

How Do Foreclosures Exacerbate Housing Downturns?

Adam Guren and Tim McQuade*

Harvard University

April 30, 2013

Abstract

The recent housing bust precipitated a wave of mortgage defaults, with over seven percent of the owner-occupied housing stock experiencing a foreclosure. This paper presents a model that shows how foreclosures can exacerbate a housing bust and delay the housing market's recovery. By raising the ratio of sellers to buyers, by making buyers more selective, and by changing the composition of houses that sell, foreclosures freeze up the market for retail (non-foreclosure) sales and reduce both price and volume. Because negative equity is necessary for default, these general equilibrium effects on prices can create price-default spirals that amplify an initial shock. To assess the magnitude of these channels, the model is calibrated to simulate the downturn. The amplification channel is significant. The model successfully explains aggregate and retail price declines, the foreclosure share of volume, and the number of foreclosures both nationwide and across MSAs. While the model can explain variation in sales across MSAs, it cannot account for the aggregate level of the volume decline, suggesting that other forces have reduced sales nationwide. The quantitative analysis implies that from 2007 to 2011 foreclosures exacerbated aggregate price declines by approximately 50 percent and declines in the prices of retail homes by approximately 30 percent.

*guren@fas.harvard.edu, tmcquade@fas.harvard.edu. We would like to thank John Campbell, Edward Glaeser, and Jeremy Stein for outstanding advice and Joe Altonji, Pat Bayer, Karl Case, Raj Chetty, Emmanuel Farhi, Veronica Guerrieri, Erik Hurst, Loukas Karabarbounis, Lawrence Katz, Guido Lorenzoni, Greg Mankiw, Chris Mayer, Atif Mian, Pascal Noel, Robert Novy-Marx, Andrei Shleifer, Alp Simsek, Francesco Trebbi, and seminar participants at Harvard University, the NBER Summer Institute Housing and the Financial Crisis session, 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.

1 Introduction

Foreclosures have been one of the dominant features of the recent housing market downturn. From 2006 through 2011, approximately 7.4 percent of the owner-occupied housing stock experienced a foreclosure.¹ Although the wave of foreclosures has subsided, foreclosures remain at elevated levels, and understanding the role of foreclosures in housing downturns remains an important part of reformulating housing policy going forward.

The behavior of the housing market concurrent with the wave of foreclosures is shown in Figure 1. Real Estate Owned (REO) sales – that is sales of foreclosed homes owned by banks and the GSEs – have made up between 20 and 30 percent of existing home sales nationally. Sales of existing homes fell 54.9 percent peak-to-trough; retail (non-foreclosure) volume fell 65.7 percent. Prices dropped considerably, with aggregate price indices plunging by a third and prices falling by a quarter for indices that exclude distressed sales. Time to sale and vacancy rates have also climbed, particularly in the retail market. Even with a slowdown in foreclosures due to lawsuits over fraudulent foreclosure practices, foreclosures have continued at a ferocious pace.

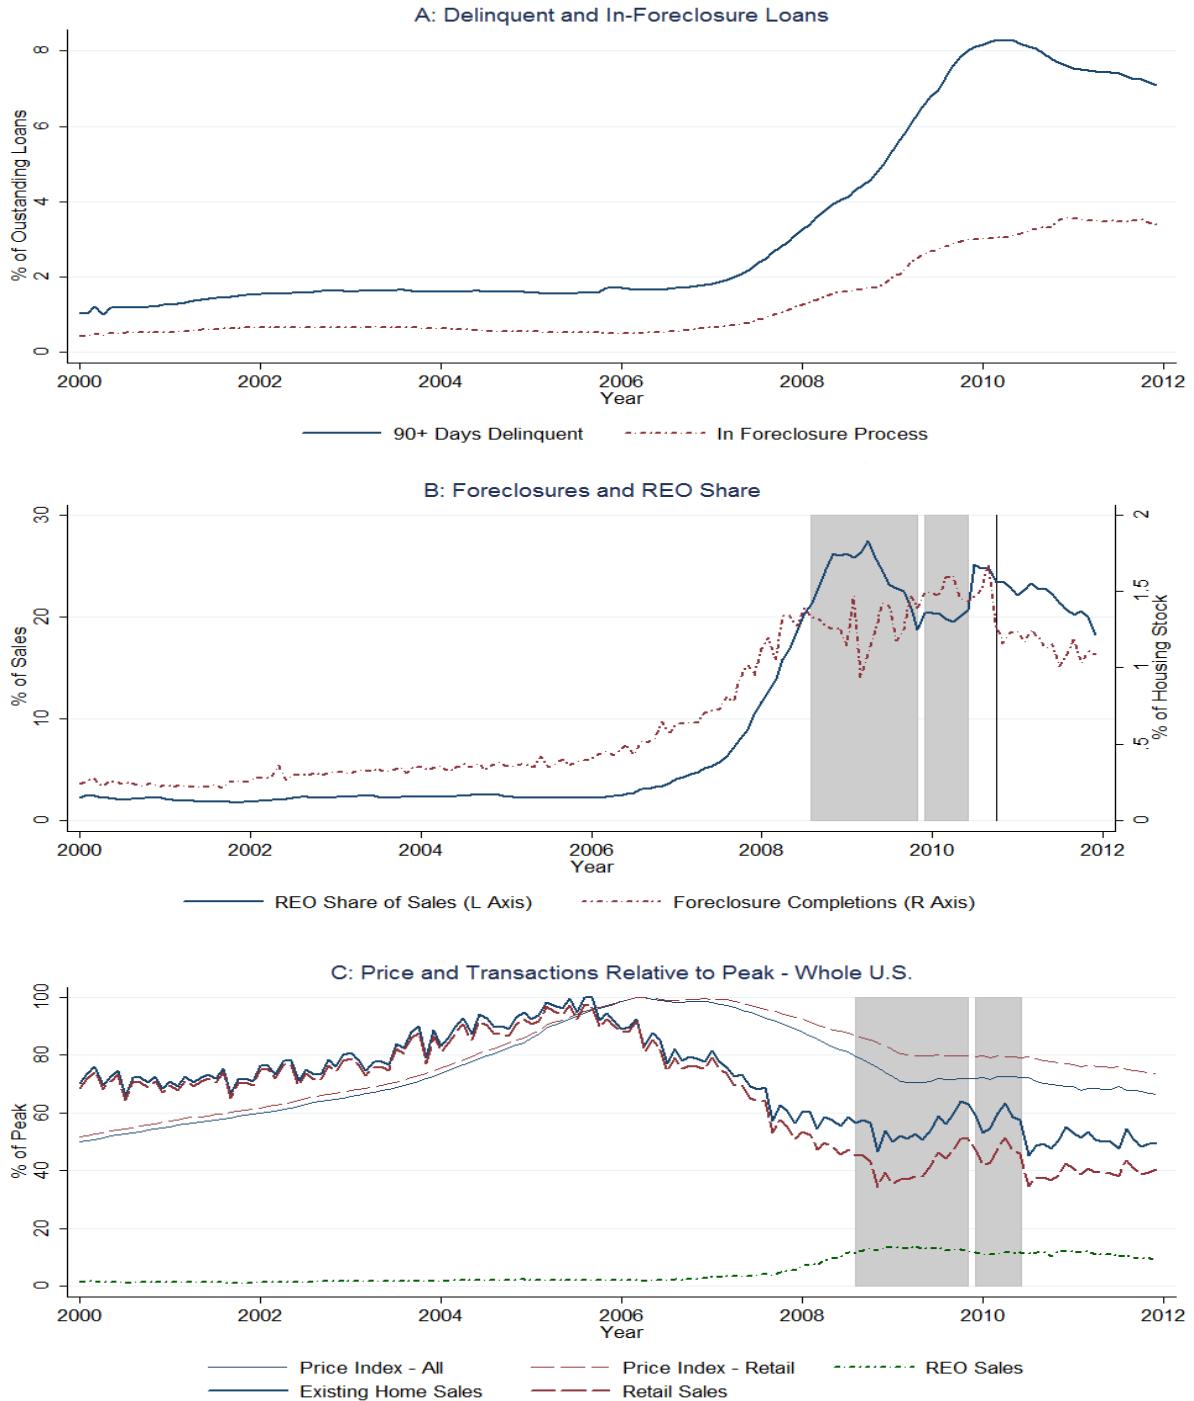
This paper presents a model in which foreclosures have important general equilibrium effects that can explain much of the recent behavior of housing markets, particularly in the hardest-hit areas. By raising the number of sellers and reducing the number of buyers, by making buyers more choosy, and by changing the composition of houses that sell, foreclosures sales freeze up the market for retail sales and reduce both price and sales. Furthermore, the effects of foreclosures can be amplified considerably because price declines induce more default which creates further price declines, generating a feedback loop. A quantitative calibration suggests that these effects can be large: foreclosures exacerbate aggregate price declines by approximately 50 percent and retail price declines by 30 percent.

Despite the importance of foreclosures in the housing downturn, economists have not closely examined how the housing market equilibrates when there are a substantial number of distressed sales. A supply and demand framework, as employed by much of the financial literature on fire sales and illiquidity, can potentially explain declining prices and volumes with demand falling relative to supply but cannot speak to the freezing up of the retail market. Such models also assume that investors can adjust their positions continuously by transacting in a liquid market, yet housing is lumpy, illiquid, and expensive. A substantial literature has sought to adapt models to fit the peculiarities of the housing market and explain the positive correlation between volume and price. For instance, search frictions as in Wheaton (1990), Williams (1995), Krainer (2001), and Novy-Marx (2009), borrowing constraints as in Stein (1995), and nominal loss aversion as in Genesove and Mayer (2001) have been shown to play important roles in housing markets. Yet no paper has explicitly examined the role of distressed sales in a model tailored to housing.

To illustrate the mechanisms through which foreclosures affect the housing market, a simple model of the housing market with exogenous foreclosures is introduced. It adds two key ingredients to an otherwise-standard search-and-matching framework with stochastic moving shocks, random search, idiosyncratic house valuations, and Nash bargaining over price: REO sellers have higher holding costs and individuals who are foreclosed upon cannot immediately buy a new house. These two additions together dry up the market for normal sales, reduce volume and price, and imply that the market only gradually recovers from a wave of foreclosures. This occurs through three main effects. First, the presence of distressed sellers increases the outside option to transacting for buyers, who have an elevated probability of being matched with a distressed

¹Data from CoreLogic. The data is described in Section 5 and Appendix A.4.

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 data appendix. The grey bars in panels B and C show the periods in which the new homebuyer tax credit applied. The black line in panel B shows when foreclosures were stalled due to the exposure of fraudulent foreclosure practices by mortgage servicers. In panel C, all sales counts are unsmoothed and normalized by the total number of existing home sales at peak while each price index is normalized by its separate peak value.

seller next period and consequently become more selective. This “choosey buyer effect” endogenizes the degree of substitutability between bank and retail sales. Second, because foreclosed individuals are locked out of the market, foreclosures reduce the likelihood that a seller will meet a buyer in the market through a “market tightness effect.” This effect emphasizes that foreclosures do not simply add supply to the market: a key feature of foreclosures is that they also reduce demand. Third, there is a mechanical “compositional effect” as the average sale looks more like a distressed sale.

The choosey buyer effect in particular is novel and formalizes folk wisdom in housing markets that foreclosures empower buyers and cause them to wait for a particularly favorable transaction. 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.’”² 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.”³ Although other papers such as Albrecht et al. (2007, 2010) and Duffie et al. (2007) have included seller heterogeneity in an asset market model, no paper that does so has generated a choosey buyer effect, which turns out to be important in explaining the disproportionate freezing up of the retail market.

To provide a more realistic treatment of the downturn, the basic model of the housing market is embedded in a richer model of mortgage default in which borrowers with negative equity may default on their mortgage or be locked into their current house despite a desire to move. This generates a new amplification channel: an initial shock that reduces prices puts some homeowners under water and triggers foreclosures, which cause more price declines and in turn further default. While reminiscent of the literature initiated by Kiyotaki and Moore (1997), the price declines here are caused by the general equilibrium effects of foreclosures. Lock-in of underwater homeowners also impacts market equilibrium by keeping potential buyers and sellers out of the market.

The richer model is used to quantitatively evaluate the extent to which foreclosures have exacerbated the ongoing housing bust. This quantitative analysis takes a two-pronged approach. First, we assess the strength of the amplification channel and its sensitivity to various parameters in the model. Second, we fit the model to data from the 100 largest MSAs to assess the empirical size of the amplification channel and test its implications across metropolitan areas. The model matches the data on the size of the price decline, the number of foreclosures, price declines in the retail market, and the REO share of sales. It also matches the heterogeneity in foreclosure discounts over the cycle found by Campbell et al. (2011). However, it falls short of explaining the full sales decline, suggesting that other forces have depressed transaction volume in the downturn. The quantitative analysis reveals that foreclosures exacerbate the aggregate price decline in the downturn by approximately 50 percent in the average MSA (or in other words account for a third of the decline) and exacerbate the price declines for retail sellers by over 30 percent.

Finally, we analyze the impact of the foreclosure crisis on welfare in our model and simulate three foreclosure-mitigating policies: slowing down foreclosures, refinancing mortgages at lower interest rates, and reducing principal. While we do not conduct a full normative analysis, the simulations of these policies highlight the trade-offs faced by policy makers.

² “Housing Market Slows as Buyers Get Picky” June 16, 2010.

³ “Buyer’s Market? Stressed Sellers Say Not So Fast” April 25, 2011.

The remainder of the paper is structured as follows. Before presenting the model, section 2 presents facts about the bust across metropolitan areas. To explain the data, the remainder of the paper develops a model of how foreclosures affect the housing market, first focusing on mechanisms and then on magnitudes. Section 3 introduces a model of exogenous defaults, and section 4 explores the intuitions and qualitative implications of the model. In section 5, the basic model is embedded in a more complete model in which negative equity is a necessary condition for default, which creates a new amplification mechanism in the form of a price-default spiral. The paper then turns to the magnitudes of the effects identified in sections 3-5. Section 6 calibrates the model and quantitatively analyzes the model's comparative statics and the strength of the price-default amplification channel. Section 7 takes the model to the national and cross-MSA data from the ongoing downturn. Section 8 considers welfare and foreclosure policy, and section 9 concludes.

2 Empirical Facts

The national aggregate time series of price, volume, foreclosure, and REO share presented in Figure 1 mask substantial heterogeneity across metropolitan areas in the severity of the housing bust and wave of foreclosures. To illustrate this, Figure 2 shows price and volume time series for four of the hardest-hit metropolitan areas. In Las Vegas, for instance, prices fell 61.5 percent, retail sales fell 84.0 percent, and the REO share was as high as 76.4 percent. Figure 2 also illustrates how foreclosure sales substitute for retail sales: retail sales rise as REO sales recede and fall as REO sales surge.

To provide more systematic facts about the heterogeneity of the bust across MSAs, we use a proprietary data set provided to us by CoreLogic supplemented by data from the United States Census. CoreLogic provides monthly data for 2000-2011 for the nation as a whole and the 100 largest MSAs, from which we drop 3 MSAs because the full data are not available for these locations at the start of the crisis. The data set includes a house price index,⁴ a house price index for retail sales only,⁵ the number of completed foreclosure auctions, sales counts for REOs, new houses, existing houses (including short sales), and the estimates of quantiles of the LTV distribution described previously. These statistics are compiled by CoreLogic using public records. CoreLogic's data covers over 85 percent of transactions nationally. We seasonally adjust the CoreLogic data and smooth the sales count series using a moving average. A complete description of the data is in appendix A.4.

By far the best predictor of the size of the bust was 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 through 2011. There is a clear downward pattern, with the notable exception of a few outliers in the lower-left of the diagrams which correspond to metropolitan areas in southern Michigan which experienced a substantial bust without a large boom.

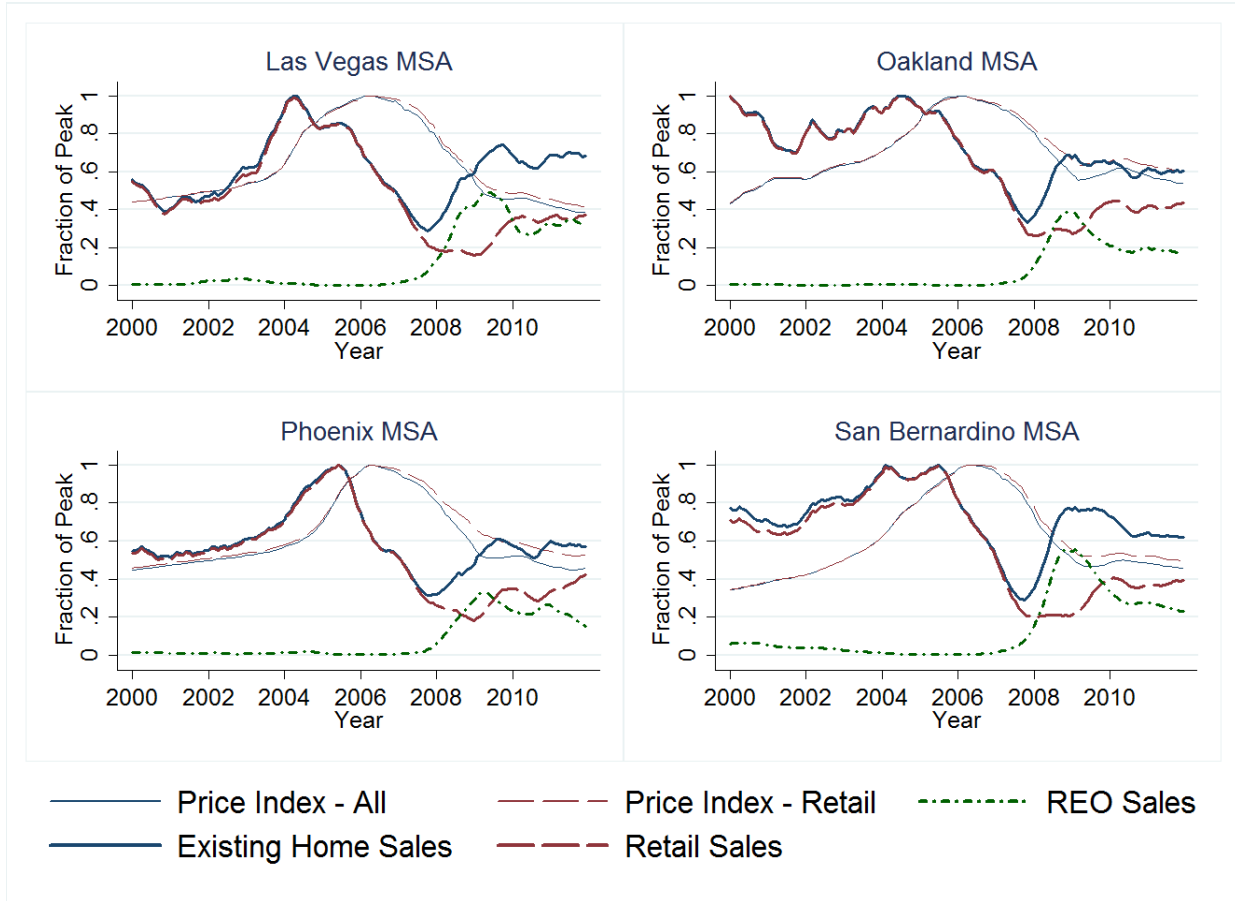
Figure 3 also reveals a more subtle fact in the data: places that had a larger boom had a more-than-proportionally larger bust. While a linear relationship between log boom size and log bust size has an r-squared of .44, adding a quadratic term that allows for larger busts in places with larger booms as illustrated in Figure 3 increases the r-squared to .57.

This paper argues that by exacerbating the downturn in the hardest-hit places, foreclosures can explain

⁴The CoreLogic price index is a widely-used repeat sales index that has behaved similarly to other cited indices in that it fell by a third during the downturn. The S&P Case-Shiller index shows similar declines to the CoreLogic index. The FHFA expanded-data index, which includes FHFA data proprietary deeds data from other sources, fell 26.7 percent.

⁵Given the small number of distressed properties prior to the downturn, price indices for distressed properties are typically not estimated. 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.

Figure 2: Price and Transaction Volume in Selected MSAs With High Levels of Foreclosures



Notes: All data is seasonally adjusted CBSA-level data from CoreLogic as described in the Appendix A.4. The sales lines are smoothed using a moving average. In panel C, all sales counts are normalized by the total number of existing home sales at peak while each price index is normalized by its separate peak value.

much of why the why the relationship between log boom size and log bust size is not linear. This explanation implies an additional reduced-form cross-sectional test: because default is closely connected to negative equity, a larger bust should occur in locations with the *combination* of a large bubble and a large fraction of houses with high loan balances – 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 are color-coded by quartiles of share of houses in the MSA with over 80 percent LTV in 2006. While the highest measured LTVs came in places that did not have a bust – home values were not inflated in 2006, so the denominator was lowest in these locations – one can see that the majority of MSAs substantially below the quadratic trend line were in the upper end of the LTV distribution.

To investigate whether the interaction of high LTV and a big bust combined is correlated with a deep

Figure 3: Boom vs. Bust Across MSAs



Note: Scatter plot of seasonally adjusted data from CoreLogic along with quadratic regression line. The data is fully described in Appendix A.4. Each data point is an MSA and is color coded to indicate in which quartile the MSA falls when MSAs are sorted by the share of homes with over 80% LTV in 2006. There are two main take-aways. First, there is a non-log-linear relationship between the size of the boom and the size of the bust, as the regression r-squared is .44 for a linear model and .57 for a quadratic model. Second, the color-coding of the points provides suggestive evidence that it is the combination of high LTV and a large bubble that is associated with a disproportionately large bust.

downturn more formally, we estimate regressions of the form:

$$\begin{aligned}
 Y = & \beta_0 + \beta_1 \max \Delta_{03-06} \log(P) + \beta_2 [\Delta_{03-06} \log(P)]^2 & (1) \\
 & + \beta_3 (Z \max \text{Share LTV} > 80\%) + \beta_4 (\Delta_{03-06} \log(P) \times Z \text{ LTV} > 80\%) \\
 & + \beta_5 (Z \% \text{ Second Mortgage, 2006}) + \beta_6 (\Delta_{03-06} \log(P) \times Z \% \text{ Second}) \\
 & + \beta_7 (Z \text{ Saiz Land Unavailability}) + \beta_8 (Z \text{ Wharton Land Use Regulation}) + \varepsilon
 \end{aligned}$$

where Z represents a z-score and the outcome variable Y is either the maximum change in log price, the maximum change in log retail prices, the maximum change in log existing home sales, the maximum change in log retail sales, the maximum REO share, or the fraction of houses that experience a foreclosure. The key coefficient is β_4 . 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, except rather than using income shocks to measure volatility, we use the size of the preceding bubble as measured by 2003-2006 price growth. 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).

