Online Appendix

Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession

Carlos Garriga and Aaron Hedlund

Contents

A Data Appendix

The detailed micro data used in the scatter plots, heat maps, and regressions throughout the paper come from several sources. This section explains the construction of each variable, provides summary statistics, and compares the aggregate dynamics of house prices and liquidity from the CoreLogic MLS housing data to those constructed from public sources. In addition, this section presents tables with the full cross-sectional regression results.

A.1 Variable Construction

Table 9: Summary of Data Sources

Variable	Source	Raw Data Aggregation
House Prices	CoreLogic MLS	Listing Level
Housing Liquidity*	CoreLogic MLS	Listing Level
Mortgage Default	Equifax	Loan Level
Adjusted Gross Income	IRS SOI	Zip Code Level
Employment	BLS QCEW	County Level
Nontradable Employment	Census CBP	County Level

*Includes both time on the market and months of supply.

Table 9 summarizes the data sources. Housing variables are constructed from listings-level CoreLogic MLS data, loan-level credit data from Equifax is used to measure mortgage default, the IRS Statistics of Income provides income data at the zip code level, and county-level employment comes from a combination of the BLS Quarterly Census of Employment and Wages (all industries) and the Census County Business Patterns (industry specific to construct nontradable employment). National Flow of Funds data is also used to construct the housing net worth shock variable originally from Mian, Rao and Sufi (2013) that appears in some of the regressions.

A.1.1 House Prices and Liquidity

The MLS data is part of the CoreLogic Real Estate Database, which contains property-level information on listings and sales—among other variables—for residential properties around the U.S drawn from organizations of real estate agents who enter properties into an electronic MLS system in order to market them. The MLS dataset is dynamic by tracking changes in each listing over time—especially whether the property is pulled off the market and re-listed, which is a common seller tactic to move their listing to the top of the search queue for potential buyers. The dataset includes a large array of fields, but most important for the analysis in this paper is the date each property went on the market, any history of de-listings and re-listings, and the final closing price and date. The analysis is restricted to single family homes and condominiums.

House Prices House prices are the MLS transaction price at closing. This data is then adjusted for inflation using the PCE and aggregated to the county level. Counties with fewer than ten transactions in a given year are excluded.

Time on the Market Time on the market is measured as the *total* number of days a property is actively on the market before selling. However, the actual selling process for a property may entail multiple episodes of listing and de-listing, each instance of which shows up as a separate entry in the raw data. For example, a property could be listed for a duration of 6 months without selling, then pulled off the market for 1 month, and subsequently re-listed and sold 2 months later. In such a case, the raw data reports a failed listing that lasts for 6 months followed by a successful listing that takes 2 months. By contrast, the view in this paper is that the property has taken 8 months to sell. More generally, to capture the effective time that properties are listed before selling or being pulled off the market for good, the measure of time on the market in this paper strings together all failed listings that are separated by less than 3 months and adds them to the terminal listing that culminates either in a sale or more permanent removal. The vast majority of de-listings and re-listings occur within a 3 month horizon, which motivates this choice of threshold for distinguishing between strategic seller behavior and genuine instances of sellers removing their property from the market (e.g. to make home improvements or wait for a better selling environment).

Months of Supply Months of supply is the number of houses on the market in each county divided by the seasonally adjusted annualized sales rate in that county. Thus, whereas time on the market is a listing-level variable that can be aggregated manually to measure housing illiquidity in a broader geography, months of supply is intrinsically a market-level illiquidity measure.

Comparisons to Publicly Available Housing Data Presently, there is no publicly available county-level data that extends back to before the crisis began in 2006 to compare with the CoreLogic MLS data.[1] However, publicly available *national* data on house prices and illiquidity shown in figure 9 serves as a useful benchmark to assess the MLS data. In particular, real house prices in the MLS follow a similar though slightly exaggerated trajectory to the paths of the inflation-adjusted house price indices from the Federal Housing Finance Agency (FHFA) and Case-Shiller. Furthermore, months of supply in the MLS closely tracks its counterparts from the Census and National Association of Realtors (NAR) for new houses and existing houses, respectively.

