non-distressed prices. Again, the results are largely robust across a range of calibrations.

Finally, the cross-cities analysis looks similar with the alternate law of motion for lender equity. Figure 8 repeats the analysis in Figure 5 and shows our simulated results from the baseline model with the alternate law of motion against actual data for 96 CBSAs (black dots) and the national model (red X). Table 22 repeats the analysis in Table 6. As with the main text, the model fits well, with the data points clustering around the 45-degree line across the spectrum of price declines. Indeed, when we regress the simulated data on the actual data we get a coefficient of 0.995, and we cannot statistically reject a coefficient of one and an intercept of zero. Panel B shows the fraction of the housing stock foreclosed upon over eight years. This is a moment used in the national calibration, so the national model is fits almost exactly, but it is not a target for the cross-section. Nonetheless, when we regress the simulated data on the real data, we get a coefficient of .981 and we cannot reject a coefficient of one and an intercept of zero. The limited heterogeneity in our model thus does a good job capturing the variation in default across cities. The other results in Table 22 are similar to the main text as well, as are the contrast between the model with and without default across cities as in Figure 10.

D.3 Cross-Cities Calibration With No Heterogeneity in Unemployment By CBSA

In the main text CBSAs differ along three dimensions: the initial loan balance distribution, the size of the permanent price decline, and their long-term unemployment rate, which scales the liquidity shock series. In this appendix, we remove heterogeneity in the liquidity series. This pedagogical exercise is meant to show how important it is to include heterogeneity in long-term unemployment by CBSA.

Tables 23 and 9 and Figures 10 and 24 reproduce Tables 6 and 7 and Figures 5 and 6 in the main text for this specification. Removing this heterogeneity does reduce the quantitative fit

Table 21: Robustness: Decomposition of Non-Distressed Price Index Effects in National Model for Alternate Law of Motion

Changed Calibration Target	Model Rel to Data	No Default	No Rationing, No Choosey Buyer	No Rationing	Full Model
Baseline	99.72%	56.57%	79.24%	82.64%	100%
10% REO Discount	108.02%	56.80%	75.45%	78.02%	100%
$\gamma_r = 1/12$	105.93%	59.84%	81.87%	85.29%	100%
$\sigma = 1/(2*52)$	99.97%	59.48%	75.63%	78.79%	100%
$\sigma = 1/(3*52)$	99.52%	52.26%	81.39%	84.97%	100%
$\zeta = .7$	100.90%	62.56%	81.42%	84.54%	100%
Cap Ratio 12.5%	102.13%	58.11%	80.32%	83.70%	100%

Note: Each row is for a different robustness check in which one calibration target is changed as indicated in the first column. Each cell indicates the fraction of the full model's peak to trough decline in the non-distressed price index accounted for each column's model relative to the full model. No rationing turns off the lender rationing effect by setting $P_t = 1$. No rationing and no choosey buyer additionally turns off the choosey buyer effect by altering the buyer's value function so she does not take into account the possibility of meeting an REO seller, which leaves only the foreclosure flag effect. No default turns off all of the effects and entirely eliminates default. This table uses the alternate law of motion.