Table 1: MSA Summary Statistics

	Unweighted Mean	SD	Min	Max	N
Max $\Delta \log(P)$	-.3398355	.2292549	-.9895244	-.0286884	97
Max $\Delta \log(P_{\text{Retail}})$	-.2784659	.1929688	-.9229212	-.0388126	97
Max $\Delta \log(\text{Sales}_{\text{Existing}})$	-.9354168	.2684873	-2.53161	-.4671416	97
Max $\Delta \log(\text{Sales}_{\text{Retail}})$	-1.174493	.3181327	-2.736871	-.5613699	97
$\frac{\text{Sales}_{\text{REO}}}{\text{Sales}_{\text{Existing}}}$.3147958	.1648681	.0834633	.795764	97
% Foreclosed	.0870826	.0719943	.0104154	.4205121	97
$\Delta \log(\text{Price})_{03-06}$.2974835	.179294	.0389295	.7288995	97
Share LTV > 80%	.1452959	.0756078	.025514	.3282766	97
Frac Second Mort, 06	.2026752	.0527425	.0259415	.2896224	97
Saiz Land Unav	.2779021	.2112399	.009317	.7964462	97
Wharton Land Reg	.2215807	.7050566	-1.239207	1.89206	97

Notes: Summary statistics for variables used in regression analysis. All data is from CoreLogic and fully described in Appendix A.4. Data is for 100 largest MSAs excluding three for which complete data are unavailable as described in the appendix.

Finally, to proxy for the housing supply elasticity we use a land unavailability index and the Wharton land use regulation index both from Saiz (2010). Table 1 shows summary statistics for our left hand side variables in the top panel and our right hand side variables in the bottom panel

The regression results are shown in Table 2. The first two columns show the impacts on price and retail price. While the additional variables do not explain all of the non-log-linearity, they have substantial predictive power. The key coefficient shows that the interaction between a large bubble and the share of homes with high LTV is correlated with large price declines, as suggested by Figure 3. These interactions also have a large effect on the REO share of sales and fraction foreclosed, suggesting that foreclosures have something to do with these trends. The coefficient on land regulation is also negative yet small, reflecting the amplification provided by a high housing supply elasticity. Having many houses with a second mortgage also reduces prices.

For sales, the regression has noticeably less predictive power and the dominant term is the constant. As discussed in the analysis of the national calibration above, this suggests that foreclosures combined with the size of the bubble will do a much worse job explaining the volume decline than the price decline, something that will be borne out in our cross-MSA simulations. The interaction between LTV and the bubble is insignificant for existing sales but significant and negative for retail sales. The pattern of REO volume largely replacing retail volume is consistent with the four markets with high levels of foreclosure in Figure 2.

3 Housing Market Model

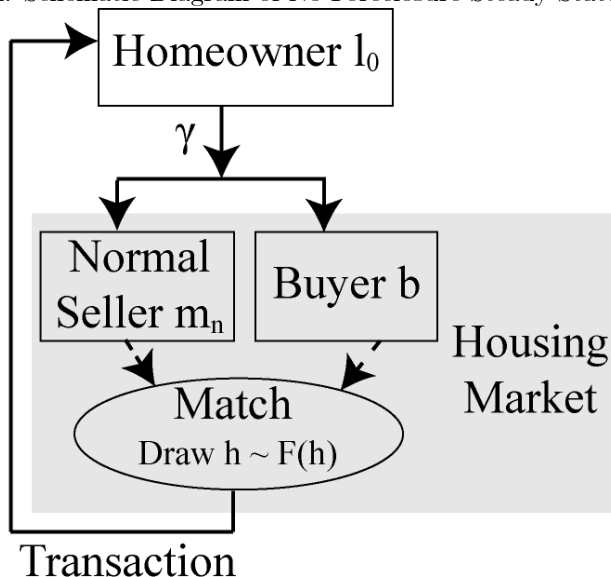
To theoretically examine the effect of forelclosures, we develop a model of the housing market in which foreclosures are exogenous. We subsequently embed this model in a framework in which default is modeled more realistically. Consequently, in this section, we focus on the mechanisms and qualitative predictions and defer a quantitative analysis of the model to section 7.

Table 2: Cross MSA Regressions on the Impact of the Size of the Bubble and Its Interaction With High LTV

Dependent Variable:	$\Delta \log(P)$	$\Delta \log(P_{\text{Retail}})$	$\Delta \log(\text{Sales}_{\text{Existing}})$	$\Delta \log(\text{Sales}_{\text{Retail}})$	$\frac{\text{Sales}_{\text{F.O.}}}{\text{Sales}_{\text{Existing}}}$	% Foreclosed
$\Delta \log(\text{Price})_{03-06}$	1.501 (0.656)**	0.884 (0.444)**	-1.932 (0.450)***	-0.493 (0.797)	-1.662 (0.477)***	-0.615 (0.189)***
$\Delta \log(\text{Price})_{03-06}^2$	-3.369 (0.785)***	-2.440 (0.529)***	1.431 (0.573)**	-1.046 (0.936)	3.016 (0.561)***	1.267 (0.229)***
Z Share LTV > 80%	0.062 (0.040)	0.064 (0.030)**	-0.008 (0.038)	0.005 (0.054)	-0.009 (0.033)	-0.008 (0.015)
$\Delta \log(P) \times Z \text{ LTV} > 80\%$	-0.314 (0.155)**	-0.314 (0.135)**	-0.146 (0.138)	-0.379 (0.182)**	0.218 (0.107)**	0.197 (0.067)***
Z Frac Second Mort, 06	-0.058 (0.033)*	-0.037 (0.023)	-0.066 (0.033)**	-0.071 (0.044)	-0.022 (0.027)	-0.019 (0.010)*
$\Delta \log(P) \times Z \text{ Second}$	0.099 (0.112)	0.036 (0.087)	0.102 (0.096)	0.049 (0.121)	0.245 (0.088)***	0.080 (0.036)**
Z Saiz Land Unav	-0.021 (0.018)	-0.008 (0.015)	-0.011 (0.018)	-0.014 (0.021)	0.003 (0.015)	-0.003 (0.006)
Z Wharton Land Reg	-0.031 (0.014)**	-0.027 (0.011)**	0.017 (0.017)	-0.001 (0.019)	0.019 (0.012)	0.008 (0.005)*
Constant	-0.417 (0.107)***	-0.282 (0.073)***	-0.551 (0.064)***	-0.944 (0.132)***	0.461 (0.081)***	0.136 (0.031)***
r^2	0.644	0.703	0.298	0.353	0.521	0.606
N	97	97	97	97	97	97

Notes: * = 10% Significance, ** = 5% Significance, *** = 1% significance. All standard errors are robust to heteroskedasticity. All data is from CoreLogic and fully described in Appendix A.4. Data is for 100 largest MSAs excluding 3 for which complete data are unavailable as described in appendix A.4. The key take-away is the coefficient on the interaction between bubble size and LTV, which shows that these two factors together were associated with a large price decline in the bust.

Figure 4: Schematic Diagram of No Foreclosure Steady State of Model



3.1 Setup

We consider a Diamond-Mortensen-Pissarides-style general equilibrium search model of the housing market. 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. Additionally, the outside options of market participants are crucial in search models, so a search framework is well-suited to formalizing the choosey buyer effect described in the introduction.

Time is discrete and the discount factor is β . There are a unit mass of individuals and a unit mass of houses, both fixed. This is a good approximation of the the downturn, in which there has been a very low level of new construction and decreased migration.⁶

The setup of the model's steady state is illustrated schematically in Figure 4. Table 3 defines the model's key variables. To simplify the analysis, we assume no default in steady state, which is approximately the case when prices are stable or rising.⁷ In steady state, mass l_0 of individuals are homeowners. Homeowners randomly experience shocks with probability γ that induce them to leave their house as in Krainer (2001) and Ngai and Tenreyro (2010).⁸ We assume that these shocks occur at a constant rate and that only individuals who receive a moving shock search for houses. This assumption turns off the amplification channel identified by Novy-Marx (2009) through which endogenous entry and exit decisions by market participants in a search model create a feedback loop which magnifies the effects of fundamental shocks.

An individual who receives a moving shock enters the housing market as both a buyer with flow utility u_b and a normal seller with holding cost m_n . Because shocks create both a buyer and a seller, the model is a closed system with a fixed population.⁹ In Section 7 we compare our model's predictions to data from

⁶We do not consider the impact of long-run changes in the homeownership rate and retirement rate on the long-run equilibrium of the market, nor do we consider the long-run impact of new construction, both of which may be affected by the downturn and are important subjects for future research.

⁷Allowing default in steady state complicates the analysis but does not substantially change the results.

⁸Moving shocks are a reduced form for a number of different life events that trigger a change in housing preference, such as the birth of children, death, job changes, and liquidity shocks.

⁹We use a closed system so that housing prices are not determined principally by the flow rates of buyers into and sellers out of the market but rather by the incentives of buyers and sellers in the market. Most moves are within-MSA (Sainai and

Table 3: Variables in Housing Market Model

Variable	Description
Endogenous Variables	
h	Stochastic Match Quality $\sim F(h)$
h_n, h_d	Cutoff h for normal, REO sellers
$S_{m,h}^B, S_{m,h}^S$	Surplus of type m seller with match quality h for buyer, seller
$p_{m,h}$	Price for type m seller with match quality h
μ	Market tightness (buyers/sellers)
$q_s(\mu), q_b(\mu)$	Prob seller meets buyer, buyer meets seller
r_m, r_d	Ratio of normal, REO sellers to total sellers
l_0, l_1	Masses of homeowners, homeowners that could foreclose
v_b, v_n, v_d, v_r	Masses of buyers, normal sellers, REO sellers, renters
Value Functions	
V_h	Value of owning home with match quality h
V_n, V_d	Value of seller for normal, REO sellers
B	Value of buyer
R	Value of renter
Parameters	
β	Discount factor
γ	Probability of moving shock
α	Probability moving shock causes foreclosure
σ	Probability of leaving renting
θ	Seller's Nash bargaining weight
χ	Probability of match in period (C-D matching function)
ξ	Exponent in C-D matching function
λ	Parameter of exponential distribution for $F(h)$
a	Shifter on exponential distribution for $F(h)$
u_b, u_r	Flow utility of being a buyer, renter
m_n, m_d	Flow utility of being a seller for normal, REO

both national and local markets, although the model as literally interpreted applies best to an metropolitan area.

As in Ngai and Tenreyro (2010), we assume that the buyer and seller act as independent agents. This means that there is no interaction between the buyer's problem or bargaining game and the seller's, and there is no structure placed on whether an individual buys or sells first. This assumption is not innocuous, as whether homeowners have sufficient liquidity to buy first may be important for market equilibrium and may affect bargaining as individuals who buy before selling holding two homes (Anenberg and Bayer, 2012). However, these effects are likely to be small relative to REOs. For instance, Springer (1996) examines several measures of seller motivation and finds that only REO sellers are distinguishable from normal sellers, and Anenberg and Bayer find that individuals who buy first sell their homes at a two percent discount, a figure that is swamped by the average REO discount.

Buyers and sellers in the housing market are matched each period. Matching is entirely random and search intensity is fixed, allowing us to focus on the effects of distressed sales rather than the search mechanism. When matched, the buyer draws a flow utility h from a distribution $F(h)$. Utility is linear and house valuations are purely idiosyncratic so that the transaction decision leads to a cutoff rule. These valuations are completely public, and prices are determined by generalized Nash bargaining. Because buyers know whether the seller is an individual or a bank in practice, symmetric information is reasonable. If the buyer (Souleles, 2009) or to MSAs with highly correlated housing prices, so the assumption of a closed system is reasonable.

and seller decide to transact, the seller leaves the market and the buyer becomes a homeowner in l_0 deriving flow utility h from the home until they receive a moving shock. If not, the buyer and seller each return to the market to be matched next period. Note that for simplicity we do not allow speculators or “flippers,” who would presumably sell quickly.

We introduce foreclosures into this basic steady state setup by adding two key ingredients. First, REO sellers have a higher holding costs, which is the case for several reasons. Mortgage servicers, who execute the foreclosure and REO sale, have substantial balance sheet concerns. In most cases, they 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 (Theologides, 2010). Furthermore, even though they are paid additional fees to compensate for the costs of foreclosure and are repaid when the foreclosed property sells, the servicer’s effective return is likely far lower than its opportunity cost of capital. Additionally, owner-occupants have much lower costs of maintenance and security. Finally, REO sellers usually leave a property vacant and thus forgo rental income or flow utility from the property.

An implicit assumption is that no deep-pocketed and patient market maker buys from distressed sellers and holds the property until a suitable buyer is found. While investors or “flippers” have bought some foreclosures, most have been sold by realtors to homeowners. This is likely due to agency problems and high transactions costs.

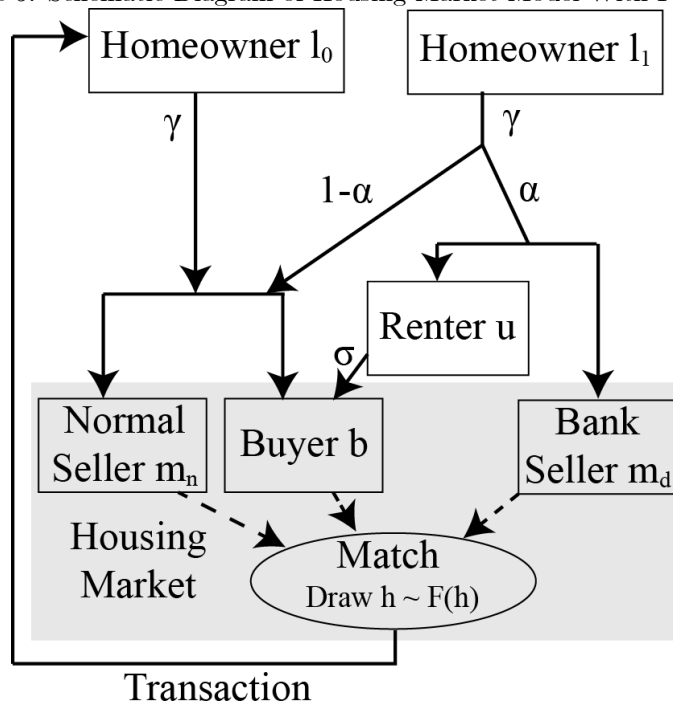
Second, individuals who experience a foreclosure are locked out of the market. This reflects the fact that a foreclosure dramatically reduces a borrower’s credit score. Indeed, many banks, the GSEs, and the FHA will not lend to someone who recently defaulted. Instead, foreclosed individuals become renters. This is supported by the data: Malloy and Shan (2011) 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. For simplicity, we assume that the rental market is segmented, and renters flow back into buying at an exogenous and fixed rate. While segmentation is a somewhat extreme assumption in the long run, it is a more reasonable approximation for the short-run effects in which we are primarily interested as conversions from owner occupied to rental units are costly and slow. Introducing an endogenous rental price and making the outflow rate covary with the price would create a force that mitigates some of the effects in the model.

Because flow utilities for foreclosures are drawn from the same distribution as for non-foreclosures, we are implicitly assuming that foreclosures are of roughly equal quality, which is likely not the case in practice (Gerardi et al., 2012). In our calibration, we are careful to use moments from studies that carefully control for quality.

While these are the only two new assumptions we make, foreclosures may have other effects. They may cause negative externalities on neighboring properties due to physical damage, the presence of a vacant home, or crime. Campbell et al. (2011) show that such effects are small and highly localized, although contagion is certainly possible in neighborhoods with high densities of foreclosures. There may also be buyer heterogeneity with respect to their willingness to purchase a foreclosure, generating an additional channel through which the REO discount widens as non-natural buyers purchase foreclosures. Finally, foreclosures may cause banks to limit credit supply, as shown theoretically by Chatterjee and Eyiungor (2011).

The two critical assumptions are introduced into the model in Figure 5. To simplify the analysis,

Figure 5: Schematic Diagram of Housing Market Model With Foreclosure



we assume away re-default.¹⁰ Instead, we consider a mass of potential defaulters and analyze how these potential defaulters flow through the system. One can think of these potential defaulters as homeowners with high mortgage balances as will be the case in section 5. These individuals have mass l_1 , and at time $t = 0$, when we introduce the exogenous foreclosure shock, we move everyone in l_0 to l_1 . Potential defaulters in l_1 also receive moving shocks with probability γ , but if they receive a moving shock it triggers a foreclosure with probability $\alpha(t)$ and is a normal moving shock with probability $1 - \alpha(t)$. If it is a normal moving shock, the homeowners becomes a buyer and a seller as in steady state. A foreclosure shock, however, causes a bank or GSE with holding cost m_d to take possession of the house and enter the housing market and the homeowner to become a renter with flow utility u_r .¹¹ Renters become buyers each period with exogenous probability σ . Because there is no re-default, all buyers, including those who were formerly renters, are added to l_0 when they buy a house, so the model gradually returns to steady state.

Buyers and sellers of both types are matched in the housing market. Let $v_b(t)$, $v_r(t)$, $v_n(t)$, and $v_d(t)$ be the masses of buyers, renters, normal sellers, and REO sellers in the market at time t . Market tightness $\mu(t)$ is equal to the ratio of buyers to sellers:

$$\mu(t) = \frac{v_b(t)}{v_n(t) + v_d(t)}. \quad (2)$$

Unlike general equilibrium search models of the labor market in which market tightness is determined principally by a free entry condition for firms posting vacancies, here market tightness is determined by flows

¹⁰We assume away re-default to keep the model consistent with the extended model in section 5. While this assumption does slightly increase the speed of convergence back to steady state over the course of the crisis, it does not substantively alter the quantitative or qualitative results.

¹¹The bank must hold a foreclosure auction, but in the vast majority of cases the auction reserve is not met and the bank takes the house as an REO. For instance, Campbell et al. (2011) report that 82 percent of foreclosures in Boston are sold as REOs rather than at auction. For simplicity we assume all houses become REO.

into renting due to default and out of renting at rate σ .

For the matching technology, we use a standard Cobb-Douglas matching function so that the number of matches when there are b buyers and s sellers is $\chi b^\xi s^{1-\xi}$. The probability a seller meets a buyer in a period with market tightness μ is given by $q_s(\mu) = \frac{\chi b^\xi s^{1-\xi}}{s} = \chi \mu^\xi$, and the probability a buyer meets a seller is $q_b(\mu) = \frac{\chi b^\xi s^{1-\xi}}{b} = \chi \mu^{\xi-1}$.

Let $V_h(t)$ be the value of being in a house with match quality h at time t , $V_m(t)$ be the value of being a seller of type m (either n or d) at time t , $B(t)$ be the value of being a buyer at time t , and $R(t)$ be the value of being a renter at time t . $V_h(t)$ is equal to the flow payoff plus the discounted expected continuation value:

$$V_h(t) = h + \beta \{ \gamma [V_n(t+1) + B(t+1)] + (1 - \gamma) V_h(t+1) \}. \quad (3)$$

The match surplus created when a buyer meets a seller of type $m = \{n, d\}$ and draws an idiosyncratic match quality of h at time t is a key value in the model. Denote this surplus by $S_{m,h}(t)$, the buyer's portion of the surplus by $S_{m,h}^B(t)$, and the seller's portion by $S_{m,h}^S(t)$. Let the price of the house sold if a transaction occurs be $p_{m,h}(t)$. The buyer's share of the surplus is equal to the value of being in the house minus the price and their outside option of staying in the market:

$$S_{m,h}^B(t) = V_h(t) - p_{m,h}(t) - u_b - \beta B(t+1). \quad (4)$$

The seller's share of the surplus is equal to the price minus their outside option of staying in the market:

$$S_{m,h}^S(t) = p_{m,h}(t) - m - \beta V_m(t+1). \quad (5)$$

Prices are set by generalized Nash bargaining with weight θ for the seller, so:

$$\frac{S_{m,h}^S(t)}{S_{m,h}^B(t)} = \frac{\theta}{1 - \theta} \quad \forall m. \quad (6)$$

Buyers and type m sellers will transact if the idiosyncratic match quality h is above a threshold value, corresponding to zero total surplus and denoted by $h_m(t)$. Because total surplus is:

$$S_{m,h}(t) = V_h(t) - (m + u_b) - (\beta B(t+1) + \beta V_m(t+1)) \quad (7)$$

the cutoff is implicitly defined by:

$$V_{h_m}(t) = m + u_b + \beta (B(t+1) + V_m(t+1)). \quad (8)$$

We can then define the remaining value functions. The value of being a type m seller 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 meet buyers with probability $q_s(\mu(t))$ and transactions occur with probability $1 - F(h_m(t))$, V_m is defined by:

$$V_m(t) = m + \beta V_m(t+1) + q_s(\mu(t)) (1 - F(h_m(t))) E[S_{m,h}^S(t) | h \geq h_m(t)]. \quad (9)$$

The most important aspect of V_m is that in a downturn $q_s(\mu)$ falls below its steady state value because foreclosures create renters rather than buyers ($\mu < 1$). The chance that a seller does not meet a buyer thus

reduces the value of being a seller.