A.1.2 Mortgage Default

The Equifax Credit Bureau Database (FRBNY Consumer Credit Panel) contains loan-level data on households' credit reports. The data is a representative 5% sample of individuals in the United States with a credit report and Social Security Number and is reported on a quarterly basis. To coincide with the other variables in the regressions, the data is aggregated to the county level. Important to the construction of the housing net worth

[1] For example, public county-level data from Zillow for days on the market begins in 2010.

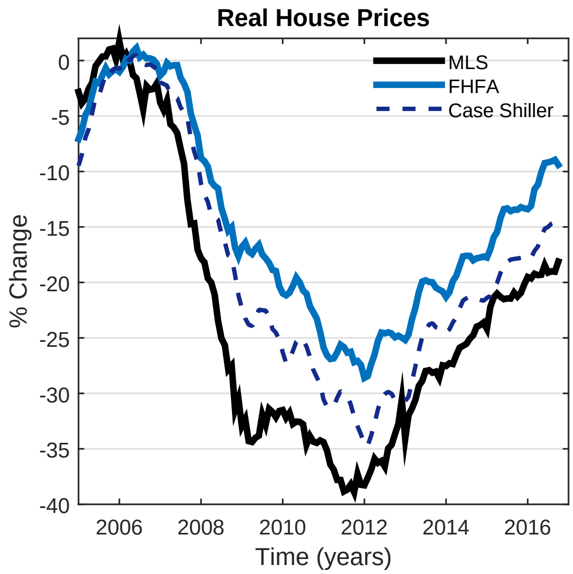

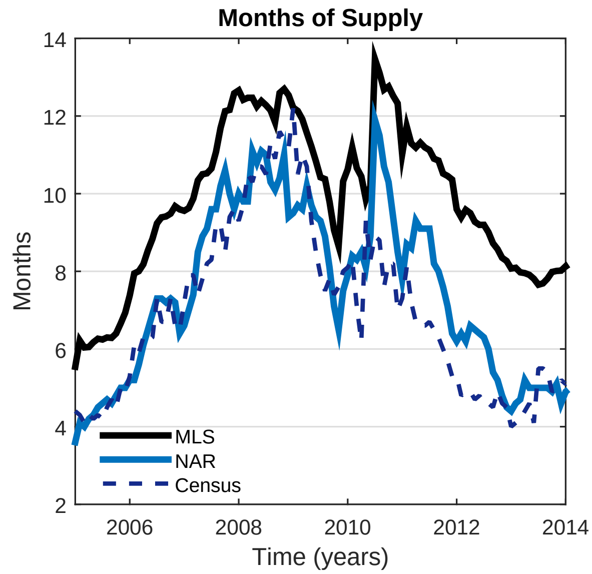

Figure 9: Comparison of national MLS data with publicly available sources. (Left) MLS average transaction price vs. FHFA purchase-only index vs. Case-Shiller index (all deflated by the PCE). (Right) MLS vs. National Association of Realtors (existing houses) vs. Census (new houses). Months of supply is the number of houses on the market divided by monthly sales.

shock discussed below, the data contain information on the size of outstanding balances for many types of household debt, especially mortgages. In addition, the data provide several measures of mortgage payment status, including whether a loan is current, past due, or severely derogatory and heading toward foreclosure, which is this paper's preferred measure of mortgage default.

A.1.3 Income

The IRS Statistics of Income (SOI) dataset provides zip code level information for selected income and tax items. The data are based on individual tax returns taken from forms 1040, 1040A, and 1040EZ filed with the IRS. If a taxpayer files returns for multiple years at any given time, only the most recent return is included. A zip-to-county crosswalk from HUD is then used to convert this data to the county level, thereby making it consistent with the other variables in this paper's analysis. From this data, the primary object of interest is adjusted gross income, but the non-wage component of income—which subsumes income from interest, dividends, and capital gains—is used to approximate county-level assets during the construction of the housing net worth shock variable described below.