Figure 8: Cross-CBSA Simulations vs. Data for Alternate Law of Motion

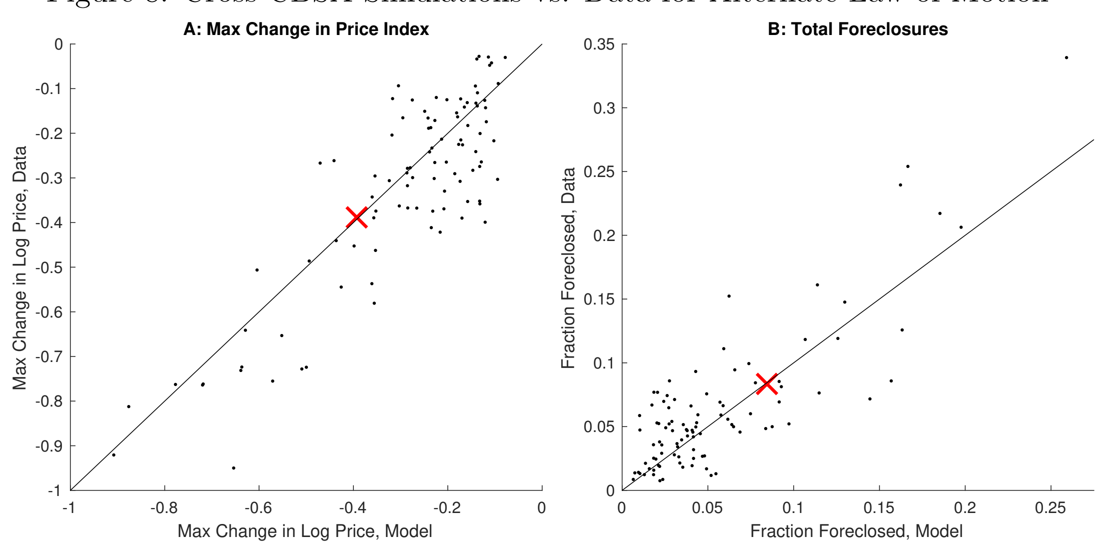

Note: Each panel shows scatter plots of data vs simulation results for 96 CBSAs in regression analysis. The red X represents the national simulation and each black dot is a CBSA. The 45-degree line illustrates a perfect match between the model and the data. The variable being plotted shown in each plot's title. The data is described in the appendix. The calibration methodology, which fits the cross-cities model only to the aggregate price decline in panel A, is described in text. The price decline is the maximum peak-to-trough change, while the fraction foreclosed which is the total from 2006 to 2013. This figure uses the alternate law of motion.

Table 22: Cross-CBSA Simulations vs. Data

Variable	$\Delta \log P$	$\Delta \log (\mathrm{P_{nd}})$	Mean REO Share	% Foreclose	% Convert
Reg Coef of	0.995	1.351	0.849	0.981	0.449
Data on Model	(0.062)	(0.085)	(0.092)	(0.068)	(0.056)
R^2	0.735	0.735	0.484	0.695	0.424

Note: Each column shows a comparison of the model and data for the given variable. The comparisons show the slope term of a regression of the actual data on the model simulated data. Standard errors are in parenthesis. This table uses the alternate law of motion.

Figure 9: No Unemployment Heterogeneity: Cross-CBSA Simulations vs. Data

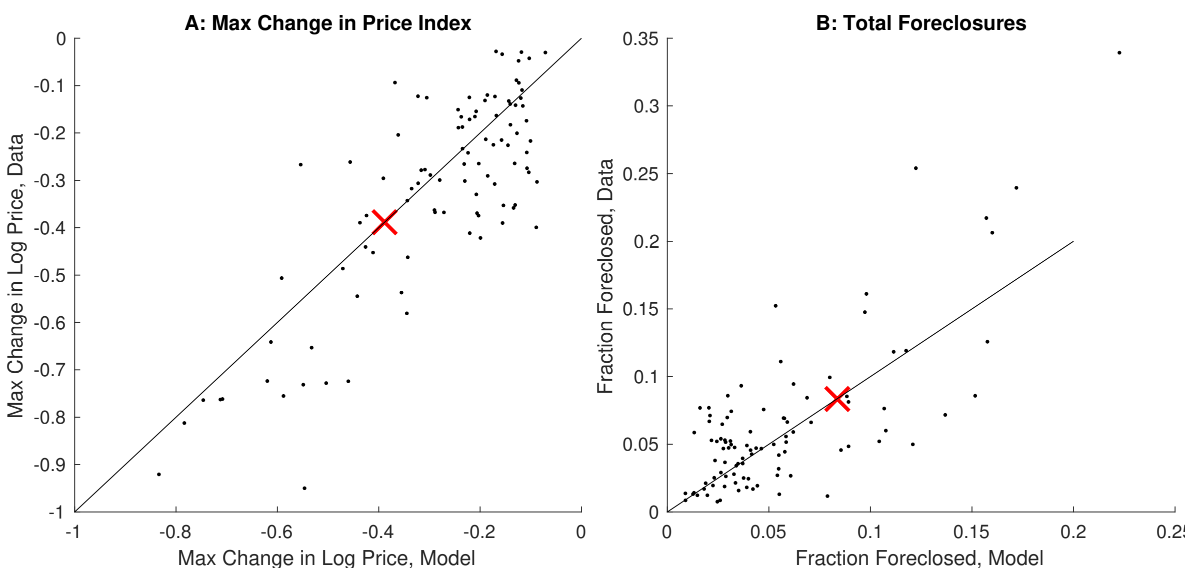

Note: Scatter plots of data vs simulation results for 96 CBSAs in regression analysis. The red X represents the national simulation and each black dot is a CBSA. The 45-degree line illustrates a perfect match between the model and the data. The variable being plotted shown in each plot's title. The calibration methodology, which fits the cross-cities model only to the aggregate price decline in panel A, is described in text. The price decline is the maximum peak-to-trough change while the fraction foreclosed which is the total from 2006 to 2013.

Table 23: No Unemployment Heterogeneity: Cross-CBSA Simulations vs. Data

Variable	$\Delta \log P$	$\Delta \log (\mathrm{P}_{\mathrm{nd}})$	Mean REO Share	% Foreclose	% Convert
National Data	-0.388	-0.298	0.190	8.34	4.35
National Model	-0.388	-0.289	0.191	8.34	7.38
Reg Coef of	0.984	1.298	0.758	1.016	0.353
Data on Model	(0.073)	(0.091)	(0.109)	(0.091)	(0.053)
R^2	0.667	0.688	0.349	0.579	0.341

Note: Each column shows a comparison of the model and data for the given variable. The top comparisons compare the data and model for the national simulation. The bottom comparisons show the slope term of a regression of the actual data on the model simulated data. Standard errors are in parenthesis.