The value of being a buyer is defined similarly, although we must account for the fact that the buyer can be matched with two types of sellers. Let the probability of matching with a type m seller conditional on a match be $r_m(t) = \frac{v_m(t)}{v_n(t)+v_d(t)}$. B is defined by:

$$B(t) = u_b + \beta B(t+1) + q_b(\mu(t)) \sum_m r_m(t) (1 - F(h_m(t))) E[S_{m,h}^B(t) | h \geq h_m(t)]. \quad (10)$$

Because of random matching, as more REO sellers enter the market the weight on REO sellers in the buyer's value function r_d rises. REO sellers are more likely to sell, so foreclosures raise the value of being a buyer. The decline in μ caused by foreclosures also raises $q_b(\mu)$, further increasing the value of being a buyer.

It is worth discussing what the implications of allowing buyers to direct their search towards foreclosures would be. A model with completely segmented REO and retail markets produces unreasonable parameter values. Intuitively, the REO and retail markets are linked by a buyer indifference condition that the probability of a match times the surplus must be the same in the REO market and the retail market. With a reasonable foreclosure discount, buyer indifference can only hold if the opportunity cost of waiting slightly longer for a distressed sale – the flow utility from being in that house – is implausibly high.

Furthermore, it is unlikely that any buyers look exclusively at one type of property. Instead, partially-directed search, in which buyers are able to direct their search to particular sub-markets in which the REO share of vacancies is higher than other sub-markets but is still not close to one, is most plausible. Examples of sub-markets include neighborhoods within a MSA or lower priced homes where there are likely to be more foreclosures. In this case, the effects we identify would be most pronounced in those sub-markets which had the highest REO share of vacancies, although there would be some spillovers because marginal individuals would switch to the REO-laden market. This is consistent with the findings of Landvoigt et al. (2012) that price declines in San Diego were stronger at the lower end of the market. We leave understanding the role of foreclosures for within-housing-market dynamics to future research.

The value of being a renter is defined as:

$$R(t) = u_r + \beta \{ \sigma B(t+1) + (1 - \sigma) R(t+1) \}. \quad (11)$$

We will assume $u_r = u_b$, so that a renter is simply a buyer without the option to buy.

The conditional expectation of the surplus given that a transaction occurs appears repeatedly in the value functions. This quantity can be simplified as in Ngai and Tenreyro (2010) by using (3) together with (7):

$$S_{m,h}(t) = V_h(t) - V_{h_m}(t) = \frac{h - h_m(t)}{1 - \beta(1 - \gamma)}.$$

The conditional expectation is

$$E[S_{m,h}(t) | h \geq h_m(t)] = \frac{E[h - h_m(t) | h \geq h_m(t)]}{1 - \beta(1 - \gamma)}. \quad (12)$$

We parameterize $F(\cdot) \sim \exp(\lambda) + a$, an exponential distribution with parameter λ shifted over by a constant a . The memoryless property of the exponential distribution implies that $E[S_{m,h}(t) | h \geq h_m^*(t)] = \frac{1}{\lambda}$. This is a fairly strong assumption. By using the exponential distribution in our simulations, we eliminate changes in the expected surplus due to changes in tail conditional expectations of the F distribution, which cannot be observed.

The model is completed with the laws of motion for the mass of sellers of type m , buyers, renters, and homeowners of type l_i . These laws of motion, which formalize Figure 5, are in appendix A.1.2.

Prices can be backed out by using Nash bargaining along with the definitions of the surpluses and (12) to get:

$$p_{m,h}(t) = \frac{\theta(h - h_m(t))}{1 - \beta(1 - \gamma)} + m + \beta V_m(t + 1) \quad (13)$$

This pricing equation is intuitive. The first term contains $h - h_m(t)$, which is a sufficient statistic for the surplus generated by the match as shown by Shimer and Werning (2007). As θ increases, more of the total surplus is appropriated to the seller in the form of a higher price. This must be normalized by $1 - \beta(1 - \gamma)$, the effective discount rate of a homeowner. The final two terms represent the value of being a seller next period, which is the seller’s outside option. 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 $h_m(t)$. Because with the exponential distribution $E[h - h_m(t)] = \frac{1}{\lambda}$, all movements in average prices work through $V_m(t + 1)$.

3.2 Numerical Methods

For reasonable parameter values, the model has a unique steady state that can be solved block recursively and studied analytically. The full derivation and existence and uniqueness proofs for the steady state can be found in appendix A.1.1. Although there are no foreclosures in steady state, the price and probability of sale for a REO seller are well defined and represent what would occur if a measure zero mass of normal sellers were instead REO sellers. For a fixed idiosyncratic valuation h , REO properties sell faster and at a discount due to the higher holding costs of distressed sellers.

The dynamics of the model, however, have no analytic solution, so we turn to numerical simulations. We solve the model using Newton’s method as described in appendix A.1.2.

Simulating the model requires choosing parameters. We defer a more rigorous quantitative analysis to section 7, which features a richer model, and focus on the mechanisms at work in this section. Consequently, for now we present simulation results using an illustrative calibration similar to the one described in section 5. We simulate a wave of foreclosures by moving everyone in l_0 to l_1 at time $t = 0$ and raising α for a period of five years. After the wave of foreclosures, the model returns to the original steady state.

4 Basic Model Results and Mechanisms

4.1 Market Tightness, Choosey Buyer, and Compositional Effects

The qualitative results in our model are caused by the interaction of three different effects: the “market tightness effect,” the “choosey buyer effect,” and the “compositional effect.” Each is crucial to understand the effect of foreclosures on the housing market.

First, because foreclosed individuals are locked out of the housing market as renters and only gradually flow back into being buyers, foreclosures reduce market tightness $\mu(t)$. This mechanically decreases the probability a seller meets a buyer in a given period and triggers endogenous responses as each party’s outside options to the transaction changes, altering the bargaining and the h s for which a sale occurs. For sellers, the reduction in market tightness reduces the value of being a seller for both types of seller, reducing prices and causing sellers to sell more frequently. The endogenous response is stronger for REO sellers who have a higher opportunity cost of not meeting a buyer. For buyers, the elevated probability of meeting a seller raises

their expected value, leading to lower prices and a shift in the cutoffs that makes buyers more choosy.¹² The market tightness effect elucidates that an important element of foreclosures is a reduction in demand relative to supply, as in a typical market a move creates a buyer and a seller while foreclosures create an immediate bank seller but a buyer only when the foreclosed upon individual's credit improves. This contrasts with some market analysts who treat foreclosures as a shift out in supply rather than a reduction in today's demand.

Second, the value of being a buyer rises because the buyer's outside option to transacting, which is walking away and resampling from the distribution of sellers next period, is improved by the prospect of finding an REO seller who will give a particularly good deal. Mathematically, as REOs make up a larger fraction of total vacancies, r_d rises and the term in the sum in (10) relating to REO sales gets a larger weight. This term is larger because REO sellers are more likely to transact both in and out of steady state. The resulting increase in buyers' outside options leads buyers to become more aggressive and demand a lower price from sellers in order to be willing to transact. In equilibrium, this leads to buyers walking away from more sales. Importantly, this effect will be most prevalent in the retail market where sellers are less desperate and therefore less willing to accommodate buyers' demand for lower prices, resulting in a freezing up of the retail market.

The choosy buyer effect is new to the literature. Albrecht et al. (2007, 2010) introduce motivated sellers into a search model, but focus on steady-state matching patterns (eg 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, and so while there are market tightness effects their model does not have a choosy buyer effect.

The market tightness effect and choosy buyer effect are mutually reinforcing. As discussed above, the market tightness effect is more pronounced for REO sellers. Because the value of being an REO seller falls by more, REO sellers become even more likely to sell relative to non-REO sellers during the downturn. This sweetens the prospect of being matched with an REO seller next period, amplifying the choosy buyer effect.

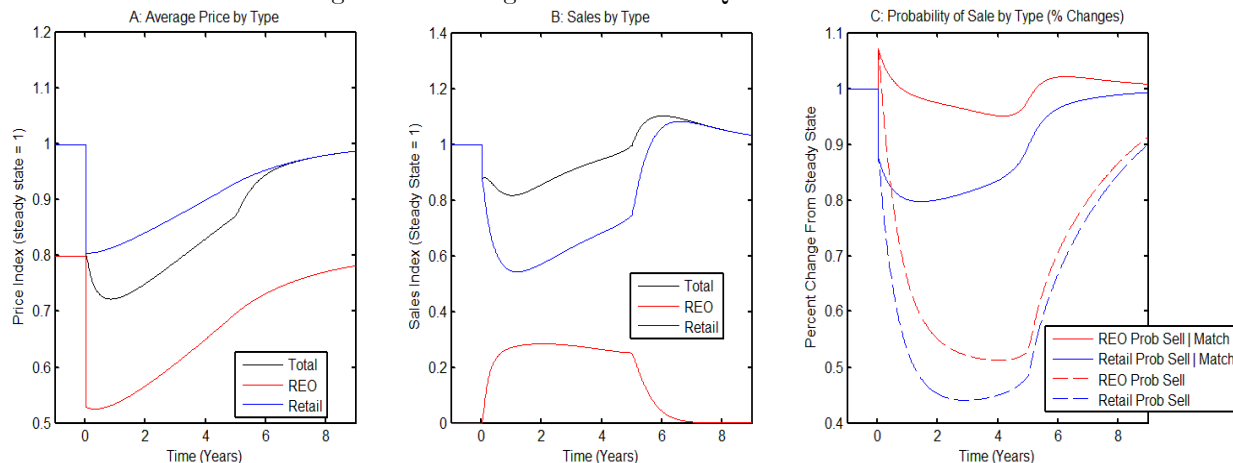
Finally, a greater share of REO sales makes the average sale look more like REO properties, which sell faster and at lower prices both in and out of steady state. Foreclosures thus cause a mechanical compositional effect that affects sales-weighted averages such as total sales and the aggregate price index.

The market tightness effect is the aspect of the model that comes closest to a standard Walrasian analysis with a single market for housing. By reducing the number of buyers relative to sellers, it is similar to an inward shift in the demand curve relative to the supply curve that reduces both prices and transaction volume. The market tightness effect does, however, asymmetrically impact REO and retail sellers due to their differential holding costs, leading to a greater freezing up of the retail market as buyers walk away from retail sellers in hopes of contacting increasingly-desperate REO sellers. These types of differential effects and further feedback loops – which stem from the choosy buyer effect and its interaction with the market tightness effect – are novel to the literature and differentiate our model from a simpler Walrasian model.

Furthermore, all three effects dissipate more slowly than in traditional asset pricing models because they depend on flows as well as stocks and lead to a sluggish return of the housing market to steady state. The choosy buyer and compositional effects last as long as foreclosures remain in the market, which is only a few months after the shock ends as these houses sell quickly. However, the market tightness effect persists

¹²In the calibration utilized here and in later sections, we set $\xi = .84$. Thus the effect on $q_s(\mu)$ significantly outweighs the effect of market tightness on $q_b(\mu)$.

Figure 6: Housing Market Model: Qualitative Results



Notes: This figure shows the results of the housing market model with exogenous foreclosures using an illustrative calibration similar to the one developed in Section 5 and a five-year foreclosure shock. Panels A and B show the average price and sales by type, with pre-downturn price and volume normalized to 1. Prices drop discretely at time zero as is standard in forward-looking models with no uncertainty. The REO discount widens, the aggregate price index is pulled towards the REO index as REOs make up a greater share of the market, and prices rise in anticipation of the end of the downturn. Retail volume plunges dramatically, but the decline is partially made up for by surging REO volume. Panel C shows the probability of sale conditional on a match and the unconditional probability of sale for each type with the pre-downturn probability normalized to one. This panel illustrates the mechanisms at work in the model, as described in the main text. The key take-away is that the probability of sale conditional on a match, which is the clearest indicator of how the behavioral responses of buyers and sellers play out in equilibrium, falls dramatically for retail and is roughly flat for REO.

for much longer as it takes several years for the renters to return to being homeowners.

4.2 Qualitative Results

Figure 6 shows the effect of a the five-year wave of foreclosures. Because the model is entirely forward looking, prices and probability of sale conditional on a match fall discretely on the impact of the shock at $t = 0$. This is typical in completely forward-looking models. The sluggish adjustment of house prices to shocks remains a puzzle for much of the literature, and a solution to this problem is outside the scope of this paper.

As shown in panel A, at $t = 0$ prices fall considerably for both REO and retail and gradually return to steady state over the next several years. The overall sales-weighted price index dips more than retail sales as foreclosures are averaged in. The price movements lead to a substantial rise in the average REO discount that falls off over time.

Prices fall due to all three effects. Recall that from (13), movements in the average price of properties sold by a type m seller are controlled by movements in $V_m(t + 1)$. The market tightness effect has a direct effect on the value of being a seller and thus brings down prices. Because this effect is stronger for REO sellers, this contributes to the larger REO discount. The choosey buyer effect has an indirect effect on $V_m(t + 1)$, as in general equilibrium increased buyer choosiness reduces the value of being a type m seller, which causes prices to fall. The effect of market tightness on the value of being a buyer operates in a similar manner. Finally, there is a pure compositional effect as REO sales become a greater share of total sales, which is shown graphically by the departure of the aggregate price index from the price index for retail sales.

As for sales, the wave of foreclosures sales causes the retail market to freeze up, with retail volume falling substantially as shown in panel B and REOs constituting a larger fraction of total sales than of total vacancies. Total volume, however, does not fall as much because much of the decline in retail sales is offset by REO sales. After the foreclosures end, sales return back to normal in a matter of months as REOs are eliminated from the market. Most of the sluggish adjustment comes from the dissipation of the accumulated renters and retail sellers, which takes several years.

The intuition behind the effects on transaction volume is more nuanced as the market tightness, choosey buyer, and composition effects have cross-cutting impacts. Panel C, which shows percent changes from steady state in the probability of sale both raw and conditional on a match, elucidates the role of each effect.¹³

Consider first the probability of sale conditional on a match, controlled by $h_m(t)$.¹⁴ The market tightness effect on the probability a seller meets a buyer raises the probability of sale conditional on a match because sellers meet buyers less frequently and thus have a greater incentive to sell when they are matched, an effect which is stronger for REO sellers. The choosey buyer effect and the effect of market tightness on the probability a buyer meets a seller both reduce the probability of sale conditional on a match as buyers become more choosy. Panel C shows that the two effects offset for REO sales as the probability of sale conditional on a match fluctuates around its steady state value, while the choosey buyer effect and the market tightness effect on buyers dominate for retail sales as the probability of sale conditional on a match falls substantially. The relative strength of these two effects for the two types of sellers thus plays an important role in freezing up the retail market.

The market tightness effect, however, plays an additional role: it mechanically reduces volume because there are fewer buyers. This causes the unconditional probability of sale and thus transaction volume to fall for both types, although it falls more for REO sellers. Note, however, that decline for retail sales is quicker and the trough lasts longer.

The compositional effect also plays an important role in determining transaction volume. Because REOs sell faster both in and out of steady state, as the average sale looks more like an REO, volume rises. This is the main reason why total volume does not fall so dramatically. It is possible for volume to rise, although for reasonable calibrations we find that the market tightness effect is strong enough relative to the compositional effect that REO sales do not make up the full shortfall in retail sales and overall volume falls.

Qualitatively, the model explains many salient features of the housing downturn. The substantial decline in both retail and REO prices is consistent with the data in Figure 1, and the widened distressed sale discount in a downturn is corroborated by Campbell et al. (2011). The freezing up of the retail market and the large share of REO sales in total sales relative to listings is borne out in the data, as are a rise in times to sale and increasing vacancy rates. The fact that REO sales replace a good deal of the lost volume in the retail market is consistent with the evidence from the hardest hit markets as shown in Figure 2.

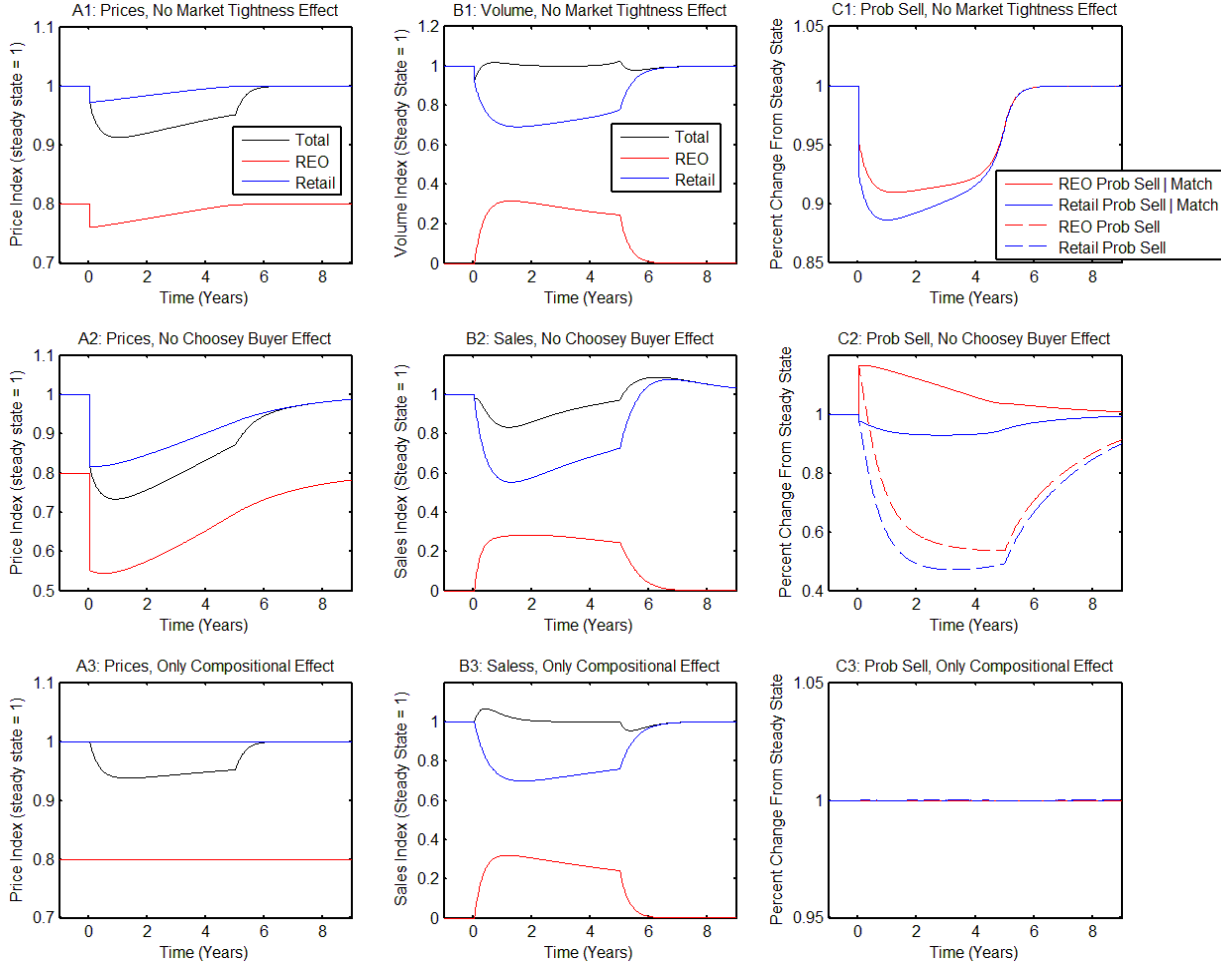
4.3 Isolating the Role of Each Effect

To further illustrate how each effect contributes to our results, Figure 7 depicts simulations identical to our main results for a wave of foreclosures except with the market tightness effect, choosey buyer effect, and both the choosey buyer and market tightness effects shut down. Although the market tightness effect plays

¹³The probability of sale conditional on a match is $\exp(-\lambda(h_m(t) - a))$ and the total probability of sale is $q_s(\mu(t)) \exp(-\lambda(h_m(t) - a))$

¹⁴Time to sale is inversely related to the unconditional probability of sale.

Figure 7: Isolating the Role of Each Effect



Notes: The top row shows price, sales, and conditional and unconditional probability of sale, all normalized to 1 for their pre-downturn values, for the case of no market tightness effect. The second row shows the same results for no choosey buyer effect. The third row shuts down both, leaving only the compositional effect. The figure shows that both the market tightness and outside option effects are critical for the qualitative results; in particular panel C2 shows that without the choosey buyer effect the probability of sale conditional on a match does not fall much for retail sellers. To shut down the market tightness effect, instead of creating renters we instead assume that distressed sale shocks create REO sellers and home-buyers. To shut down the choosey buyer effect, we modify the buyer's value function so that agents behave as if the probability they will hit a distressed seller is zero regardless of the presence of distressed sellers in the market. Because we calibrate to a steady state with no distressed sales, the steady state of these modified models replicates the steady state of our full model. See appendix A.2.1 for full details on these models.

an outsized role, all three effects are necessary for our results.

The market tightness effect generates a significant fraction of the price and volume declines. Row two also illustrates that the market tightness effect increases the conditional probability of sale for REO sellers during the downturn. Market tightness effects also cause total volume to decline because of the mechanical decrease in matching probabilities.

However, the choosey buyer effect plays an essential role in freezing up the retail market. As can be seen from row two, with no choosey buyer effect the conditional probability of sale for retail sellers essentially remains flat. On the other hand, from row one we can see that when only the choosey buyer effect is present there is a non-trivial decrease in this conditional probability. This freezing up is even more pronounced when both market tightness and choosey buyer effects are present due to their interaction.

The compositional effect mainly reduces the aggregate price index, as shown in row 3 of Figure 7. It also increases total volume slightly because REO sales sell faster.

5 An Extended Model of Default

Foreclosures are not random events. With few exceptions, negative equity is a necessary but not sufficient condition for foreclosure (Foote et al., 2008). This is because a homeowner with positive equity can sell his or her house, pay off the mortgage balance, and have cash left over without having to default. Homeowners with negative equity, however, are not able to pay the bank and thus default if they experience with a liquidity shock.

The previous section showed that foreclosures have general equilibrium effects that cause prices – and thus homeowner equity – to fall. In a world in which negative equity leads to foreclosure, this will cause more foreclosures and price declines, generating a feedback loop that amplifies the effects of an initial decline in house prices.¹⁵

In this section we embed the housing market component of the exogenous default model developed in the section 3 into a model in which negative equity is a necessary but not sufficient condition for default. Subsequent sections provide a rigorous quantitative analysis of the extended model and analyze welfare and foreclosure policy using the model.