A.1.4 Employment

The Quarterly Census of Employment and Wages (QCEW) from the BLS publishes a quarterly count of employment and wages reported by employers covering more than 95 percent of U.S. jobs at various degrees of geographic disaggregation down to the county level.

This paper relies on the QCEW for its measure of **total employment** in each county. As a supplement, the Census County Business Patterns (CBP) gives detailed industry-specific county-level employment data. This paper then constructs **nontradable employment** for each county by assigning industries according to their 4-digit NAICS classification using the criteria in Mian and Sufi (2014).

A.1.5 Housing Net Worth Shock

Constructing the housing net worth shock from Mian et al. (2013) and Mian and Sufi (2014) requires a pre-crisis measure of county-level net worth in 2006: $NW_{06} = A_{06} + H_{06} - D_{06}$, where NW is net worth, A is financial assets, H is housing wealth, and D is total household debt. No data exists for the county-level *stock* of assets, but it can be approximated as each county's share of total asset *income* (i.e. IRS non-wage income a^i_{06} for county i) multiplied by the Flow of Funds aggregate stock of financial assets: $A_{06} = \frac{a^i_{06}}{\sum_j a^j_{06}} A^{FOF}_{06}$. Housing wealth is calculated as the county-level 2006 average MLS price multiplied by the number of *owner-occupied* units in the county interpolated between the Census 2000 and 2010 values: $H_{06} = \sum_j p^j_{06} n^{j,own}_{06}$. To construct the measure of debt and correct for possible under-reporting in small counties, the cumulative debt balance *per borrower* in 2006 for each county is scaled up by the *total* number of households in the county interpolated from the Census 2000 and 2010 values: $D_{06} = \sum_j d^j_{06} n^{j,tot}_{06}$. Table 10 shows that the constructed aggregates from the merged sample mirror the direct Flow of Funds measures.[2]

Table 10: Aggregate Net Worth Ratios

Ratio	Merged Sample	Flow of Funds
A_{06}/NW_{06}	0.83	0.76
H_{06}/NW_{06}	0.32	0.34
D_{06}/NW_{06}	0.15	0.20

The housing net worth shock between 2006 and 2011 in county i is then defined as the percentage house price change multiplied by the housing net worth share, i.e. $\%\Delta\text{HNW}^i_{06,t} = \%\Delta\text{Prices}^i_{06,t} \times \frac{H^i_{06}}{NW^i_{06}}$.[3]

A.2 Descriptive Statistics and Figures

This section provides supplemental information to what section I covers regarding the behavior of key variables during the housing crash. In particular, table 11 provides summary statistics from the peak of the pre-crisis to its trough for housing, credit, and macroeconomic variables of interest. The data reveal both the aggregate severity of the crisis and the large degree of variation.

[2] Flow of Funds: (A_{06}) FL154090005Q. (H_{06}) LM155035015Q. (D_{06}) FL154190005Q. (NW_{06}) FL152090005Q. The cleaned merged data are winsorized above the 95th percentile of H^i_{06}/NW^i_{06} to correct for exaggerated 2006 housing values in the MLS relative to Census.

[3] The 2006–2009 time window in Mian et al. (2013) stops prior to the house price trough.