Figure 10: No Unemployment Heterogeneity Boom vs. Bust in Baseline and No Default Models

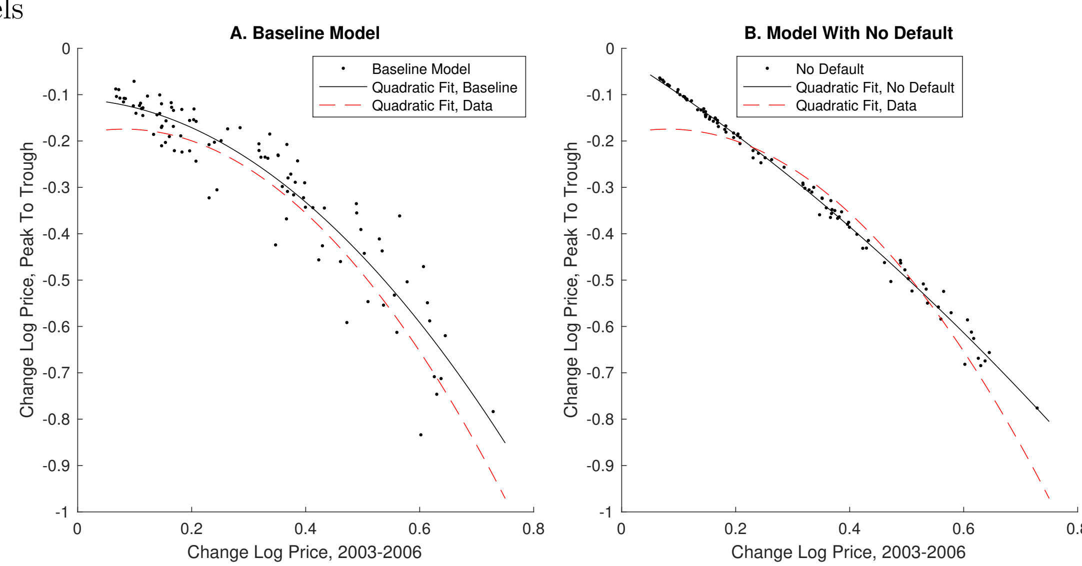

Note: The left panel shows the size of the boom vs. the model simulated data for the baseline calibrated model with default, while the right panel shows the same plot for the no default model. The black solid line is a best quadratic fit. The red dashed line shows the best quadratic fit to the actual data.

Table 24: No Unemployment Heterogeneity Model vs. Data: Interaction of Size of Boom With High LTV Share

$\Delta \log(\text{P}) \times$ Z LTV > 80%	$\Delta \log(\text{P})$	$\Delta \log(\text{P}_{\text{nd}})$	Mean REO Share	% Foreclosed
Baseline	-0.700	-0.139	0.298	0.154
	(.018)***	(0.008)***	(0.018)***	(.008)***
No Default	-.102	-.102		
	(.004)***	(.004)***		
Data	-0.310	-0.336 139	0.205	0.235
	(0.123)**	(0.113)***	(0.075)***	(0.054)***

Notes: * = 10% Significance, ** = 5% Significance *** = 1% significance. All standard errors are robust to heteroskedasticity. Each column shows estimates of (1) with the constant suppressed. P_{nd} is a non-distressed only price index. The mean REO share of sales volume is the average from 2008 to 2012, and the fraction foreclosed is the fraction of the housing stock foreclosed upon over the first 8 years of the downturn. All data is from CoreLogic and described in the appendix. Data is for 97 largest CBSAs excluding Syracuse NY, Detroit MI, Warren MI, and, for columns 3-5, Birmingham AL. All dependent variables are peak-to-trough maximums with the exception of percentage foreclosed upon.

of the model somewhat. The r-squared of regressing the actual data on the simulated data for aggregate price is 0.67 rather than 0.74 with heterogeneity by CBSA. The r-squared for other variables falls as well. The regression coefficients are also further from one. However, the non-linearity in boom-size relative to the size of the bust is quite similar in Figure 10 and Figure 6, suggesting that heterogeneity in long-term unemployment does not drive this non-linearity

We conclude that adding heterogeneity in liquidity shocks by CBSA based on the CBSA's unemployment rate helps the quantitative fit but is not crucial for generating the non-linearity in the model.

E Policy Appendix

E.1 Costs of Government Policies

In the text, we report per-household costs of various foreclosure policies.