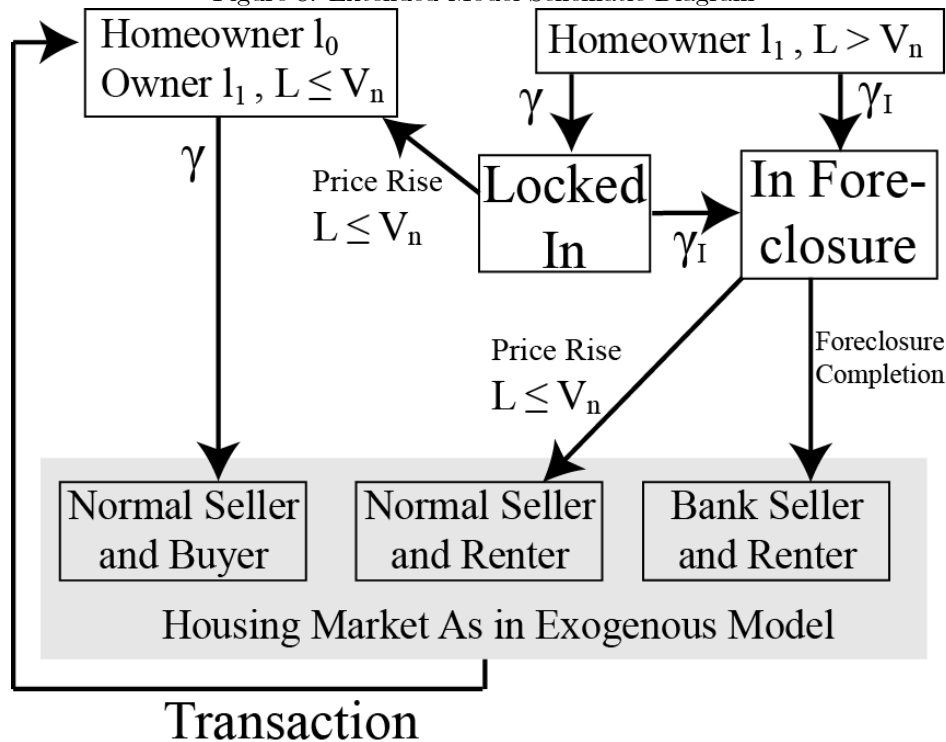
5.1 Default in the Extended Model

We model default as resulting primarily from shocks that cause homeowners with negative equity to be unable to afford their mortgage payments, the so-called “double trigger” model of mortgage default. While “ruthless” or “strategic default” by borrowers has occurred, much of the literature on default argues that strategic default has contributed surprisingly little to foreclosures, particularly at low levels of negative equity.¹⁶ Bhutta et al. (2011) use a method of controlling for income shocks to estimate that the median non-prime borrower does not strategically default until their equity falls to negative 67 percent. Even among non-prime borrowers in Arizona, California, Florida, and Nevada who purchased homes with 100 percent financing at the height of the bubble – 80 percent of whom defaulted within 3 years – over 80 percent of the defaults were caused by income shocks. Similarly, Foote et al. (2008) show that in the Massachusetts

¹⁵In all cases we have considered, each additional round of feedback is smaller than the previous one generating a convergent series and a unique dynamic equilibrium, although in principle the feedback could be strong enough to generate a divergent series.

¹⁶Relevant papers that analyze the default decision and conclude that a “ruthless exercise” option model of default is insufficient include Deng et al. (2000), Bajari et al. (2009), Elul et al. (2010), and Campbell and Cocco (2011).

Figure 8: Extended Model Schematic Diagram



housing downturn of the early 1990s, the vast majority of individuals who default have negative equity but most individuals with negative equity do not default. Consequently, the largest estimate of the share of defaults that are strategic is 15 to 20 percent.¹⁷ To keep the model tractable, we thus do not model strategic default, nor do we model the strategic decision of the bank to foreclose or short sales.¹⁸

Modeling negative equity requires that homeowners have loan balances. We assume that homeowners in l_1 have a distribution of loan balances L defined by a CDF $G(L)$.¹⁹ So that no foreclosures occur without an additional shock, in general we assume that $G(L)$ has continuous support on $[0, V_n^*]$, where the steady state value of being a normal seller which is equal to the expected price net of the costs of sale. We assume away re-default so that we do not need to worry how new home purchases affect $G(L)$.

To incorporate liquidity shocks into our model, we assume that they occur to individuals with negative equity at Poisson rate γ_I . All other shocks are taste shocks that occur at Poisson rate γ , so that liquidity shocks are in addition to normal shocks.

Liquidity and taste shocks have different effects depending on the equity position of the homeowner. Homeowners with any shock with $L \leq V_n(t)$ have positive equity enter the housing market as a buyer and seller. To keep the model tractable, we assume that buyers and sellers are identical once they pay off their loan balance. Homeowners with $L > V_n(t)$ have negative equity net of moving costs and default if they

¹⁷This estimate comes from Experian-Oliver Wyman. Guiso et al. (2009) analyze a survey that asks people whether they strategically defaulted and find that 26 percent of defaults are strategic.

¹⁸“Fishing” – that is listing a home for a high price and hoping that someone who overpays for it will come along as in Stein (1995) – and short sales are unusual because they require sellers to find a buyer who will pay a minimum price, which affects bargaining. Modeling short sales and their effect on market equilibrium is an important topic for future research.

¹⁹We are agnostic as to the source of the loan balance distribution and leave this unmodeled. $G(L)$ is fixed over time because principal is paid down slowly, particularly by those in the upper tail of the loan balance distribution who are relevant for the size of the feedback loop.

Table 4: Variables Used In Extended Model

Variable	Description
Endogenous Variables	
L	Loan Balance $\sim G(L)$
w	Mass of locked in homeowners
f	Mass of homeowners in foreclosure process
Parameters	
γ_I	Probability of a liquidity shock
δ	Probability above-water homeowner becomes renter
b_a, b_b	Parameters of Beta distribution for $G(L)$

experience an income shock because they cannot pay their mortgage or sell their house. Defaulters enter the foreclosure process.²⁰ Although foreclosure is not immediate and some loans in the foreclosure process do “cure” before they are foreclosed upon, for simplicity we assume that foreclosure occurs immediately. We alter this and introduce foreclosure backlogs in Section 8. We also assume that income shocks are a surprise, so an underwater homeowner expecting an income shock cannot list their house with the hope of getting a high-enough price that they can pay off their loan before the bank forecloses. While this may happen infrequently, it is unlikely very unlikely that a desperate seller would receive such a high price.

Finally, homeowners with negative equity who receive a taste shock would like to move but owe more than their house is worth. Consequently, they are “locked in” their current house until prices rise to the point that they have positive equity.²¹ We assume that once they do not move when they get a taste shock, these homeowners make accommodations and thus do not immediately move when they reach positive equity. Instead, once they have positive equity they become indistinguishable from households in l_0 who are waiting for a moving shock. In our context, lock-in further reduces sales in markets with many foreclosures.

Figure 8 shows our formalization of default in the extended model in a schematic diagram. The extended model only alters the mechanisms through which homeowners default and enter the housing market, and consequently the housing market component of the model is exactly as in section 3 and is thus not depicted in Figure 8. This structure preserves all of the key intuitions developed in section 4. Because there is still no default in steady state, the steady state remains the same. Because the modified model has an identical housing market to the model in section 3, the Bellman equations and cutoff conditions are unchanged. The new laws of motion are in Appendix A.1.4, and the additional parameters introduced in the extended model are listed in Table 4.

5.2 Starting the Downturn

An exogenous shock is required to generate an initial price drop. We introduce the exogenous shock in two different ways.

First, we assume that due to both tighter lending practices and income shocks a fraction $\delta(t)$ of individuals who sell their house as a normal seller after receiving a taste shock cannot buy a house and instead transition from owning to renting. This generates an initial market tightness effect that reduces prices, putting some individuals underwater and triggering a price-default spiral. Such a shock fits most naturally in the model.

²⁰It is also possible for banks to possess the house and rent it to the homeowner or to offer a short sale, in which the bank accepts a sale at a price below the outstanding loan balance. While these options have become more popular in recent years due in large part to political pressure, for most of the crisis the banks simply foreclosed on borrowers.

²¹Formally, define $w(t)$ as the mass of individuals who are locked in at time t . The distribution of loan balances in $w(t)$ will be the same as $G(L)$ truncated below at $V_n(t)$.

We use a 5-year increase in $\delta(t)$ to perform a sensitivity analysis in section 6.

However, when we take the model to the data in section 7, it is clear that the main shock in the recent episode is a bursting housing bubble. Our primary interest is understanding the amount of overshooting of prices caused by the presence of foreclosures and not realistically modeling the bursting of the bubble absent foreclosures. Consequently, we introduce a bursting bubble into the model with a permanent decline in a , the minimum flow utility of housing. While the flow utility of housing clearly did not fall overnight, the source of the initial drop in prices is immaterial to our results and reducing a is the simplest way to generate such a price drop. The reduction in a should thus be seen as a stand-in for a number of factors that have depressed home prices, from changes in credit availability to belief disagreements to irrational exuberance.²²

With a permanent decline in a and a constant hazard of an income shock, all individuals whose loan balance is above the new steady-state price level will eventually default. This does not seem realistic – many homeowners will eventually pay down their mortgage and avoid default. Rather than modeling the dynamic deformation of the $G(\cdot)$ distribution over time, we instead assume that after 5 years the hazard of income shocks γ_I gradually subsides over the course of a year.²³

With both exogenous shocks, defaults due to negative equity and the resulting foreclosures amplify the effects of shocks in the housing market due to a price-foreclosure feedback loop. Due to the forward-looking nature of agents in the model, this spiral is capitalized into the prices of retail sales and REO sales when the shock occurs, with further gradual declines as the REO share of volume in the market increases.

6 Quantitative Analysis: Calibration and Amplification

Having analyzed the mechanisms at work in our extended model, we now turn to the model’s implied magnitudes by conducting a quantitative analysis. In this section, we calibrate the model and examine the strength of the amplification channel. In section 7, we take the model to the data on the ongoing downturn.

6.1 Calibration

In order to simulate the model numerically, we must choose parameters. As mentioned previously, we parameterize the distribution of idiosyncratic valuations $F(\cdot)$ as an exponential distribution with parameter λ shifted by a . We parameterize the loan balance distribution $G(L)$ as a beta distribution with parameters b_a and b_b scaled to have support on $[0, V_n^*]$. This flexible distribution allows us to approximate the loan balance distribution in various locations on the eve of the crisis as described below. This gives 12 exogenous parameters to calibrate for the basic housing market model – a , γ , m_d , m_n , u_r , u_b , θ , β , σ , λ , χ , and ξ – and three parameters to calibrate unique to the extended model – γ_I , and b_a and b_b . We also must choose the initial shock.

Our calibration procedure proceeds in three steps. We take care to calibrate to pre-downturn moments whenever possible in order to make our tests out-of sample, although in some cases we have no choice but to choose a parameter using data from the housing bust. First, we set γ , u_r , u_b , β , χ , ξ , and γ_I independently to match several moments. Second, we choose a , m_n , m_d , θ and λ so that the steady state of the model matches additional targets. Third, we calibrate b_a and b_b to the appropriate geographic unit that we are considering.

²²While a bursting bubble is the most likely source of a large price change that would put many homeowners underwater, the type of foreclosure crisis we describe could be created by any type of large negative price shock.

²³Formally, after 5 years γ_I falls by 5% of its previous value every month, taking roughly a year to settle at zero.

This leaves two variables that we do not choose through calibration: σ , the probability of leaving the rental market, and the initial shock. Although there are several guidelines regarding how long banks deny mortgages to individuals who default,²⁴ there is no good data on this parameter in practice. Consequently, we pursue a two-pronged approach. First, to understand the impact of σ and the exogenous shock, we examine the response of the model to different values of each. Second, we use data on the size of the bust across housing markets to select a preferred calibration of these two parameters, as described in Section 7. In the remainder of this section, we describe the moments to which we calibrate the model in our three steps.

Step 1: Exogenous Parameter Choices

We choose u_r , u_b , β , and γ so that one period is equivalent to one week, although the results are insensitive to the time interval chosen. We set the annual discount rate equal to 5%, so that the discount factor is $\beta = 1 - \frac{.05}{52}$. We assume $u_r = u_b = 0$ so that buyers and renters are identical in their flow utility. Buyers, however, have the option to buy which has considerable value so $B > R$. This assumption is equivalent to assuming that buyers and renters pay their flow utility in rent in a perfectly competitive rental market.

We set the probability of moving houses in a given week to fit a median tenure for owner occupants of approximately 9 years from the American Housing Survey from 1997 to 2005, so $\gamma = \frac{.08}{52}$.

We set $\xi = .84$ using estimates from Genesove and Han (2012), who use National Association of Realtors surveys to estimate the contact elasticity for sellers with respect to the buyer-to-seller ratio. χ is then a constant chosen to make sure the probability of matching never goes above 1 for either side of the market. We choose $\chi = .5$, which in our simulations leads to matching probabilities on $[0, 1]$. The results are robust to alternate choices of χ .

The one parameter that we need data from the downturn to choose is γ_I . We set $\gamma_I = 8.6\%$ annually using national data from CoreLogic on the incidence of foreclosure starts for houses with negative equity as described in appendix A.4.

Step 2: Matching the Pre-Downturn Steady State

We then fit the following five aggregate moments from the housing market prior to the housing bust to the model's steady state to set θ , a , λ , m_n , and m_d :

1. The mean house price for a retail sale, which we set equal \$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. Our results are approximately invariant to the mean house price as long as the residual variance is rescaled proportionally.
2. The residual variance of house prices due to the idiosyncratic preferences of buyers. We set this equal to \$10,000.
3. The REO discount in terms of mean prices, which we set equal to a quality-adjusted 20% based on the average discount in good times from Campbell et al. (2011), whose results are consistent with a literature surveyed by Clauretje and Daneshvary (2009). In Section 6 we also consider a 10% discount, closer to the estimates of Zillow (2010) and Clauretje and Daneshvary (2009).
4. Time on the market for retail houses, which we set to 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

²⁴Three years of good credit are needed to get a Federal Housing Administration loan, and according to Fannie Mae guidelines issued in 2010, individuals are excluded from getting a mortgage for two to seven years if they are foreclosed upon, depending on the circumstances. However, these guidelines are not ubiquitous.

Table 5: Calibrated Parameter Values

a	γ	m_d	m_n	u_r	u_b	θ
3.399	$\frac{0.08}{52}$	-0.403	-0.091	0	0	0.087
β	λ	χ	ξ	γ_I	b_a	b_b
$1 - \frac{0.95}{52}$	3.490	0.5	0.84	$\frac{.086}{52}$	0.898	1.223

(2012) and Springer (1996), likely because of imperfect adjustment for withdrawn listings and re-listings. Our results are not sensitive to this number.

- Time on market for REO houses, which we set to 15 weeks. Springer (1996) analyzes various forms of “seller motivation” such as relocation and financial distress using data from Texas from 1991-3. He finds that a foreclosure sales are the only motivated sellers that have significantly reduced time on the market. His estimate is that time on the market is reduced by .2135 log points or 23.7%. However, REO sales are almost never withdrawn from the market, whereas retail sales are frequently withdrawn (Anenberg, 2012). We also attempt to adjust so our number excludes extremely rundown properties that sit on the market for several years.

These moments provide a unique mapping to θ , a , λ , m_n , and m_d , as described in appendix A.1.3.

The calibration procedure results in the parameter values listed in Table 5, with all prices and dollar amounts in thousands of dollars. Two main things are of note about the calibration. First, $m_d < m_n$, so the flow cost of being a REO seller is higher than the flow cost of being a regular seller. This is due to the fact that distressed sales take less time to sell and trade at a discount in steady state. Second, θ is quite low in order to rationalize the 20 percent discount for REO sales in steady state. This means the buyer will get a majority of the surplus and the value of being a buyer in the buyer’s market of the downturn will be high.²⁵

Step 3: Geographically-Specific Parameters

To calibrate the two parameters of the loan balance distribution b_a and b_b at the national and MSA level we use proprietary data from CoreLogic on seven quantiles of the combined loan-to-value distribution for active mortgages in 2006. These LTV estimates are compiled by CoreLogic using public records, with the LTV estimates supplemented using CoreLogic’s valuation models.²⁶ 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. We construct the empirical CDF of the loan balance distribution and then fit a beta distribution with parameters b_a and b_b to the empirical distribution using a minimum distance method described in Appendix A.4. The fit is quite good across the 50th to 100th percentiles of the LTV distribution.] Table 5 shows the resulting b_a and b_b for the nationwide loan balance distribution.

²⁵A low θ is consistent with the logic of directed search and Genesove and Han’s (2011) estimate of $\xi = .84$. In directed search models θ is determined endogenously through price posting behavior. The buyer’s bargaining power is the elasticity of contacts with respect to market tightness ξ and so θ should be below $1 - \xi = .16$. Intuitively the seller has very low bargaining power because the seller’s contact rate is more sensitive to a change in market tightness than the buyer’s.

²⁶Although the CoreLogic estimates of negative equity and the loan-to-value distribution are most cited, some have argued that they do not fully capture the extent of negative equity. Recent estimates by Zillow use credit report data instead of public records and get a higher figure for negative equity than CoreLogic (methodological differences are described here <http://blogs.wsj.com/developments/2012/05/24/negative-equity-more-widespread-than-previously-thought-report-says/>). Beyond issues of data sourcing, Kortewig and Sorensen (2012) argue that traditional methods of estimating house price indices under-estimate the variance of the house price distribution and thus under-estimate the number of loans with high LTV.

Table 6: Sensitivity Analysis: Time Out of Market For Renters

	Price Index Decline	Total Sales Decline
$\sigma = 1/52$ (1 year out)	4.0%	4.4%
$\sigma = 1/65$ (1.25 years out)	4.9%	4.8%
$\sigma = 1/78$ (1.5 years out)	5.8%	5.2%
$\sigma = 1/91$ (1.75 years out)	6.6%	5.5%

6.2 Strength of Amplification Channel and Comparative Statics

Having calibrated the model, we now gauge the potential magnitude of the amplification channel and elucidate key comparative statics in the extended model. Our initial shock will be one in which a fraction $\delta(t)$ of individuals who receive a taste shock transition from owning a home to renting, as described previously. This causes an initial price decline since it reduces the number of buyers but not the number of sellers. In particular, we use a shock of $\delta(t) = .10$ for five years.²⁷ Holding the initial shock constant, we then vary the average weeks out of the market for a renter, the steady state discount on REO sales, and the loan balance distribution. Specifically, we consider average times out of the market of 1, 1.25, 1.5, and 1.75 years and REO discounts of 10% and 20%. For the loan balance distribution, we consider beta distributions fitted to match the national data as well as data from New York, which had a low share of high LTV homes, and Las Vegas, which had a high share.

Since our initial shock is a market tightness effect, the strength of this effect will be governed by σ . Table 6 reports the price index decline and total volume decline generated by the initial shock in the absence of any defaults.

Increasing the average length of time for which individuals who transition to renting stay out of the housing market leads to greater percentage decreases in both the price index and total volume. The effects of such a shock as small as the one we consider are relatively modest, but our key question remains the potential strength of the amplification channel when we allow for defaults. Table 7 reports the results from varying the REO discount and loan balance distribution in addition to σ . Rather than reporting levels, each entry reports the percentage amplification of the aggregate price index decline and sales decline generated by defaults over and above the decline created by the initial shock in a price-volume pair.

Table 7 shows that foreclosure spirals can significantly amplify an initial shock. The strength of this spiral grows as we increase the amount of time individuals who default spend out of the housing market due to a more persistent market tightness effect. The shape of the loan balance distribution also plays a critical role in determining the strength of the amplification. A greater proportion of individuals with high LTV ratios implies that a given initial shock will put a greater fraction of the market underwater. This leads to greater numbers of foreclosures, more powerful general equilibrium effects, and in turn even more foreclosures. This illustrates the fragility created by the combination of a housing bubble and reduced down payment requirements.

Table 7 also shows that a lower steady state REO discount dampens the spiral. This is in part due to a compositional effect which ameliorates the effect of a given number of REOs on the price index. Additionally, though, the choosy buyer effect is weaker since waiting for an REO sale is no longer as attractive, and so the retail market does not freeze up as much. Finally, the amplification channel can generate significant *total* volume declines, greater than we saw in section 3 with exogenous defaults for a given shock size. The

²⁷Raising the size of the initial shock within reasonable parameter ranges magnifies the strength of the amplification channel, but the increase is mild.

Table 7: Sensitivity Analysis: Loan Balance Distribution and REO Discount

$\sigma = 1/52$	REO Disc. 10%	REO Disc. 20%	$\sigma = 1/65$	REO Disc. 10%	REO Disc. 20%
New York	(4.8%, 2.8%)	(6.2%, 3.4%)	New York	(6.6%, 4.0%)	(8.4%, 4.9%)
National	(22.3%, 11.9%)	(29.4%, 14.6%)	National	(30.0%, 17.3%)	(38.5%, 20.5%)
Las Vegas	(46.0%, 23.7%)	(60.2%, 28.7%)	Las Vegas	(62.0%, 34.6%)	(79.3%, 40.6%)
$\sigma = 1/78$	REO Disc. 10%	REO Disc. 20%	$\sigma = 1/91$	REO Disc. 10%	REO Disc. 20%
New York	(8.6%, 5.5%)	(10.8%, 6.5%)	New York	(10.9%, 7.2%)	(13.6%, 8.4%)
National	(39.3%, 23.9%)	(49.7%, 27.8%)	National	(51.0%, 32.2%)	(63.9%, 36.9%)
Las Vegas	(83.6%, 49.0%)	(105.6%, 56.4%)	Las Vegas	(115.4%, 69.5%)	(145.5%, 79.1%)

Note: This table shows comparative statics with respect to three important variables: σ , the Poisson probability of a renter becoming a buyer, the REO discount, and the loan balance distribution. The table shows how these variables affect the degree of amplification of the initial shock to prices created by adding foreclosures. The first entry in each pair is the percentage increase in the price index decline generated by defaults over and above that created by the initial shock. The second entry in each pair is the percentage increase in the total sales decline. For instance, when $\sigma = 1/52$ and the REO discount is 20%, the table has an entry of (60.2%, 28.7%) for the 2006 Las Vegas loan balance distribution. Given a price index decline of 4.0% and volume decline of 4.4% from the initial shock alone, this indicates that the full price index and volume declines are respectively 6.4% and 5.7%.

reason is that in the extended model, relative to section 3, the number of individuals who become locked-in during the downturn can be substantial. Note that this implies a greater REO share of vacancies which strengthens both the compositional effect and the choosy buyer effects and thus feeds back into further declines in both the overall and non-distressed price indexes.