Table 11: Summary Statistics (Peak-Trough 2006–2011)

	Obs	Mean	St. Dev.	10th	Median	90th
House Prices (%Δ)	7,570	−36.64	20.71	−63.59	−35.73	−11.63
Months Supply (ΔMonths)	7,261	10.57	11.98	0.53	6.83	26.05
Time on Market (ΔDays)	7,269	53.71	38.33	10.42	50.78	101.37
Mortgage Default (Δpp)	7,047	5.32	4.99	0.90	3.84	12.08
Adjusted Gross Income (%Δ)	7,519	−6.52	11.40	−16.00	−6.26	1.43
Employment* (%Δ)	1,496	−5.92	4.54	−12.04	−6.07	−0.57
Nontradable Employment* (%Δ)	1,496	−5.26	7.45	−10.94	−6.50	2.96

*County level. All other statistics are at the zip-code level. Δ = change; %Δ = percent change. Sources: (Income) IRS Statistics of Income deflated by PCE. (Employment) BLS Quarterly Census of Employment and Wages. (Nontradable Employment) Census County Business Patterns. (House Prices, Months Supply, Time on Market) CoreLogic MLS. House prices deflated by PCE. (Mortgage Default) Equifax. MLS statistics are sales-weighted. All other statistics are population-weighted.

Figure 10 supplements the empirical evidence in section IV that supports the mechanisms described in the model whereby house prices and liquidity transmit to the rest of the macroeconomy. Specifically, the left column of panels is consistent with the findings in Mian et al. (2013) and Mian and Sufi (2014) that associate larger house price declines with worse macroeconomic outcomes. However, the remaining two columns show that prices provide only a limited view of the negative housing spillovers during the crisis. In particular, deteriorations in housing liquidity measured either by rising months of supply or time on the market are associated with significant declines in income and employment at the county level throughout the country.

Recall the heat maps from section I that show a clear geographic pattern between the decline in house prices, the fall in income, the spike in mortgage default, and the rise in housing *illiquidity* at the county level, using months of supply as the measure. The heat maps in figures 11 and 12 confirm that the same patterns arise when using time on the market as the measure of illiquidity instead. Furthermore, the map in figure 13 reveals a strong association between the decline in income and employment, suggesting that both are good proxies for the poor performance of the macroeconomy during the housing bust.