The present-value cost of government equity injections is computed as:

$$\sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t \psi_G \iota_t l_t \int 1\left[L_t > V_t^n\right]\left(V_t^d - L_t\right) dG\left(L_t\right),$$

where ψ_G is the fraction of book losses covered by the government. Recall that

$$\iota_t l_t \int 1\left[L_t > V_t^n\right]\left(V_t^d - L_t\right) dG\left(L_t\right)$$

are the total book losses experienced by the lender in a given period t, where ι_t is the incidence of liquidity shocks, only homeowners who are underwater default, and V_t^d is the value that the lender can recover from the home. We set $r = \beta^{-1} - 1$.

The cost of financing the principal reduction is equal to the cost of the baseline 25% equity injection plus the cost of principal reductions of Φ dollars up to an LTV limit of 115. The cost of the principal reduction is given by:

$$\int min\left(\Phi, L_0 - 1.15V_0^n\right) 1\left[L_0 > 1.15V_0^n\right] dG\left(L_0\right),$$

where $G\left(L_0\right)$ is the initial loan balance distribution and V_0^n is the price a homeowner can receive for her home at the onset of the crisis.

Finally, we describe how we implement the policy in which the government purchases distressed homes from lenders and then holds them off the market for a period of time. First, we introduce a new state variable v_t^g, denoting the stock of homes held by the government. The laws of motion for v_t^g and the adjusted law of motion for v_t^d become:

$$\begin{aligned}
v_{t+1}^d &= \sigma_g v_t^g + v_t^d \left[1 - q^s\left(\theta_t\right)\left(1 - F_t\left(h_t^d\right)\right)\right] \\
v_{t+1}^g &= \iota_t \sum_i l_t^i 1\left[L_t^i > V_t^n\right] + \left(1 - \sigma_g\right) v_t^g,
\end{aligned}$$

where σ_g is the rate at which the government re-introduces homes to the market. Finally, we let

$$V_t^g = m_d + \beta\left(1 - \sigma_g\right) E_t V_{t+1}^g + \beta \sigma_g E_t V_{t+1}^d$$

denote the value of a home in the government stock. Note that the government pays the flow cost m_d, reflecting the higher depreciation associated with vacant homes. Upon re-introducing the home to the market, the government receives V_t^d, the value of a distressed home. The present-value cost of this policy to the government along the perfect foresight path is:

$$\sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t \iota_t l_t \int 1\left[L_t > V_t^n\right]\left(V_t^g - V_t^d\right) dG\left(L_t\right).$$

E.2 Payment Reductions

In this appendix, we consider policies which offer payment relief by subsidizing lower interest payments holding the principal balance fixed. This is a policy intervention that our model is not as well-suited to consider since we abstract from household balance sheets. Nonetheless, we can use results from the empirical literature on how payment reductions affect the incidence of default in partial equilibrium and feed this into the model to get some understanding of the aggregate effect of the policy, taking into account equilibrium effects. Because this exercise is not fully internal to our model, though, it is more difficult to compare to the other policies we consider on even footing.

Specifically, we use results from Fuster and Willen (2017), who provide quasi-experimental empirical evidence on the effect of payment reductions on default by exploiting plausibly exogenous differences in the timing of rate resets among a sample of homeowners with ALT-A hybrid adjustable rate mortgages between 2005-2008. They document that a 0.50% reduction in the interest rate reduces the default hazard by 10% while a 1.0% reduction leads to a 20%

Table 25: Effects of Payment Reductions

Rate Reduction	Decline in C_ι	$\Delta \log P_{nd}$	Total Foreclosures
No Rate Reduction	0%	-0.289	8.34%
0.5% Rate Reduction	10%	-0.274	7.51%
1.0% Rate Reduction	20%	-0.259	6.67%
2.0% Rate Reduction	40%	-0.231	5.00%

Notes: The table shows the peak-to-trough decline in log non-distressed prices and total foreclosures from 2006 to 2013 under each policy intervention.

decline in the default hazard. We introduce this evidence into our model by assuming that the government covers a certain percentage of the interest that underwater homeowners owe to lenders and reducing C_ι, the scaling parameter of the liquidity shock series $\iota_t = C_\iota Unemp_t$ and hence the default hazard, by the amount indicated by Fuster and Willen's analysis.

Our results are in Table 25. A 0.5% rate reduction leads to a 5.27% smaller non-distressed price decline, a 1.0% rate reduction to a 10.40% smaller decline, and a 2.0% rate reduction to a 21.06% smaller decline. Intuitively, payment reductions empirically reduce the incidence of default, which when added to the model reduces the severity of the equilibrium price-default spiral. Note that this policy is imperfectly targeted as with principal forgiveness. Homeowners who would never have defaulted receive benefits, which decreases the cost-effectiveness of the policy.