7 Cross-MSA Quantitative Analysis

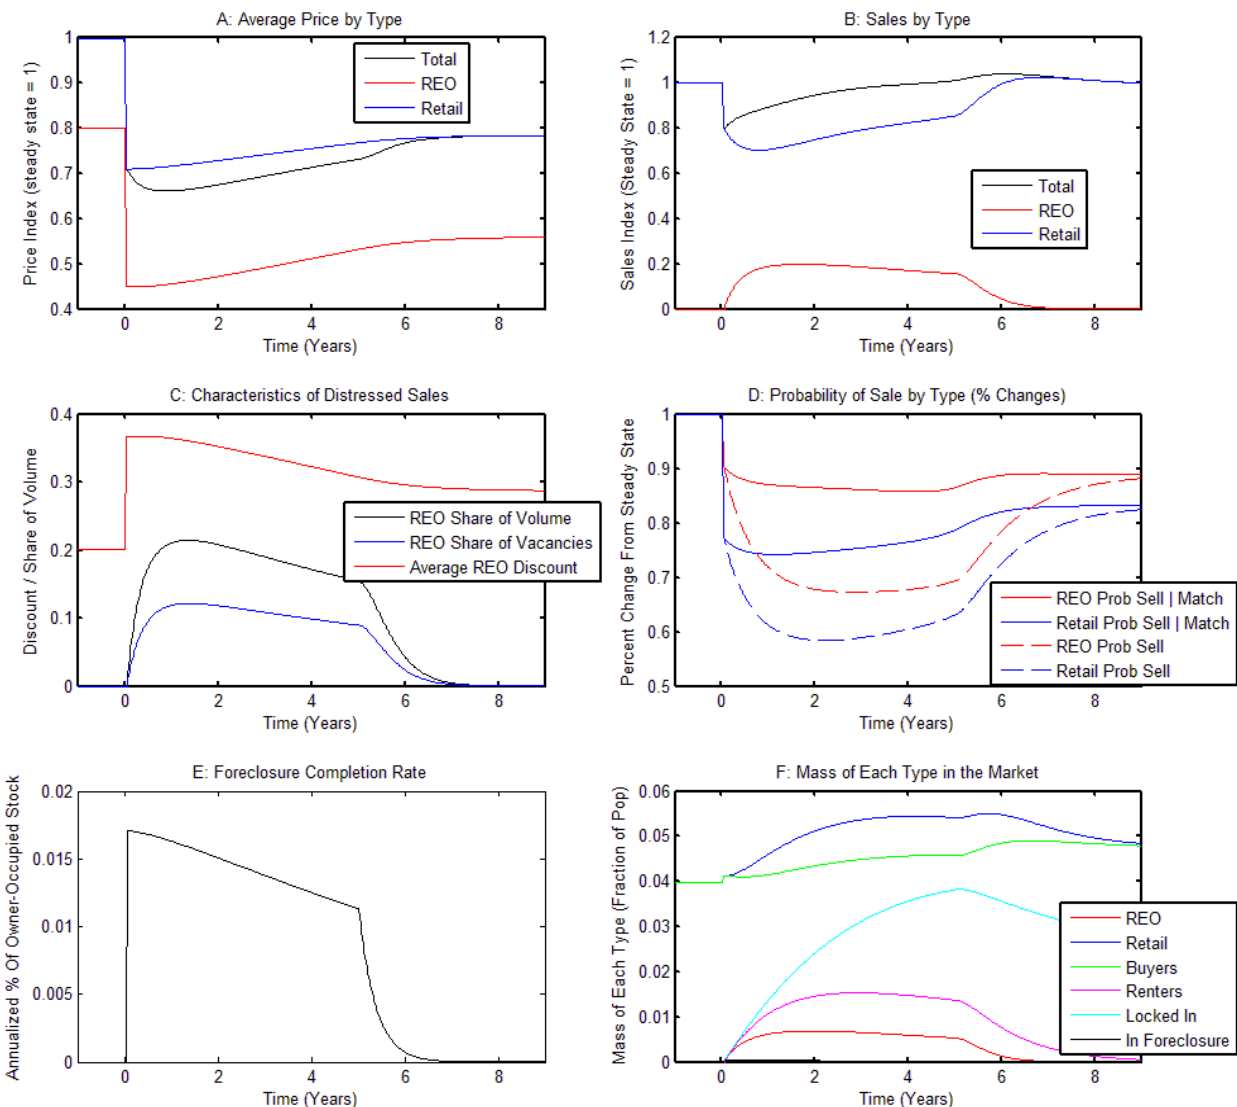
In order to assess quantitatively the role of foreclosures in the crisis and to test the model's performance, we calibrate the model to national and cross-MSA data described in section 2.

Because the size of the preceding bubble is the single best predictor of the size of the ensuing bust, we use a permanent shock to prices that operates through reducing flow utilities a to start the downturn. We assume income shocks last for 5 years, after which they gradually taper off. Recall that there are two parameters that are left to calibrate: σ , which controls how long the average renter takes to return to owner-occupancy, and the size of the permanent shock to prices. To calibrate these parameters, we use the aggregate national data and the cross-MSA data together. We first set a grid of σ s. For each σ , we choose the nationwide permanent shock to prices so that with the nationwide loan balance distribution the model exactly matches the maximum log change in the national house price index. We then simulate the model for each MSA and for each value of σ by assuming that the permanent shock to prices in the MSA is equal to the nationwide permanent shock to prices multiplied by the relative size of the bubble in the MSA as measured by $\frac{\Delta \log(P_{03-06}^{MSA})}{\Delta \log(P_{03-06}^{National})}$. In other words, we assume that the relative amount of housing price appreciation from the bubble that is permanently lost is the same in each MSA.²⁸ We then calculate the unweighted sum of squared distances between the maximum change in the log price index in the data and the model for each value of σ and choose the σ with the minimum sum of squared distances.

This methodology yields a sum of squared distances function that is a smooth and convex function of σ . Intuitively too high of a σ will cause the model to over-predict the size of price declines in high-LTV but

²⁸For instance, the maximum $\Delta \log(P)$ in Las Vegas is was 1.52 times as big as the nation-wide price index. Below we find the permanent price decline for the nation is 21.5%. Thus the permanent decline in Las Vegas is 32.8%.

Figure 9: National Calibration With Permanent Price Drop



Note: This figure shows the results of the extended model calibrated to match the national and MSA data as in section 7. It uses a permanent shock to housing values and an LTV distribution corresponding to the national housing market. Panels A and B show the average price and sales by type, with pre-downturn price and volume normalized to 1. Panel C shows the REO discount, share of vacancies, and share of volume. Panel D shows the probability of sale conditional on a match and the unconditional probability of sale for each type with the pre-downturn probability normalized to 1. Panel E shows the annualized fraction of the owner occupied housing stock that is foreclosed upon at each point in time. Panel F shows the mass of each type of agent in the market. The overall results are similar to the qualitative results outlined previously.

low-bubble MSAs, while too low of a σ will cause the model to under-predict the size of price declines in high-bubble MSAs. The optimal σ is the one that does best across distribution of bubble size.

The sum of squared distances has a unique minimum, which corresponds to $\frac{1}{\sigma} = 1.05$ years out of the market for the average renter and $\Delta_{\max} P^{National} = 21.5\%$ (a log point decrease of .242).²⁹ A 21.5% price drop from the peak implies that about 2/3 of the price gains between 2003 and 2006 were permanently lost when the bubble burst.

Figure 9 shows the time series price, sales, the characteristics of distressed sales, probability of sale, foreclosures, and the mass of each type in the market for the resulting national simulation. The qualitative patterns closely match those described in the simpler model in Section 4. Recall that there is a one-to-one mapping between the unconditional probability of sale and time to sale, which move in opposite directions.

Figure 10 shows the results of the cross-MSA simulations by plotting the simulated results against the data in six panels that show the maximum change in log price, log retail price, log sales, log retail sales, REO share, and fraction foreclosed. In each figure, the 45-degree line is drawn in to represent a perfect match between the model and the data. The small dots represent MSAs while the large X represents the national calibration.

The calibration procedure does well in matching declines in the aggregate price index across the bubble-size spectrum, as indicated by the fact that the data points are clustered around the 45-degree line. As with the regressions, the extreme outliers are in greater Detroit (the two points in the lower right), which has had a large bust without a preceding boom, and Stockton (the point in the lower middle), which had a much larger bust than boom. Despite these outliers, we cannot reject a coefficient of one when regressing the simulated results on the data, and relative to a case with no default where the entire national price decline is permanent, adding default to the model increases the r-squared of the simulation by 20 percent.³⁰ These results suggest that default can explain some of the nonlinearity in Figure 3.

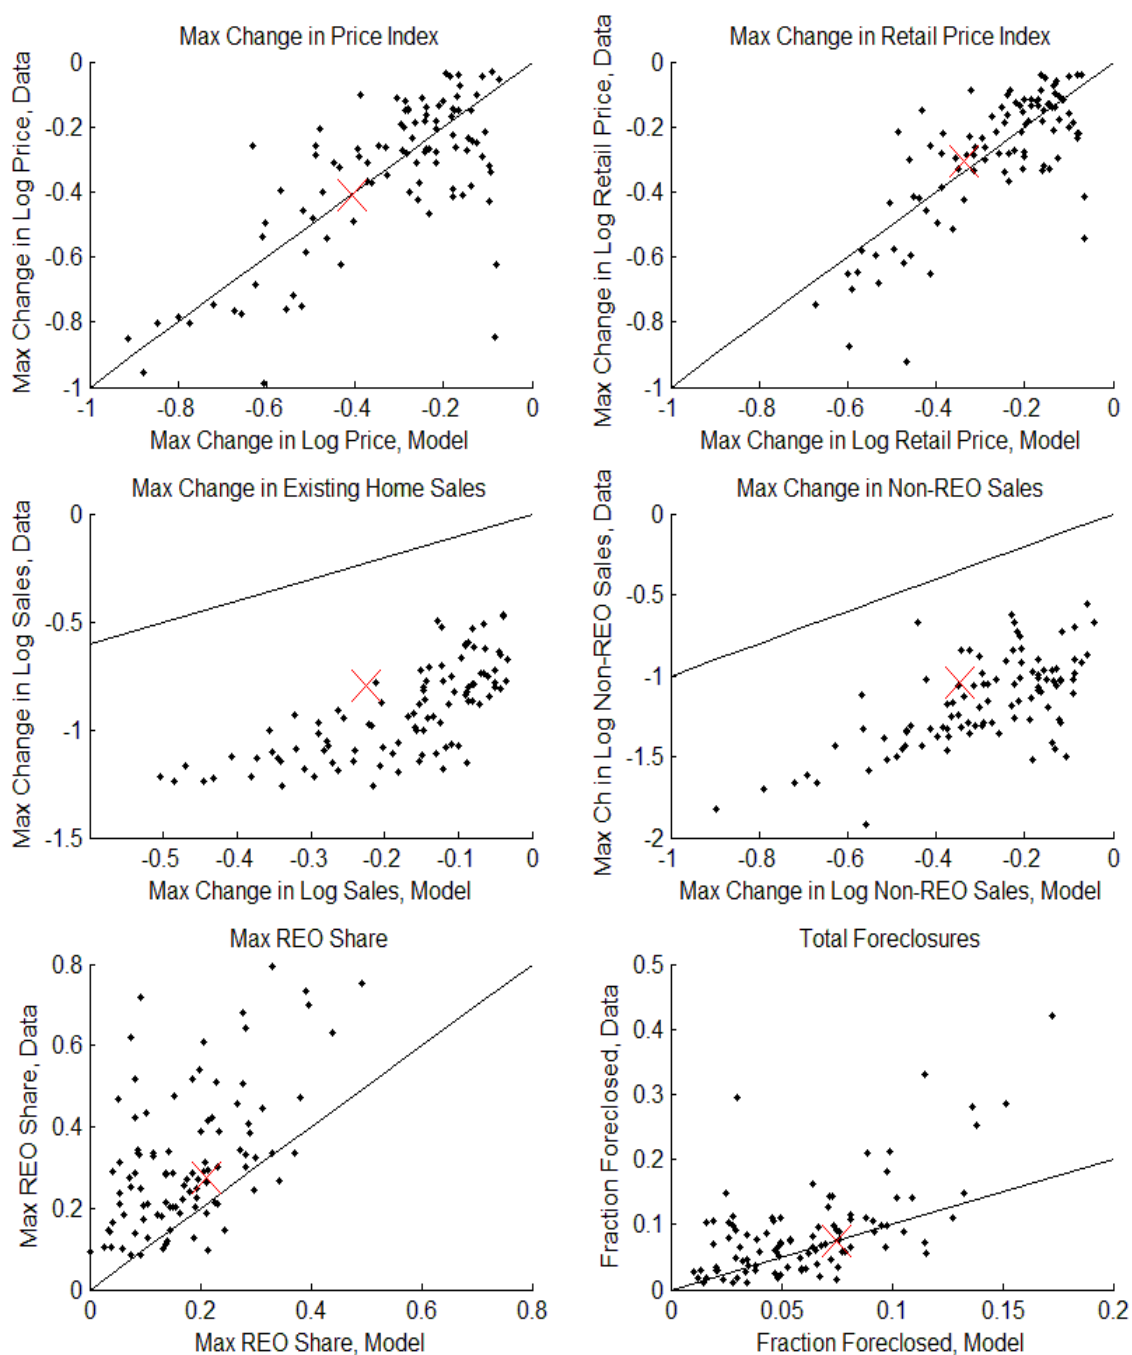
More importantly, the model cannot be rejected for additional outcomes beyond the aggregate price index that was used for calibration. For the change in log retail price, the national data closely matches the model: for the most part, the points lie around the 45-degree line, with a few more exceptions where the model under-estimates the change in log retail price such as Stockton and Las Vegas. Nonetheless, a statistical test confirms that we cannot reject that the model matches the data. The model also comes close to matching the data for total foreclosures over 5 years and the model cannot be statistically rejected, although there are some extreme outliers where there were many more foreclosures than the model predicted. Again, some of these are in Greater Detroit and Stockton, but there are a few other hard-hit markets like Las Vegas and the Central Valley in California where the model under-predicts the number of foreclosures. Finally, although it is not in Figure 10, the national calibration predicts a maximum REO discount of 36.7%. This is slightly above the maximum foreclosure discount for Boston of 35.4% reported by Campbell et al. (2011), so the model can explain time variation in REO discounts.

However, as foreshadowed by the regressions, the model consistently under-predicts the decline in sales. In Figure 10, the data cluster roughly parallel to the 45-degree line for both retail and total sales, although this is only statistically significant for retail sales. This means that the model does a good job of capturing differences in the size of the maximum sales decline across locations but that volume has fallen nationwide for reasons beyond the model. Potential unmodeled forces reducing volume include the tightening of credit

²⁹ 1.05 years out of the market for the average renter may be a bit on the short side, but because we have a constant Poisson probability of leaving renting, the distribution of times out of the market has a very thick tail.

³⁰ This r-squared is calculated as 1 minus the squared distance between the simulated maximum log price decline and its counterpart in the data divided by the squared distance between the data and its mean. This is the r-squared of a regression of the model on the data without a constant.

Figure 10: Cross-MSA Simulations vs. Data



Note: Scatter plots of data vs. simulation results for 97 MSAs in regression analysis. The red X represents the national simulation and each black dot is an MSA. The 45-degree line illustrates a perfect match between the model and the data. The variable being plotted shown in each plot's title. Data is fully described in appendix A.4. The calibration methodology described in text and appendix A.4. The figure shows the model performs well for prices and the number of foreclosures but is off by a constant for volume.

markets, credit constraints and losses on levered properties reducing the purchasing power of buyers (Stein, 1995; Ortalo-Magne and Rady, 2006), a decline in household formation and immigration, a reluctance on the part of retirees to sell their house in a down market, nominal loss aversion (Genesove and Mayer, 2001), increasing returns to scale in matching (Ngai and Tenreyro, 2010), and a reduction in the number of transactions by speculators who flip houses quickly. The cause of the massive decline in volume in the housing downturn is an important subject for future research.

Because it under-estimates the sales decline, the model also under-estimates the REO share in locations that had extremely high amounts of foreclosures, although when we include a constant in a regression of the simulated results on the data we cannot reject a coefficient of one. Because the vast majority of sellers also become buyers, a decline in sales would strengthen the choosey buyer effect, as REO sellers would take up a greater fraction of the market. It would, however, have a much smaller impact on market tightness, because a reduction in the number of buyers and sellers would reduce both the numerator and the denominator. We expect the overall magnitude of the combined general equilibrium effect of foreclosures to be similar.

What do these figures imply about the quantitative extent to which foreclosures exacerbate housing downturns? In the national data, the permanent price decrease that would occur without default is 21.5% (.24 log points) and with default is 33.5% (.41 log points). This implies that the general equilibrium effects of foreclosures together with the compositional effects on the price index induced by a high REO share made the downturn 56% worse than it would have been in the absence of foreclosure. Equivalently, foreclosures account for 36% of the price decline. This figure is larger in MSAs with larger busts, more default, and a bigger price-default spiral.

The 56% figure, however, includes compositional effects and is thus not the best measure of how much the general equilibrium effects of foreclosure reduce the price a *retail* seller would get if they wanted to sell at the bottom of the market. This is the relevant price for determining negative equity and thus defaults. An alternate metric of the extent to which foreclosures exacerbate downturns, then, is the decline in the retail-only price index, which is 28.7% (.34 in log points) with default and 21.5% without default. The price decline in the retail market is thus 34% worse than it would have been in the absence of foreclosures.

Perhaps surprisingly, these quantitative results are not dramatically changed with an REO discount of 10% in steady state as suggested by Clauretje and Daneshvary (2009) and Zillow (2010). In this case, the same calibration procedure implies a permanent price decline of 22.4% and an average time out of the owner-occupied market for foreclosures of 1.3 years. Intuitively, with the compositional effect weakened by a smaller foreclosure discount, the calibration implies a slightly larger permanent price decline and a stronger market tightness effect. With a 10% steady state discount, the model implies that foreclosures exacerbate the aggregate price decline by 50%. See Appendix A.2.2 for details.

These magnitudes are larger than those implied by other papers. Mian et al.'s (2012) empirical study comes closest to our results. By comparing neighborhoods in states that require judicial approval of foreclosure with neighborhoods just over a border in states that do not, they find that foreclosures were responsible for 20 to 30% of the decline in prices. Our analogous figure of 36% is only slightly higher, likely because we consider market-wide effects that comparing neighborhoods only partially picks up. Calomiris et al. (2008) use a panel VAR to analyze the effect of foreclosures on housing market equilibrium and find that foreclosures would reduce prices by 5.5 percentage points in a foreclosure wave, about half what we find. However, they simulate the impulse response to a wave of foreclosures without a bursting bubble that puts a substantial fraction of homeowners under water, which dramatically increases the size and length of the foreclosure wave. Using a calibrated macro model that focuses on how foreclosures can constrict credit

supply, Chatterjee and Eyigungor (2011) find that foreclosures account for 16% of the overall price decline.³¹

8 Welfare and Policy Implications

8.1 Welfare

To evaluate welfare we adopt a utilitarian social welfare function that equally weights all agents. We can construct social welfare as the discounted sum of individual flow utilities:

$$W = \sum_{t=0}^{\infty} \beta^t \left\{ v_n(t) m_n + v_d(t) m_d + (l_1(t) + f(t)) (h_n^* + 1/\lambda) + \frac{f(t)}{\phi f(t) + 1} c + q(t) \right\}$$

where $q(t)$ follows the law of motion:

$$q(t+1) = (1 - \gamma) q(t) + v_b(t) q_b(\mu(t)) \sum_m r_m(t) (1 - F(h_m(t))) (h_m(t) + 1/\lambda), \quad q(0) = 0.$$

where $q(t)$ denotes the expected flow housing services generated to homeowners in l_0 at time t . We also assume that a foreclosure completion entails certain costs to the bank, such as legal fees and lost revenue from interest payments. A 2008 report by Standard & Poors estimates these costs of foreclosure in excess of the loss on the sale to be approximately \$10,000, so that we set $c = -10$. Finally, we suppose that individuals who receive a taste shock but are unable to move due to negative equity receive no flow utility from their mismatched house.

There are competing effects of foreclosures on social welfare. First, total welfare is decreased relative to the steady state since securing a foreclosure completion is costly, foreclosed homes are sold by REO sellers with higher holding costs, and all homes take longer to sell. More significantly, welfare is decreased by the fact that a number of homeowners receive no flow utility from housing because they are excluded from the housing market for a period of time due to default or locked into a house that does not suit their needs. Second, buyers who do participate in the market are on average purchasing homes which they value more than homes purchased in steady state. Because these buyers stay in these houses for a median of 9 years, this generates a substantial positive effect on welfare that is consistent with anecdotal evidence of the downturn being a “buyer’s market.” Note that the decline in prices which accompanies the downturn has no direct impact on welfare since it operates simply as a transfer from sellers to buyers that has no effect on social welfare. Ultimately, when all of the various effects net out, social welfare falls, but the decline is modest.

This likely understates policy makers’ perception of the social impact of foreclosures. The welfare calculation uses a utilitarian framework and a high discount rate. Given that the downturn is temporary and a number of individuals actually benefit from the housing downturn in the form of increased housing services relative to steady state, it is not surprising we find a modest decline in welfare. However, it is still the case that a substantial mass of individuals are substantially worse off for several years. To the extent that policy makers adopt a Rawlsian short-term perspective, the social impact of foreclosures could be large.