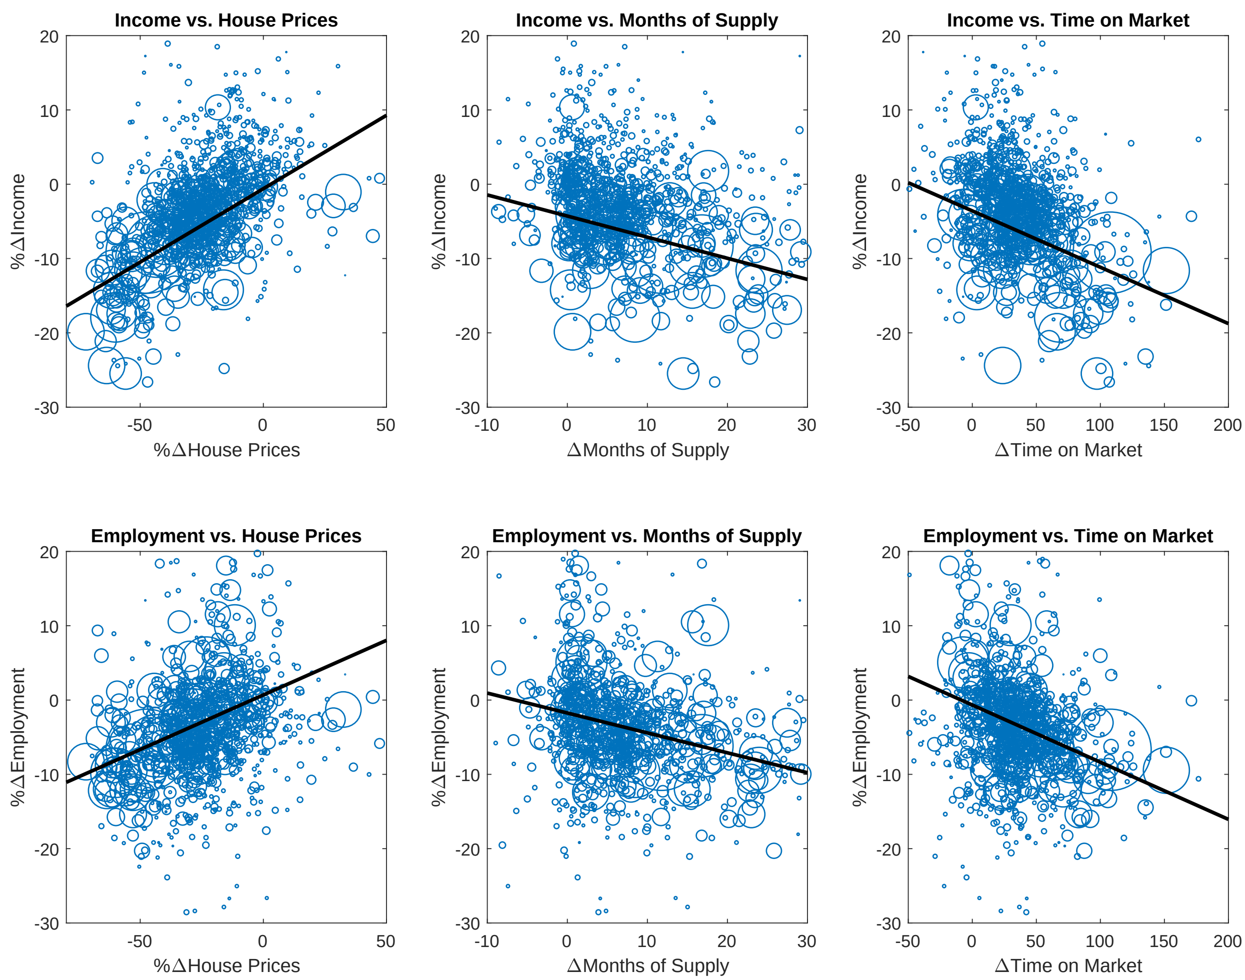

Figure 10: Housing, income, and employment from 2006–2011. Larger circles represent more populous counties. Sources: (Housing) CoreLogic MLS data. (Income) IRS Statistics of Income. (Employment) BLS Quarterly Census of Employment and Wages. Financial variables are deflated by the PCE.