³¹Our results also relate to an empirical literature that examines the effects of REO sales on the sale prices of extremely nearby houses. These papers typically find that a single REO listing reduce the prices of neighboring properties by 1%, with a nonlinear effect (Campbell et al, 2011). The literature is divided as to the mechanism. Anenberg and Kung (2012) arguing that REOs increase supply at an extremely-local level consistent with our market tightness effect, although at a very local level. By contrast, Gerardi et al. (2012) argue that the owners of distressed property reduce investment in their home, effects for which we attempt to control. These papers typically include fine geographic fixed effects and consequently do not pick up the search-market-level effects that are substantial in our model.

Table 8: Effects of Foreclosure Policies

Policy	Baseline	$\phi = 4000$	$\phi = 6000$	7% \rightarrow 4%	$L \downarrow \$2K$	$L \downarrow \$5K$
$\max \Delta \log (P)$	-414	-400	-387	-382	-407	-396
$\max \Delta \log (P_{\text{Retail}})$	-343	-344	-347	-327	-339	-334
$\max \Delta \log (\text{Sales}_{\text{Existing}})$	-226	-205	-189	-210	-223	-217
$\max \Delta \log (\text{Sales}_{\text{Retail}})$	-352	-322	-295	-314	-339	-321
$\max \frac{\text{Sales}_{\text{REO}}}{\text{Sales}_{\text{Existing}}}$	21.44	17.49	12.61	17.47	20.63	19.41
% Ever Foreclosed	7.55	7.54	7.48	6.10	7.27	6.86

Note: This table shows the impact of various policies on market equilibrium. The first and second rows show the maximum change in log aggregate and retail price, the third and fourth rows show the maximum change in log existing sales and retail sales, the fourth row shows the maximum REO share, and the sixth row shows the fraction of homes foreclosed upon. The first column shows the baseline, which is as in Section 7. The second and third columns show the effects of slowing down foreclosures, as $\frac{1}{\phi}$ foreclosures can be processed each week. $\phi = 4000$ and $\phi = 6000$ correspond to a maximum of 1.3 and .9 percent of the housing stock being foreclosed upon per year, respectively. The fourth column shows the effect of reducing interest rates for all homeowners from 7 to 4 percent, a generous estimate of the potential effects of refinancing mortgages. The sixth and seventh columns show the effects of reducing principal by \$2,000 and \$5,000 for all homeowners, corresponding to a \$100 billion principal reduction that is either untargeted or only targeted at under water homeowners.

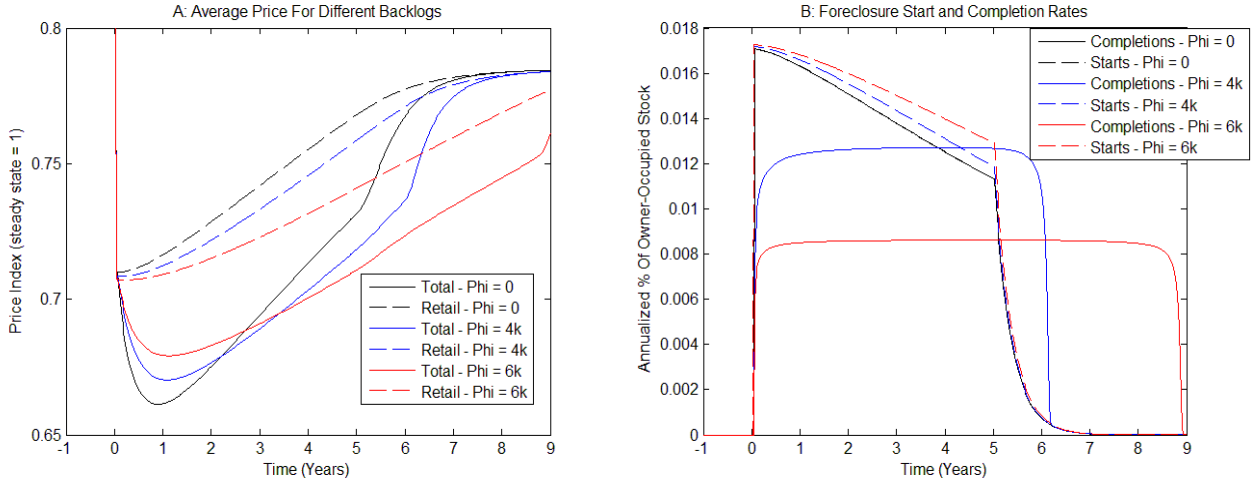
Most importantly, by focusing only on the housing market, the model misses a number of other potential normative implications of foreclosures. As discussed by Iacoviello (2005), house price declines can have pecuniary externalities because a collapse in home prices destroys wealth in the form of home equity and can impede borrowing by households and firms, creating a financial accelerator effect similar to Kiyotaki and Moore (1997). Moreover, lock-in due to negative equity can impede labor mobility (Ferriera, Gyourko and Tracy, 2010) and exacerbate structural unemployment and can increase the effective risk faced by households since housing consumption is not adjustable (Chetty and Szeidl, 2007). Additionally, banks may be forced to realize substantial losses on foreclosed properties, which impacts their balance sheets through the well-documented net-worth channel in financial intermediation, potentially leading to problems in the interbank repo market, cash hoarding by banks, and a freezing up of credit. Finally, the presence of substantial numbers of foreclosed homes can have negative externalities on communities (Campbell et al., 2011) and can reduce residential investment, construction employment, and consumption (Mian et al, 2011).

While we leave detailed analyses of these important issues to future research, we believe the discussion in this section illustrates there is value in understanding the effectiveness of various policies in ameliorating the foreclosure crisis. We thus conduct a basic positive analysis of three policies that have been proposed that fit into our model: delaying foreclosure, refinancing mortgages at lower interest rates, and reducing principal. To assess the maximum potential impact of each policy, we introduce the policy at time 0. The results of our policy simulations, discussed in the following subsections, are shown in Table 8. We use the national loan balance distribution and compare the housing market under each policy to a baseline of no policy that is shown in column 1.

8.2 Delaying Foreclosure

A simple and low-cost policy that has been proposed is slowing down the pace of foreclosures. To incorporate sluggish foreclosure into our model, we assume that when a homeowner defaults the bank begins foreclosure

Figure 11: Policy: Various Sized Backlogs



Note: The figure shows the effect of prices and foreclosure start and completion rates for three different backlogs. The model is calibrated to the national calibration developed in section 7. The pre-downturn price level is normalized to one. One in ϕ houses can be foreclosed upon each week, so $\phi = 4,000$ corresponds to 1.3 percent of the housing stock being foreclosed upon per year and $\phi = 6,000$ corresponds to .9 percent. The figure shows that slowing down foreclosures extends the downturn and makes prices remain low for longer. Although the overall price index rises, this is because of a compositional effect and the retail price index falls.

proceedings but that only $\frac{1}{\phi}$ foreclosures can be processed by the system each week.³² While in the foreclosure process, it is possible for prices to rise and the house to no longer be in negative equity. If this happens, the foreclosure “cures” and the homeowner lists their house as a normal seller but subsequently become a renter because of the liquidity shock they experienced. In table Table 8, we compare the baseline of $\phi = 0$ to cases when $\phi = 4,000$ and $\phi = 6,000$ so that the maximum annual pace of foreclosure is given by 1.3 percent per year and .9 percent per year, respectively. To further elucidate the effects of foreclosure backlogs, Figure 11 shows the aggregate and retail price indices and foreclosure starts and completions for the three values of ϕ .

A tighter foreclosure pipeline has different implications for the prices of retail homes versus the overall price index. In particular, the maximum decline in the overall price index falls because the compositional effects of foreclosure are weakened. However, the maximum retail price decline is greater and the price declines last for longer because the foreclosure crisis is extended: as panel B shows, even though foreclosure starts fall off after 5 years, with $\phi = 4,000$ the wave of foreclosures lasts over 6 years and with $\phi = 6,000$ it lasts nearly 9 years. This increases the duration of both the market tightness and choosey buyer effects which gets capitalized into lower retail prices.

The effect of foreclosure backlogs in our model is consistent with the argument that delaying foreclosures does not substantially prevent foreclosures in the long run and only draws out the pain. However, there may be benefits to delaying foreclosure that are not captured by the pure backlog story. For instance, if one expects household formation to pick up and boost demand in the near future, delaying foreclosures from

³²Formally, we assume that if $f(t)$ homes are in the foreclosure process pending approval only $\frac{f(t)}{\phi f(t)+1}$ can be processed in a given period. We choose this function as a smooth approximation to $\min\left\{f, \frac{1}{\phi}\right\}$, which processes up to $\frac{1}{\phi}$ foreclosures each period. Such an approximation is necessary for the numerical implementation.

a period of low demand to a period of higher demand could limit price declines. Similarly, slowing down foreclosures could cause banks to offer more mortgage modifications or short sales, reducing the number of delinquencies that result in a foreclosure.

In fact, the empirical evidence on states with judicial approval of foreclosure – in which backlogs are much larger (Mian et al., 2012) – suggests that slowing down foreclosures might reduce the incidence of foreclosure. Adding a judicial state dummy to regression (1) leads to a judicial dummy coefficient $+0.08$ log points for the aggregate price index and $+0.05$ for the non-distressed index even with a full set of controls, as shown in appendix A.3. Our model cannot generate such a dramatic price increase by adding a narrow foreclosure pipeline – the only way to get an effect of this order of magnitude is to reduce the incidence of foreclosures. The welfare effects of policies that limit the ability of lenders to foreclose by slowing down foreclosures are, however, unclear, as lenders may respond to a diminished ability to foreclose by increasing interest rates on mortgages or denying mortgages to credit-worthy borrowers.

8.3 Interest Rate or Payment Reductions

Another much-discussed policy that is being implemented with the Home Affordable Refinance Program is to refinance the mortgages of underwater borrowers, many of whom are stuck at extremely high interest rates due to an inability to refinance, at today’s low interest rates. This could reduce defaults because some individuals who are currently unable to meet their monthly payment may be able to pay a reduced monthly payment.

To simulate this intervention, we reduce γ_I , the hazard of default for individuals who are underwater, from 8.6 percent to 7.1 percent. Appendix A.4 uses an estimate of the effect of reducing monthly payments on defaults from Bajari et al. (2010) to show that this reduction in γ_I is equivalent to reducing interest rates from 7 to 4 percent – a generous estimate of what is possible purely through refinancing. The results are shown in column 4 of Table 8.

Although foreclosures still play an important role in exacerbating the downturn, refinancing has a substantial effect both because mechanically fewer foreclosures occur at a given level of negative equity and because the amplification mechanism is weaker. The size of these effects, however, depends critically on the effect of reducing interest rates on default. While we calibrate to the existing evidence from Bajari et al. (2010), their estimates are not causal. Understanding the impact of interest rate reductions on default is an important subject for future research.

8.4 Principal Reduction

The final policy we simulate is a \$100 billion principal reduction.³³ The results are shown in Columns 5 and 6 of Table 8. Column 5 assumes that the government cannot target underwater homeowners gives every mortgage holder a \$2,000 principal reduction. Column 6 assumes that the government can target the approximately 20 million individuals who have had negative equity during the crisis so that principal is reduced by \$5,000 for each underwater homeowner.

Principal reduction provides a direct way to reduce negative equity and the price-default amplification. The targeted principal reduction has a significant ameliorating effect on the crisis, although it is not as effective as the interest rate reduction. The smaller principal reduction, however, has an effect that is

³³We assume that raising the funds for the intervention does not affect housing demand, housing prices, or the rate at which liquidity shocks occur.

much smaller. The government’s ability to target homeowners in need is thus crucial to the effectiveness of principal reduction.

Beyond the government’s ability to target under water homeowners, costs that we do not model may limit the effectiveness of principal reduction. Chief among these is moral hazard: if people expect that underwater mortgages will be bailed out with principal reductions, they may be more likely to become delinquent on their mortgage. Similarly, strategic default may be elevated. The empirical relevance and size of such moral hazard effects is an important subject for future research.

8.5 Other Policies

Our policy simulations reveal the trade-offs faced by policy makers and the parameters that future research on anti-foreclosure policy should consider. In addition to the policies simulated here, there are a number of other policies that require a richer model to be given full justice and that we hope will be analyzed by future research.

First, a policy maker might try to stimulate additional buyer demand. To have a substantial effect on market tightness, one would have to stimulate entry by new homeowners. Such a policy is outside the scope of our model as it would require endogenous household formation or an endogenous buy-rent decision. Nonetheless, our model does suggest that any increase in new home ownership would have to be permanent; an intervention that boosts new home ownership for a few months at the expense of demand in subsequent months would not have a lasting effect. The short-lived effects of the 2009 new homeowner tax credit suggests that it is difficult to generate a long-lasting effect.

Second, our model cannot consider the conversion of owner-occupied housing to rental housing without an endogenous rent-buy margin. In particular, the conversion of REOs to rental properties has been discussed. While such a policy would reduce rental prices and REO inventories, with endogenous tenure choice it is possible that renting becomes much more appealing, drawing away buyers and further freezing up the owner-occupied market. There is also the potential for rent-seeking behavior by investors who seek to buy REO properties in bulk and convert them to rental homes.

9 Conclusion

This paper argues that foreclosures play an important role in exacerbating housing downturns due to their general equilibrium effects. We add foreclosure to a simple search model of the housing market with two types of sellers by making two additional assumptions: banks selling foreclosed homes have higher holding costs than retail sellers and homeowners who are foreclosed upon cannot immediately purchase another home.

With these assumptions, foreclosures alter market behavior by reducing the number of buyers in markets, which makes sellers and particularly REO sellers desperate to sell, and by raising the probability that a buyer meets a REO seller who sells at a discount, which makes buyers more selective. Foreclosures also alter the composition of transactions, making the average sale look more like a foreclosure sale. These effects all create downward pressure on price but have opposing effects on volume as sellers want to sell faster but buyers are more choosy. Sales fall disproportionately in the retail market, helping to explain how foreclosures freeze up the market for non-distressed homes.

We then embed our basic model of the housing market in a richer model which allows for endogenous defaults and homeowner lock-in. We elucidate the potential for spirals in which foreclosures lower prices, putting more homeowners underwater, leading to more defaults and therefore even more foreclosures. A

sensitivity analysis demonstrates that such a spiral can operate as a powerful amplification channel of shocks, especially when the proportion of homeowners in the market with high LTV ratios is high. A calibration of the full model to cross-market data is successful in matching both the average level of the price decline of the housing bust and a significant proportion of the cross-sectional variation in prices. The model matches the cross-sectional pattern of volume declines but is unable to fully account for the level. A quantitative exercise shows that foreclosures exacerbate the price declines in downturns on the order of 50 percent overall and 33 percent in the retail market.

An alternative explanation for the freezing up of the retail market during the housing bust is nominal loss aversion as documented by Genesove and Mayer (2001). The housing bubble may have created a reference point for homeowners such that when the bubble burst, they were not willing to sell for less than what they *perceived* the true value of their homes to be. If this were not the case for banks, the retail market would disproportionately freeze up. However, loss aversion would have to be extreme to explain a freezing up of the retail market for several years. Consequently, while it may not be able to fully account for the freezing up of the retail market, loss aversion may have played a role in the housing downturn and may be able to explain the volume declines that our model cannot capture. Note also that in a model in which nominal loss aversion is an operative channel, foreclosures may actually aid in price discovery.³⁴

Credit constraints and capital losses on levered houses could also explain some of the freezing up of the retail market and the decline in volume that our model cannot explain. Ortalo-Magne and Rady (2006) present a model in which homeowners use equity extracted from their previous house to purchase their next house. With down payment requirements as in Stein (1995), moderate swings in housing prices can generate large swings the purchasing power of potential homeowners. This may cause some homeowners not to move at all, creating effective lock-in of non-under-water borrowers and helping to freeze up the retail market. The substantial decline in household formation is another factor that could explain the decline in volume that our model cannot explain.

Our analysis suggests several directions for future research. First, it would be interesting to endogenize the decision to enter the housing market. Chetty and Szeidl present an (S,s) model of consumption commitments based on the decision to move homes which could be embedded into a general equilibrium model of the housing market. This would allow an analysis of how market forces affect the decision to move and elucidate why sales remain depressed and more people are not taking advantage of the buyer's market by trading up. Second, an endogenous rent-buy margin would allow the rental market to be less segmented and would allow for analyses of several additional policies, as we describe in Section 8. Finally, the addition of supply considerations would allow an analysis of how the dynamics of new construction and conversion of owner-occupied housing to renter-occupied are affected by foreclosures.

³⁴Thanks to Ed Glaeser for this insight.

10 References

- Anenberg, Elliot. 2012. "Information Frictions and Housing Dynamics." Federal Reserve Board of Governors Working Paper.
- Anenberg, Elliot, and Patrick Bayer. 2012. "Endogenous Sources of Cycling in Housing Markets: The Joint Buyer-Seller Problem." Duke University Working Paper.
- Anenberg, Elliot, and Edward Kung. 2012. "Estimates of the Size and Source of Price Declines Due to Nearby Foreclosures." Duke University Working Paper.
- Albrecht, James, Pieter Gautier, and Susan Vroman. 2010. "Directed Search in the Housing Market." Georgetown University Working Paper.
- Bajari, Patrick, Sean Chu, and Minjung Park. 2010. "An Empirical Model of Subprime Mortgage Default from 2000 to 2007." University of Minnesota Working Paper.
- Bhutta, Neil, Jane Jokko, and Hui Shan. 2011. "Consumer Ruthlessness and Mortgage Default During the 2007-2009 Housing Bust." Federal Reserve Board Working Paper.
- Calomiris, Charles, Stanley Longhofer, and William Miles. 2008. "The Foreclosure-House Price Nexus: Lessons from the 2007-2008 Housing Turmoil." NBER Working Paper 12918.
- Campbell, John and Joao Cocco. 2011. "A Model of Mortgage Default." NBER Working Paper 17516.
- Campbell, John, Stefano Giglio, and Parag Pathak. 2011. "Forced Sales and House Prices." *American Economic Review*, 101(5): 2108-31.
- Chatterjee, Satyajit, and Burcu Eyigungor. 2011. "A Quantitative Analysis of the U.S. Housing and Mortgage Markets and the Foreclosure Crisis." Federal Reserve Bank of Philadelphia Working Paper No. 11-26.
- Chetty, Raj and Adam Szeidl. 2007. "Consumption Commitments and Risk Preferences." *Quarterly Journal of Economics* 122(2): 831-877.
- Clauretje, Terrence and Nasser Daneshvary. 2009. "Estimating the House Foreclosure Discount Corrected for Spatial Price Interdependence and Endogeneity of Marketing Time." *Real Estate Economics* 37(1): 43-67.
- Deng, Yongheng, John M. Quigley, and Robert Van Order. 2000. "Mortgage Terminations, Heterogeneity and the Exercise of Mortgage Options." *Econometrica* 68(2): 275-307.
- Diaz, Antonia, and Belen Jerez. 2009. "House Prices, Sales and Time on the Market: A Search-Theoretic Framework." Universidad Carlos II de Madrid Working Paper 09-25.
- Duffie, Darrell, Nicolae Garleanu, and Lasse Pedersen. 2007. "Valuation in Over-the-Counter Markets." *Review of Financial Studies*, 20: 1865-1900.
- Duffie, Darrell, Nicolae Garleanu, and Lasse Pedersen. 2007. "Valuation in Over-the-Counter Markets." *Review of Financial Studies*, 20: 1865-1900.
- Elul, Ronel, Nicholas S. Souleles, Souphala Chomsisengphet, Dennis Glennon, and Robert Hunt. 2010. "What 'Triggers' Mortgage Default?" Federal Reserve Bank of Philadelphia Working Paper No. 10-13.
- Fernando Ferreira, Joseph Gyurko, and Joseph Tracy. 2010. "Housing Busts and Household Mobility." *Journal of Urban Economics*, 68(1): 34-45.

- Foote, Christopher L., Kristopher Gerardi, Lorenz Gotte, and Paul S. Willen. 2009. "Reducing Foreclosures." Federal Reserve Bank of Boston Discussion Paper No. 09-2.
- Foote, Christopher L., Kristopher Gerardi, and Paul S. Willen. 2008. "Negative Equity and Foreclosure: Theory and Evidence." *Journal of Urban Economics*, 64(2): 234-245.
- Genesove, David and Lu Han. 2012. "Search and Matching in the Housing Market." *Journal of Urban Economics*, 72(1):31-45.
- Genesove, David and Christopher Mayer. 2001. "Loss Aversion and Seller Behavior: Evidence From the Housing Market." *Quarterly Journal of Economics* 116(4): 1233-1260.
- Gerardi, Kristopher, Eric Rosenblatt, Paul S. Willen, and Vincent W. Yao. 2012. "Foreclosure Externalities: Some New Evidence." Federal Reserve Bank of Boston Working Paper.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2009. "Moral and Social Constraints to Strategic Default on Mortgages." NBER Working Paper 15145.
- Iacoviello, Matteo. 2005. "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle." *American Economic Review* 95: 739-764.
- Kiyotaki, Nobuhiro, and John Moore. 1997 "Credit Cycles." *Journal of Political Economy* 105(2): 211-248.
- Korteweg, Arthur and Morten Sorensen. 2011. "Estimating Loan-to-Value and Foreclosure Behavior." Columbia University Working Paper.
- Krainer, John. 2001. "A Theory of Liquidity in Residential Real Estate Markets." *Journal of Urban Economics* 49: 32-52.
- Lamont, Owen and Jeremy C. Stein. 1999. "Leverage and House-Price Dynamics in U.S. Cities." *RAND Journal of Economics* 30(3): 498-514.
- Landvoigt, Tim, Monika Piazzesi, and Martin Schneider. 2012. "The Housing Market(s) of San Diego." Stanford University Working Paper.
- Malloy, Raven and Hui Shan. 2011. "The Post-Foreclosure Experience of U.S. Households." Federal Reserve Board Working Paper.
- Mayer, Christopher, Karen Pence, and Shane M. Sunderlund. "The Rise in Mortgage Default." *Journal of Economic Perspectives* 23(1): 27-50.
- Mian, Atif and Amir Sufi. 2009. "The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis." *Quarterly Journal of Economics* 124(4): 1449-1496.
- Mian, Atif and Amir Sufi. 2011. "House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis." *American Economic Review* 101(5): 2132-2156.
- Mian, Atif, Amir Sufi, and Francesco Trebbi. 2011. "Foreclosures, House Prices, and the Real Economy." NBER Working Paper No. 16685.
- Ngai, L. Rachel and Silvana Tenreyro. 2010. "Hot and Cold Seasons in the Housing Market." Centre for Economic Performance Discussion Paper No. 922.
- Novy-Marx, Robert. 2009. "Hot and Cold Markets." *Real Estate Economics*, 37(1): 1-22.

- Ortalo-Magne, Francois, and Sven Rady. 2006. "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints." *Review of Economic Studies*, 73: 459-485.
- Piazzesi, Monika and Martin Schneider. 2009. "Momentum Traders in the Housing Market: Survey Evidence and a Search Model." *American Economic Review*, 99(2): 406-411.
- Sainai, Todd and Nicholas S. Souleles. 2009. "Can Owning a Home Hedge the Risk of Moving?" University of Pennsylvania Working Paper.
- Springer, Thomas M. 1996. "Single-Family Housing Transactions: Seller Motivations, Price, and Marketing Time." *Journal of Real Estate Finance and Economics*, 13: 237-254.
- Shimer, Robert and Ivan Werning. 2007. "Reservation Wages and Unemployment Insurance." *Quarterly Journal of Economics*, 122(3): 1145-1185.
- Stein, Jeremy. 1995. "Prices and Trading Volume in the Housing Market: A Model with Down-Payment Effects." *Quarterly Journal of Economics*, 110(2): 379-406.
- Theologides, Stergios. 2010. "Servicing REO Properties: The Servicer's Role and Incentives." In *REO & Vacant Properties: Strategies for Neighborhood Stabilization*.
<http://www.bos.frb.org/commdev/REO-and-vacant-properties/index.htm>.
- Wheaton, William. 1990. "Vacancy, Search, and Prices in a Housing Market Matching Model." *Journal of Political Economy*, 98(6): 1270-1292.
- Williams, Joseph. 1995. "Pricing Real Assets With Costly Search." *Review of Financial Studies*, 8(1): 55-90.
- Zillow Real Estate Research. 2010. "Price Differences Between Foreclosures and Non-Foreclosures." Available online at <http://www.zillow.com/blog/research/2010/01/26/price-differences-between-foreclosures-and-non-foreclosures/>.

A Appendices

A.1 Derivations and Proofs

A.1.1 Steady State

Denoting steady state values with a star superscript, the no default assumption means that $l_1^* = 0$, $v_d^* = 0$, $v_r^* = 0$, $v_n^* = v_b^*$ and $l_0^* = 1 - v_n^*$. This implies $\mu^* = 1$, $q_s(\mu^*) = q_b(\mu^*) = \chi$, $r_n^* = 1$, and $r_d^* = 0$. Because inflows into being a seller and outflows from being a seller are equal in steady state,

$$\begin{aligned} (1 - v_n^*)\gamma &= v_n^*\chi(1 - F(h_n^*)) \\ v_n^* &= \frac{\gamma}{\gamma + \chi(1 - F(h_n^*))} \quad l_0 = \frac{\chi(1 - F(h_n^*))}{\gamma + \chi(1 - F(h_n^*))}. \end{aligned}$$

Replacing the conditional expectations of surpluses with differences of cutoffs as in (12) and setting $r_n^* = 1$, $r_d^* = 0$, and $q_s^* = q_b^* = \chi$ yields simplified steady state value functions:

$$\begin{aligned} V_h^* &= \frac{h + \beta\{\gamma(V_n^* + B^*)\}}{1 - \beta(1 - \gamma)} \\ B^* &= \frac{u_b}{1 - \beta} + \frac{(1 - \theta)\chi(1 - F(h_n^*))E[h - h_n^*|h \geq h_n^*]}{(1 - \beta)(1 - \beta(1 - \gamma))} \\ V_m^* &= \frac{m}{1 - \beta} + \frac{\theta\chi(1 - F(h_m^*))E[h - h_m^*|h \geq h_m^*]}{(1 - \beta)(1 - \beta(1 - \gamma))}, m \in \{n, d\} \\ R^* &= \frac{u_r + \beta\sigma B^*}{1 - \beta(1 - \sigma)}. \end{aligned}$$

With everything in terms of the cutoffs, a two-equation system that pins down h_n^* and h_d^* . Subtracting the cutoff condition (8) at the distressed and non-distressed cutoffs gives:

$$V_{h_n^*}^* - V_{h_d^*}^* = (m_n - m_d) + \beta[V_n^* - V_d^*].$$

Plugging in the steady state values and manipulating yields an equation that implicitly defines the difference of the cutoffs:

$$\begin{aligned} h_n^* - h_d^* &= (m_n - m_d) \frac{1 - \beta(1 - \gamma)}{(1 - \beta)} \\ &+ \theta \frac{\beta}{1 - \beta} \chi \{(1 - F(h_n^*))E[h - h_n^*|h \geq h_n^*] - (1 - F(h_d^*))E[h - h_d^*|h \geq h_d^*]\}. \end{aligned} \quad (14)$$

The second equation comes from evaluating the cutoff condition (8) at h_n^* :

$$V_{h_n^*}^* = m_n + u_b + \beta[B^* + V_n^*]. \quad (15)$$

Equations (14) and (15) define a system that can be solved for h_d^* and h_n^* . All of the other steady-state variables are written in terms of these cutoffs.

Proposition 1 *If $a < m_n + u_b + \frac{\beta(1-\gamma)\chi E[h-a]}{1-\beta(1-\gamma)}$, there exists a unique steady state of the model.*

Proof. *Because there are no REO sellers in steady state,*

$$V_n^* + B^* = \frac{m_n + u_b}{1 - \beta} + \frac{\chi(1 - F(h_m^*))E[h - h_m^*|h \geq h_m^*]}{(1 - \beta)(1 - \beta(1 - \gamma))}$$

The cutoff condition for h_n^* is:

$$\begin{aligned}\frac{h_n^* + \beta [V_n^* + B^*]}{1 - \beta(1 - \gamma)} &= m_n + u_b + \beta [V_n^* + B^*] \\ \frac{h_n^*}{1 - \beta(1 - \gamma)} &= m_n + u_b + \beta \left(\frac{1 - \beta(1 - \gamma) - \gamma}{1 - \beta(1 - \gamma)} \right) [V_n^* + B^*] \\ h_n^* &= (1 - \beta(1 - \gamma)) (m_n + u_b) + \beta(1 - \beta)(1 - \gamma) [V_n^* + B^*]\end{aligned}$$

Plugging in for $V_n^* + B^*$ and re-arranging yields:

$$\begin{aligned}h_n^* &= (1 - \beta(1 - \gamma)) (m_n + u_b) + \beta(1 - \beta)(1 - \gamma) \left[\frac{m_n + b}{1 - \beta} + \frac{\chi(1 - F(h_n^*)) E[h - h_n^* | h \geq h_n^*]}{(1 - \beta)(1 - \beta(1 - \gamma))} \right] \\ h_n^* &= m_n + u_b + \frac{\beta(1 - \gamma)\chi(1 - F(h_n^*)) E[h - h_n^* | h \geq h_n^*]}{1 - \beta(1 - \gamma)}\end{aligned}\quad (16)$$

We want to find a unique solution to this equation on $h_n^* \in [a, \infty)$. As $h_n^* \rightarrow \infty$, the RHS of equation (16) approaches $m_n - u_b$. As $h_n^* \rightarrow a$, the RHS of equation (16) approaches $m_n + u_b + \frac{\beta(1 - \gamma)\chi E[h - a]}{1 - \beta(1 - \gamma)}$. Thus as long as $a < m_n + u_b + \frac{\beta(1 - \gamma)\chi E[h - a]}{1 - \beta(1 - \gamma)}$, since both the RHS and LHS of equation (16) are continuous in h_n^* , by the intermediate value theorem we know there exists a solution on $[a, \infty)$ to equation (16). Furthermore, we know that the LHS is strictly increasing in h_n^* , while the RHS is strictly decreasing in h_n^* since:

$$\frac{d}{dx} (1 - F(x)) E[h - x | h \geq x] = -(1 - F(x)) < 0.$$

This implies that the solution to equation (16) is unique.

Finally note that

$$h_d^* - \theta \frac{\beta}{1 - \beta} \chi (1 - F(h_d^*)) E[h - h_d^* | h \geq h_d^*]$$

is a monotonically increasing function of h_d^* and thus, given the solution h_n^* to equation (16), there exists a unique solution to equation (14). ■

Note that given our assumptions it is generally the case that $a < m_n + u_b + \frac{\beta(1 - \gamma)\chi E[h - a]}{1 - \beta(1 - \gamma)}$ and thus that there is a unique equilibrium. This is because β is close to 1 and γ is close to 0, so the denominator of the fraction is very small. Uniqueness would only be a concern with a very low discount factor or high moving probability.

Due to higher holding costs and balance sheet concerns, an REO seller should be more willing to sell the property conditional on being matched with a buyer than a normal seller. We show this is always the case in steady state:

Lemma 2 For a given h , the probability of sale for a distressed seller is higher than the probability of sale for a non-distressed seller.

Proof. Note again that:

$$\frac{d}{dx} (1 - F(x)) E[h - x | h \geq x] = -(1 - F(x)) < 0.$$

Suppose that $h_n^* \leq h_d^*$. Then:

$$\begin{aligned}(h_n^* - h_d^*) &- \theta \frac{\beta}{1 - \beta} \chi \{ (1 - F(h_n^*)) E[h - h_n^* | h \geq h_n^*] - (1 - F(h_d^*)) E[h - h_d^* | h \geq h_d^*] \} \\ &< 0 \\ &< (m_n - m_d) \frac{1 - \beta(1 - \gamma)}{(1 - \beta)}.\end{aligned}$$

which contradicts equation (14). It must therefore be that $h_n^* > h_d^*$, which indicates that distressed sellers

are more likely to sell than non-distressed sellers. ■

We use the Nash bargaining condition to back out steady state prices. We have for $m \in \{n, d\}$ and a given h :

$$\frac{\theta}{1-\theta} = \frac{S_m^S}{S_m^B} = \frac{p_m^*(h) - m - \beta V_m^*}{V_h^* - p_m^*(h) - \beta B^*}$$

And so

$$\begin{aligned} \theta [V_h^* - p_m^*(h) - \beta B^*] &= (1-\theta) [p_m^*(h) - m - \beta V_m^*] \\ p_m^*(h) &= \theta V_h^* - \beta \theta B^* + (1-\theta)m + (1-\theta)\beta V_m^* \\ &= \theta V_h^* + m + \beta V_m^* - \theta [m + \beta V_m^* + \beta B^*] \\ &= \theta [V_h^* - V_{h_m^*}] + m + \beta V_m^* \\ &= \frac{\theta(h - h_m^*)}{1-\beta(1-\gamma)} + m + \frac{\beta m}{1-\beta} + \frac{\beta \theta \chi (1-F(h_m^*)) E[h - h_m^* | h \geq h_m^*]}{(1-\beta)(1-\beta(1-\gamma))} \\ &= \frac{\theta(h - h_m^*)}{1-\beta(1-\gamma)} + \frac{m}{1-\beta} + \frac{\beta \theta \chi (1-F(h_m^*)) E[h - h_m^* | h \geq h_m^*]}{(1-\beta)(1-\beta(1-\gamma))} \end{aligned} \quad (17)$$

It is also always the case that distressed properties sell for less than non-distressed properties:

Proposition 3 *Distressed sales trade at a constant discount, in the sense that $p_n^*(h) - p_d^*(h) = \Delta$ for all $h \geq h_n^* > h_d^*$ for some constant $\Delta > 0$.*

Proof. Taking the difference of prices we get:

$$\begin{aligned} p_n^*(h) - p_d^*(h) &= \frac{\theta(h_d^* - h_n^*)}{1-\beta(1-\gamma)} + \frac{m_n - m_d}{1-\beta} \\ &\quad + \frac{\beta \theta \chi \{(1-F(h_n^*)) E[h - h_n^* | h \geq h_n^*] - (1-F(h_d^*)) E[h - h_d^* | h \geq h_d^*]\}}{(1-\beta)(1-\beta(1-\gamma))} \end{aligned}$$

Using equation (14) this simplifies to:

$$\begin{aligned} p_n(h) - p_d(h) &= \frac{\theta(h_d^* - h_n^*)}{1-\beta(1-\gamma)} + \frac{m_n - m_d}{1-\beta} - \frac{m_n - m_d}{1-\beta} + \frac{h_n^* - h_d^*}{1-\beta(1-\gamma)} \\ &= \frac{1-\theta}{1-\beta(1-\gamma)} (h_n^* - h_d^*) \equiv \Delta > 0 \end{aligned}$$

■

A.1.2 Dynamics

This appendix describes how we solve the dynamic model with exogenous defaults presented in section 3.

First, the laws of motion simply add inflows and subtract outflows according to figure 5:

$$\begin{aligned} l_0(t+1) &= (1-\gamma)l_0(t) + v_b(t)q_b(\mu(t)) \sum_m r_m(t) e^{-\lambda(h_m(t)-a)} \\ l_1(t+1) &= (1-\gamma)l_1(t) \\ v_n(t+1) &= \gamma l_0(t) + \gamma(1-\alpha)l_1(t) + v_n \left[1 - q_s(\mu(t)) e^{-\lambda(h_n(t)-a)} \right] \\ v_d(t+1) &= \gamma \alpha l_1(t) + v_d \left[1 - q_s(\mu(t)) e^{-\lambda(h_d(t)-a)} \right] \\ v_b(t+1) &= \gamma l_0(t) + \gamma(1-\alpha)l_1(t) + v_r(t)\sigma + v_b(t) \left[1 - q_b(\mu(t)) \sum_m r_m(t) e^{-\lambda(h_m(t)-a)} \right] \\ v_r(t+1) &= \gamma \alpha l_1(t) + (1-\sigma)v_r(t) \end{aligned}$$

To generate the dynamic path of V_h , we expand the sum in equation (3) and collecting terms:

$$V_h(t) = \frac{h}{1 - \beta(1 - \gamma)} + \beta \sum_{j=1}^{\infty} [(\beta(1 - \gamma))^{j-1} \{\gamma[V_n(t+j) + B(t+j)]\}].$$

The sum in the second term can be written recursively as:

$$\Gamma(t) = \gamma[V_n(t) + B(t)] + \beta(1 - \gamma)\Gamma(t+1) \quad (18)$$

so that

$$V_h(t) = \frac{h}{1 - \beta(1 - \gamma)} + \beta\Gamma(t+1).$$

The entire dynamic system can thus be written recursively. The cutoff rule then simplifies to:

$$\frac{h_m(t)}{1 - \beta(1 - \gamma)} + \beta\Gamma(t+1) = u_b + \beta B(t+1) + m + \beta V_m(t+1). \quad (19)$$

The full general equilibrium model is made up of 13 endogenous variables and 13 equations. The endogenous variables are the cutoffs h_n and h_d , the masses v_n , v_d , v_b , v_r , l_0 , and l_1 , and value functions Γ , V_{m_d} , V_{m_r} , B , and R . From these values, all of the other endogenous parameters of the model can be determined. Substituting out the conditional expectation of the surplus using (12), using the exponential distribution, and using the definitions of q_b and q_s from the matching function gives the dynamic system:

$$\begin{aligned} V_n(t) &= \beta V_n(t+1) + m_n + \chi\mu(t) \frac{\theta}{\lambda[1 - \beta(1 - \gamma)]} e^{-\lambda(h_n(t)-a)} \\ V_d(t) &= \beta V_d(t+1) + m_d + \chi\mu(t) \frac{\theta}{\lambda[1 - \beta(1 - \gamma)]} e^{-\lambda(h_d(t)-a)} \\ B(t) &= \beta B(t+1) + u_b + \chi\mu(t) \frac{1 - \theta}{\lambda[1 - \beta(1 - \gamma)]} \sum_m r_m(t) e^{-\lambda(h_m(t)-a)} \\ R(t) &= u_r + \beta \{\sigma B(t+1) + (1 - \sigma)R(t+1)\} \\ \Gamma(t) &= \gamma[V_n(t) + B(t)] + \beta(1 - \gamma)\Gamma(t+1) \\ l_0(t+1) &= (1 - \gamma)l_0(t) + v_b(t) \chi\mu(t) \sum_m r_m(t) e^{-\lambda(h_m(t)-a)} \\ l_1(t+1) &= (1 - \gamma)l_1(t) \\ v_n(t+1) &= \gamma l_0(t) + \gamma(1 - \alpha)l_1(t) + v_n \left[1 - \chi\mu(t) e^{-\lambda(h_n(t)-a)}\right] \\ v_d(t+1) &= \gamma\alpha l_1(t) + v_d \left[1 - \chi\mu(t) e^{-\lambda(h_d(t)-a)}\right] \\ v_b(t+1) &= \gamma l_0(t) + \gamma(1 - \alpha)l_1(t) + v_r(t)\sigma + v_b(t) \left[1 - \chi\mu(t) \sum_m r_m(t) e^{-\lambda(h_m(t)-a)}\right] \\ v_r(t+1) &= \gamma\alpha l_1(t) + (1 - \sigma)v_r(t) \\ \frac{h_n(t)}{1 - \beta(1 - \gamma)} + \beta\Gamma(t+1) &= u_b + \beta B(t+1) + m_n + \beta V_n(t+1) \\ \frac{h_d(t)}{1 - \beta(1 - \gamma)} + \beta\Gamma(t+1) &= u_b + \beta B(t+1) + m_d + \beta V_d(t+1) \end{aligned}$$

where $\mu(t) = \frac{v_b(t)}{v_n(t) + v_d(t)}$ and $r_m(t) = \frac{v_m(t)}{v_n(t) + v_d(t)}$. We solve this system using Newton's Method as implemented in DYNARE, which guesses that the model returns to steady state at time T , solves a system of $13T$ equations, and checks that the model is in fact within ε of the steady state at time T . Solving the model with endogenous defaults of section 5 is performed similarly although the laws of motion are modified as described in appendix A.1.4.

We then back out prices as in the steady state. Prices are defined by (13). The mean price for a type

m seller is then:

$$\bar{p}_m = E_h [p_{m,h}(t) | h_m \geq h_m] = \frac{\theta}{\lambda} \frac{1}{1 - \beta(1 - \gamma)} + m + \beta V_m(t+1) \quad (20)$$

and the overall mean price in a price index is:

$$\begin{aligned} \bar{p} = & FVol_N \left[\frac{\theta}{\lambda} \frac{1}{1 - \beta(1 - \gamma)} + m_n + \beta V_n(t+1) \right] \\ & + FVol_D \left[\frac{\theta}{\lambda} \frac{1}{1 - \beta(1 - \gamma)} + m_d + \beta V_d(t+1) \right] \end{aligned} \quad (21)$$

where $FVol_m$ is the fraction of total volume accounted for by type m sellers.

A.1.3 Calibration

As described in the main text, we use five aggregate moments from the housing market prior to the crash to set a , λ , m_n , and m_d . These moments are the average price of a normal home in steady state \bar{p}_n^* , the variance of the residual price distribution $\sigma_{p_n^*}^2$, the discount for a distressed sale in terms of mean prices $\frac{\bar{p}_n - \bar{p}_d}{\bar{p}_n}$, and the time on the market for a normal sale T_n^* and a distressed sale T_d^* .

Using the expressions for the price and the probability of sale in the main text along with properties of the exponential distribution, these moments are:

$$\begin{aligned} \bar{p}_n^* &= \frac{\theta}{\lambda [1 - \beta(1 - \gamma)]} + \frac{m_n}{1 - \beta} + \frac{\beta \theta \chi e^{-\lambda(h_n^* - a)}}{\lambda(1 - \beta)(1 - \beta(1 - \gamma))} \\ \sigma_{p_n^*}^2 &= \frac{\theta^2}{\lambda^2 [1 - \beta(1 - \gamma)]^2} \\ \frac{\bar{p}_n - \bar{p}_d}{\bar{p}_n} &= \frac{m_n - m_d}{\bar{p}_n(1 - \beta)} + \frac{\beta \theta \chi \{e^{-\lambda(h_n^* - a)} - e^{-\lambda(h_d^* - a)}\}}{\lambda \bar{p}_n(1 - \beta)(1 - \beta(1 - \gamma))} = \frac{h_n^* - h_d^*}{\bar{p}_n [1 - \beta(1 - \gamma)]} \\ T_n^* &= \frac{1 - \chi \exp(-\lambda(h_n^* - a))}{\chi \exp(-\lambda(h_n^* - a))} = \frac{1}{\chi} \exp(\lambda(h_n^* - a)) - 1 \\ T_d^* &= \frac{1 - \chi \exp(-\lambda(h_d^* - a))}{\chi \exp(-\lambda(h_d^* - a))} = \frac{1}{\chi} \exp(\lambda(h_d^* - a)) - 1 \end{aligned}$$

Plugging the second and fourth equations into the first gives:

$$\bar{p}_n^* = \sigma_{p_n^*} + \frac{m_n}{1 - \beta} + \frac{\beta}{1 - \beta} \frac{\sigma_{p_n^*}}{T_n^* + 1}$$

which implicitly defines m_n as a function of known parameters and observable moments.

We then define a six equation system with six variables – a , θ , λ , m_d , h_n^* , and h_d^* – that we use to calibrate the remainder of the model. Taking the square root of the second equation and rearranging gives λ as a function of θ and observable moments:

$$\lambda = \frac{\theta}{\sigma_{p_n^*} [1 - \beta(1 - \gamma)]}.$$

An expression for a is obtained by inverting the fourth equation $h_n^* = a + \frac{1}{\lambda} \ln(\chi(T_n^* + 1))$ and then plugging into the cutoff condition for n :

$$\begin{aligned} V_{h_n^*} &= m_n + u_b + \beta [B^* + V_{m_d}^*] \\ h_n^* &= (1 - \beta(1 - \gamma)) m_n + \beta(1 - \gamma) \left[m_n + \frac{\chi e^{-\lambda(h_n^* - a)}}{\lambda(1 - \beta(1 - \gamma))} \right] \end{aligned}$$

Plugging in for h_n^* and solving gives:

$$a = (1 - \beta(1 - \gamma))m_n + \beta(1 - \gamma) \left[m_n + \frac{1}{(T_n^* + 1)\lambda(1 - \beta(1 - \gamma))} \right] - \frac{1}{\lambda} \ln(\chi(T_n^* + 1))$$

The equations for λ and a , the moments for T_n^* , T_d^* , and the discount,

$$\frac{\bar{p}_n - \bar{p}_d}{\bar{p}_n} = \frac{h_n^* - h_d^*}{\bar{p}_n [1 - \beta(1 - \gamma)]} \quad T_n^* = \frac{1}{\chi} \exp(\lambda(h_n^* - a)) - 1 \quad T_d^* = \frac{1}{\chi} \exp(\lambda(h_d^* - a)) - 1,$$

along with (14),

$$h_n^* - h_d^* = (m_n - m_d) \frac{1 - \beta(1 - \gamma)}{(1 - \beta)} + \frac{\theta}{\lambda} \frac{\beta}{1 - \beta} \chi \left\{ e^{-\lambda(h_n^* - a)} - e^{-\lambda(h_d^* - a)} \right\},$$

form the six equation system, which we solve numerically.

Although all of the variables are jointly determined, we have found that θ and the gap between m_d and m_n are principally determined by the gap in time on the market and the REO discount while a and λ are principally determined by the moments of the price distribution.

A.1.4 Extended Model

For the extended model, the housing market is unchanged and so the value functions are unchanged. Only the laws of motion differ, as described by Figure 8. As described in the text, we have two different exogenous shocks. First, we assume that a fraction $\delta(t)$ of individuals who sell due to taste shocks become renters instead of buyers and shock $\delta(t)$. This leads to the following laws of motion:

$$\begin{aligned} l_0(t+1) &= (1 - \gamma)l_0(t) + v_b(t)q_b(\mu(t)) \sum r_m(t)(1 - F(h_m(t))) + w(t) \frac{G(V_n(t)) - G(V_n(t-1))}{1 - G(V_n(t-1))} \\ l_1(t+1) &= (1 - \gamma)l_1(t) \\ w(t+1) &= (1 - \gamma_I)w(t) + (\gamma - \gamma_I)l_1(t)(1 - G(V_n(t))) - w(t) \frac{G(V_n(t)) - G(V_n(t-1))}{1 - G(V_n(t-1))} \\ f(t+1) &= \gamma_I l_1(t)(1 - G(V_n(t))) + \gamma_I w(t) \\ &\quad + f(t) \left(1 - \frac{1}{\phi f(t) + 1} \right) - f(t) \frac{G(V_n(t)) - G(V_n(t-1))}{1 - G(V_n(t-1))} \\ v_n(t+1) &= \gamma l_0(t) + \gamma l_1(t)G(V_n(t)) + (w(t) + f(t)) \frac{G(V_n(t)) - G(V_n(t-1))}{1 - G(V_n(t-1))} \\ &\quad + v_n[1 - q_s(\mu(t))(1 - F(h_n(t)))] \\ v_d(t+1) &= \frac{f(t)}{\phi f(t) + 1} + v_d[1 - q_s(\mu(t))(1 - F(h_d(t)))] \\ v_b(t+1) &= (1 - \delta)[\gamma l_0(t) + \gamma l_1(t)G(V_n(t))] + \\ &\quad + v_r(t)\sigma + v_b(t) \left[1 - q_b(\mu(t)) \sum r_m(t)(1 - F(h_m(t))) \right] \\ v_r(t+1) &= \delta \left[\gamma_I l_0(t) + \gamma_I l_1(t)G(V_n(t)) + w(t) \frac{G(V_n(t)) - G(V_n(t-1))}{1 - G(V_n(t-1))} \right] \\ &\quad + f(t) \frac{G(V_n(t)) - G(V_n(t-1))}{1 - G(V_n(t-1))} + \frac{f(t)}{\phi f(t) + 1} + (1 - \sigma)v_r(t) \end{aligned}$$

Second, we assume that a falls permanently and that γ_I declines gradually after 10 years. This is the same as setting $\delta = 0$ above, shocking a , and using the following auto-regressive process for a :

$$\gamma_I = \tau_t \bar{\gamma}_I \text{ where } \tau_t = \alpha \tau_{t-1} \text{ and } \tau_1 = 1$$

We assume $\alpha = 1$ for five years when it falls to $\alpha = .95$.

These laws of motion simply add inflows and subtract outflows. A fraction $(1 - G(V_n(t)))$ of individuals who receive taste shocks default and the same fraction of individuals with taste shocks become locked in. A fraction δ of individuals who would become buyers and sellers become a buyer and a renter instead. A mass $\frac{f(t)}{\phi f(t)+1}$ experiences a foreclosure completion. The final added complexity is accounting for the mass of individuals who were locked in in period $t-1$ but are no longer locked in in period t or who were in foreclosure in period $t-1$ but are no longer in foreclosure in period t due to rising prices. Because only individuals with a loan balance above $V_n(t-1)$ are locked in at time $t-1$, this mass is a fraction $\frac{G(V_n(t))-G(V_n(t-1))}{1-G(V_n(t-1))}$ of the mass $w(t)$ and $f(t)$, respectively.

These laws of motion replace the laws of motion in appendix A.1.2 above. The rest of the equations are the same, yielding a 15 equation and 15 unknown dynamic system.

A.2 Details Omitted From Main Text

A.2.1 Details of Isolating Each Effect

As described in the main text we perform three experiments to isolate the role of each driving force in our model. We provide the details of these experiments here:

1. To shut down the market tightness effect, we assume that a homeowner who defaults is not forced to become a renter for a certain random amount of time, but can instead immediately re-enter the housing market as a buyer. The law of motion for the stock of buyers in the market then becomes

$$v_b(t+1) = \gamma(l_0(t) + l_1(t)) + v_b(t) \left[1 - q_b(\mu) \sum r_m(t) e^{-\lambda(h_m(t)-a)} \right],$$

while all other equations remain unchanged. Note that the stock of renters will always be zero in this experiment and $\mu(t) = 1$ for all t .

2. To shut down the choosey buyer effect, we assume that the buyer believes every seller he meets will be a retail seller. That is, even though there may well be distressed sellers in the market, the buyer fails to take their presence into account when determining his optimal market behavior. Along these lines, for this experiment we modify the Bellman equation of the buyer's value function to read:

$$B(t) = \beta B(t+1) + b + q_b(\mu(t)) \frac{1 - \theta}{\lambda[1 - \beta(1 - \gamma)]} e^{-\lambda(h_n(t)-a)}.$$

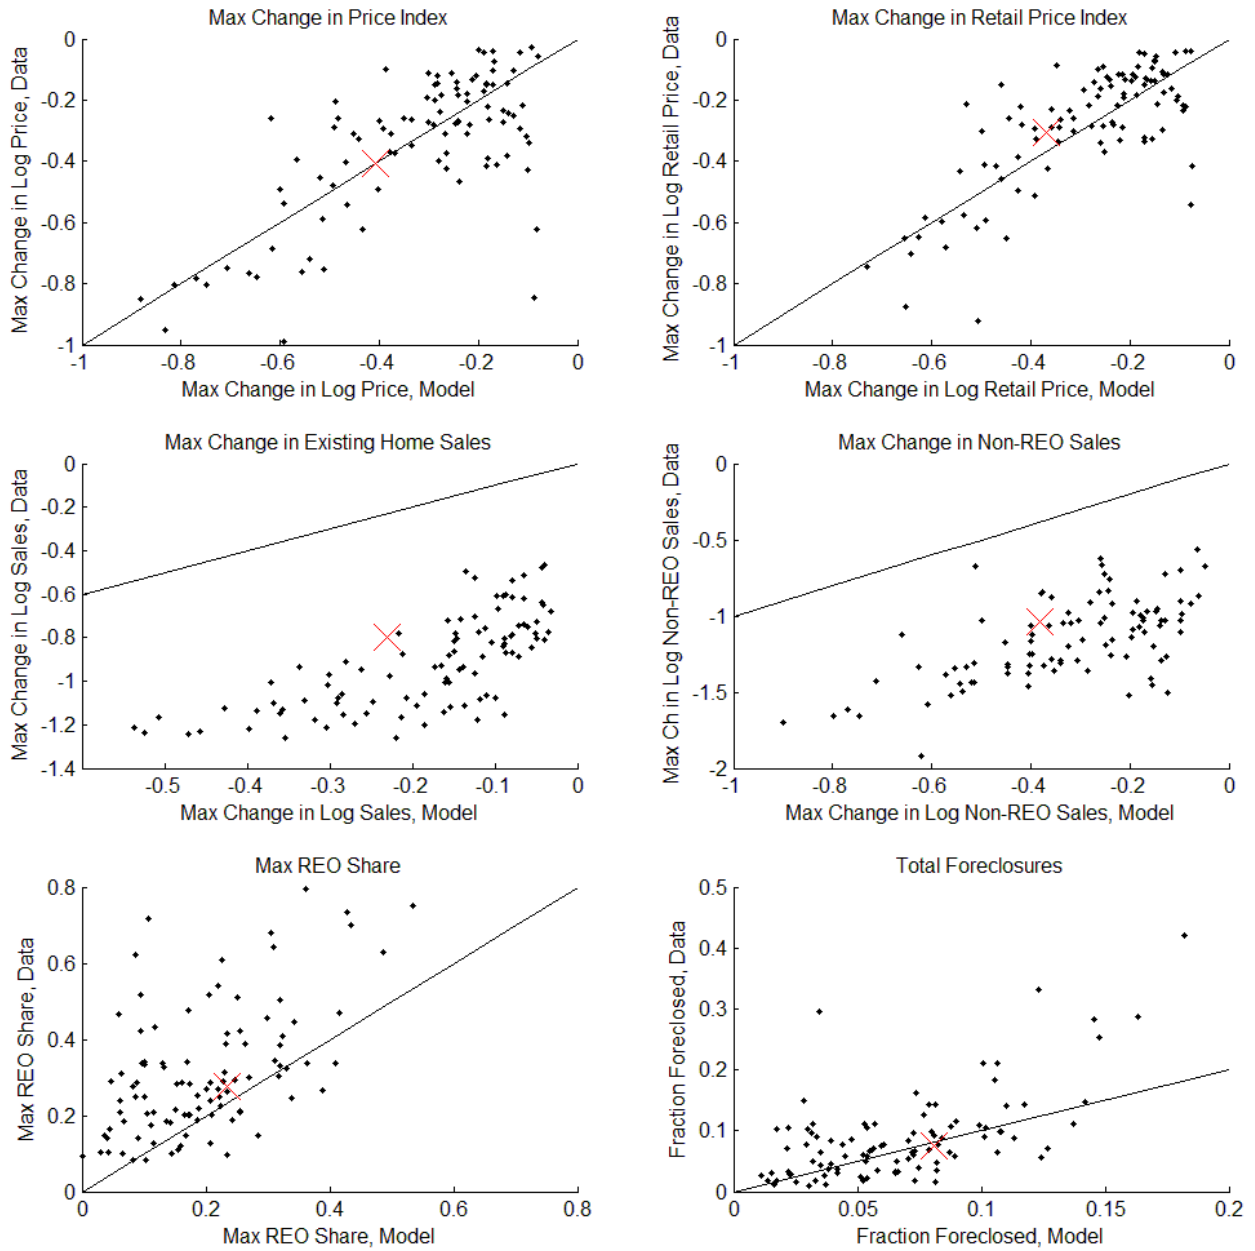
Again, we leave all other equations unchanged.

3. Finally, we run an experiment in which we include only compositional effects. For this experiment, we shut down both the market tightness and choosey buyer effects by modifying the law of motion for the stock of buyers and the Bellman equation for the buyer's value function in the manner described above.

A.2.2 Cross-Markets Analysis With 10% REO Discount

Figure 12 shows the results of the same calibration procedure in 7 for a 10% REO discount instead of a 20% REO discount. The lower REO discount weakens the compositional effect whereby a large REO share reduces the aggregate price index by mechanically placing more weight on properties that sell at a discount. The lower discount also weakens the choosey buyer effect, since the benefit of waiting for a foreclosure is reduced somewhat (though it still grows substantially in the downturn). Consequently, to match the non-linearity in price declines relative to the size of the preceding boom, the non-compositional effects of foreclosure – the market tightness effect – must be larger. This results in a longer average time out of the market of 1.3 years relative to 1.05 years for a 20% steady state discount. Additionally, the permanent price decline is 22.4%, which is slightly bigger than the 21.5% for a 20% discount. Foreclosures thus exacerbate the downturn by 50%. The permanent price decline in the model remains high in order to fit the non-linearity shown in Figure 3.

Figure 12: Cross-MSA Simulations vs. Data: 10% REO Discount



Note: Scatter plots of data vs. simulation results for 97 MSAs in regression analysis for a 10 percent discount. The red X represents the national simulation and each black dot is an MSA. The 45-degree line illustrates a perfect match between the model and the data. The variable being plotted shown in each plot's title. Data is fully described in appendix A.4. The calibration methodology described in text and appendix A.4.

Table 9: Judicial vs. Non-Judicial States

Dependent Variable	$\Delta \log (P)$	$\Delta \log (P_{\text{Retail}})$
Data	0.084 (0.035)**	0.053 (0.025)**
Model With Backlogs	0.014 (0.005)***	0.002 (0.002)
Model No Backlogs	0.008 (0.004)*	0.004 (0.002)*
N	45	45

Notes: * = 10% Significance, ** = 5% Significance *** = 1% significance. All standard errors are robust to heteroskedasticity. Every reported coefficient is for the judicial state dummy in a regression that includes a linear and quadratic term for change in log price 2003-2006, z score for share with LTV over 80 percent and its interaction with change in log price 2003-2006, and z score for share with second mortgage and its interaction with change in log price 2003-2006. These regressions do not include the Saiz (2010) variables, which are not available at the state level. The columns differ by dependent variables. The rows differ by data source: the first row shows the actual CoreLogic data, the second row uses simulated dependent variable data from a model in which judicial states have a backlog, and the third row uses simulated dependent variable data from a model in which judicial states have no backlog. Every regression has 45 states as described in Appendix A.4. We use data from Mian et al. (2012), which they obtained from RealtyTrac.com, to categorize states as judicial foreclosure or non-judicial foreclosure states. Using this methodology, Connecticut, Delaware, Florida, Illinois, Indiana, Kansas, Kentucky, Louisiana, Maine, Maryland, Massachusetts, Nebraska, New Jersey, New York, North Dakota, Ohio, Pennsylvania, South Carolina, and Vermont are judicial states.

Interestingly, because the compositional effect is reduced, the extent to which foreclosures exacerbate the retail price decline is increased from 28.7% to 37.5%. Intuitively, with the compositional effect weakened the retail price declines must be stronger. Unfortunately, this results in a worse average fit for retail price declines, as shown in panel B of Figure 12. Thus while the model fit for overall price declines is roughly comparable to the 20% steady state price decline case, we prefer the 20% case shown in the main text.

A.3 Judicial vs. Non-Judicial States

Table 9 shows the coefficient on judicial state of running a regression similar to equation (1) with a judicial state dummy as described in the table note.³⁵ The first row shows the actual data, while the second row shows the results of a model with a backlog of $\phi = 3,400$ for judicial states and no backlog for non-judicial states and the third row shows the results of a model with no backlog for judicial or non-judicial states. Adding backlogs to the model is a step in the right direction, but the model is still an order of magnitude short of the data.

As a result, we speculate that this implies that backlogs cannot be the whole story in judicial states – there must be some reduction in the incidence of foreclosures as banks respond to the long foreclosure timelines. While it is possible that $\phi > 3,400$ in judicial states – something we cannot simulate because of numerical issues – the results of higher backlogs for a national calibration shown in Figure 11 suggests that even with a much narrower foreclosure pipeline it is not possible to get that judicial states have a log price decline that is .084 smaller from foreclosure backlogs alone.

A.4 Data Sources and Calculations

Data

The main data source is proprietary data from CoreLogic, which we supplement with data from the U.S. Census, Saiz (2010), and Mian et al. (2012).

³⁵As with the main text, the results are largely unchanged if we used weighted least squares and weight by the owner-occupied housing stock.

CoreLogic provides us with a monthly data set for the nation, 50 states, and the 100 largest MSAs³⁶ that for 2000-2011 includes:

- The CoreLogic home price index and non-distressed home price index estimated from public records. We refer to these as the aggregate and retail price indices. The CoreLogic non-distressed price index differs slightly from the retail price index in the model because it excludes short sales, which we count as non-REO sales.
- 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 represented by the data, we combine short sales and resales into a non-REO existing home sales measure which we call retail sales. We calculate existing home sales by adding REO and retail 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.
- First lien originations and first lien refinancings estimated using public records.
- Over-90-day-delinquent loans, loans in foreclosure, and active loans estimated using a mortgage-level database. We use the raw counts to construct the share of active loans that are over 90 days delinquent and in foreclosure.
- The mean number of days on the market for listed homes and closed sales estimated using Multiple Listing Service data.

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 state and county-level sales counts, auctions counts, days on the market, and REO

³⁶By 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 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.

share, we smooth the data using a 5 month moving average (2 months prior, the current month, and 2 months post) to remove any blips 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, state, and county level, which we aggregate to the MSA level using MSA definitions.³⁷ 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% 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 its peak in March 2006. To approximate the size of the bubble, we average the seasonally-adjusted price index for March-May 2001, March-May 2003, March-May 2006, and March-May 2011 to calculate the change in log prices for 2001 to 2006, 2003 to 2006, and 2006 to 2011. We use these variables in our regressions and to estimate the relative size of the shock for each geographical area.

We also estimate the maximum log change in seasonally-adjusted prices, smoothed and seasonally-adjusted volume, and seasonally-adjusted time to sale as well as the maximum REO share for each geographical area. We estimate the minimum value between March 2006 and December 2011 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 December 2012 and dividing by the owner-occupied housing stock in 2006 as calculated from the ACS adjusted for CoreLogic's approximately 85% coverage, which is assumed to be constant across locations. Again, our results are not sensitive to the choice of dates.

From the 100 MSAs and 50 states, we drop two MSAs and five states. The Birmingham, Alabama MSA is dropped because a major county stopped reporting to CoreLogic in the middle of the downturn, and the Syracuse New York MSA is dropped because loan balance distribution data is not available for this MSA in 2006. Maine, Vermont, and South Dakota are dropped because loan balance distribution data is not available for these states in 2006. For the cross-state analysis, we focus on the continental U.S. and omit Alaska and Hawaii.

We finally merge data from Saiz (2010) into the MSA 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 MSA. We are able to match every MSA we have data on except for Sacramento and Honolulu.

Loan Balance Distribution Calibration

We use a minimum-distance methodology to calibrate the loan balance distribution for each geography. From the 7 quantiles given to us by CoreLogic and the Census Data on the number of owner-occupied homes without a mortgage, we construct a CDF of 6 points: the fraction of loans with under 50% LTV, under 60% LTV, under 70% LTV, under 80% LTV, and under 100% LTV. We then construct a norm for the distance between the Beta distribution and the empirical CDF. Because the upper tail of the distribution is most critical for our amplification channel, we weight the under 50%, under 60%, and under 70% parts of the distribution by .1 and the under 80%, under 90%, and under 100% by .2. We then choose b_a and b_b , the parameters of the Beta distribution, to minimize this norm. The resulting fit is close enough that our results are robust to alternate weightings of the norm.

Sources of Calculations in the Text

All figures in the introduction are tabulated from the CoreLogic data as described above.

For the calibration of the housing market model, the median tenure for owner occupants of approximately 9 years comes from table 3-9 of the American Housing Survey reports for 1997-2005. The 20% REO discount

³⁷The 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.

comes from Campbell et al.'s (2011) online appendix. They report an average discount over 1987-2009 of 26%. In table A6, they estimate this by year and show that in current housing cycle it was as low as 22.6% in 2005 and as high as 35.4% in 2009. 20 percent is thus a reasonable discount.

To determine γ_I , the incidence of income shocks for houses in negative equity, we divide the seasonally adjusted number of foreclosures by the maximum seasonally adjusted number of homes in negative equity in the CoreLogic data. The mean annual incidence is $\gamma_I = 8.6\%$.

To get that interest rates decrease the hazard of default for under-water borrowers from 8.6 percent to 7.1 percent, we use data from Bajari et al. (2010) combined with standard mortgage amortization schedules. We begin by assuming that all mortgages are at 7 percent interest rates and will be refinanced to 4 percent. The average mortgage is somewhat below 7 percent, but we choose 7 percent to reflect that some ARMs reset at quite high rates and because we want to simulate the largest possible impact of a refinancing. Assuming houses are bought with 20 percent down in steady state, a 7 percent mortgage has a monthly payment of \$1,217.50 while a 4 percent mortgage has a monthly payment of \$873.67 according to standard amortization schedules and formulas. Bajari et al. report that the mean loan in their data set has a payment-to-income ratio of .312. Assuming the \$1,217.50 monthly payment matches this ratio, monthly income is \$3,902.24. A reduction of the monthly payment to \$873.67 reduces the payment-to-income ratio by .088. Bajari et al. estimate that a one standard deviation change in the payment-to-income ratio – equivalent to a .124 change – reduces the hazard of default by 17.5 percent. Assuming linearity a reduction of .088 will reduce the hazard of default by 12.435 percent. Given the initial hazard of 8.6 percent, this implies a default hazard with a reduced interest rate of 7.1 percent.

For the principal reductions, we assume a \$100 billion dollar principal reduction. 21.5 million households potentially under water implies an approximately \$5,000 principal reduction for each house. If all 50 million households with a mortgage received the reduction, this would be only a \$2,000 principal reduction for each house.