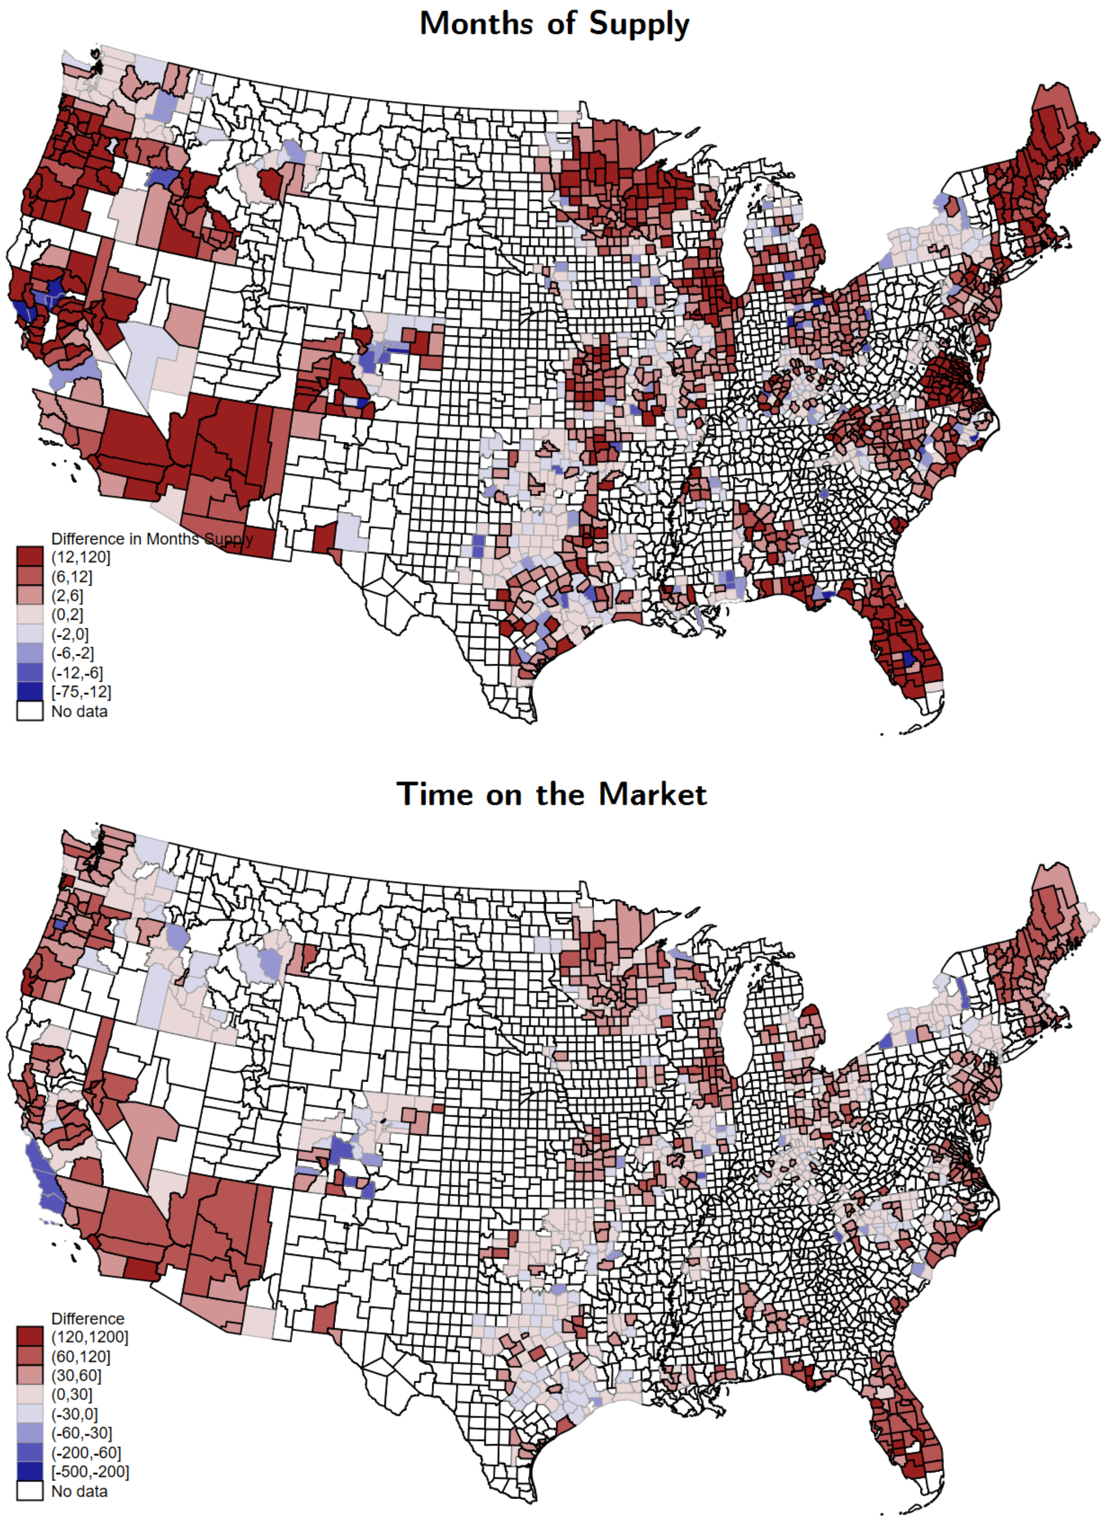

Figure 11: The *change* in months of supply and time on the market (measured in days) from 2005 to 2008. Source: CoreLogic MLS data.

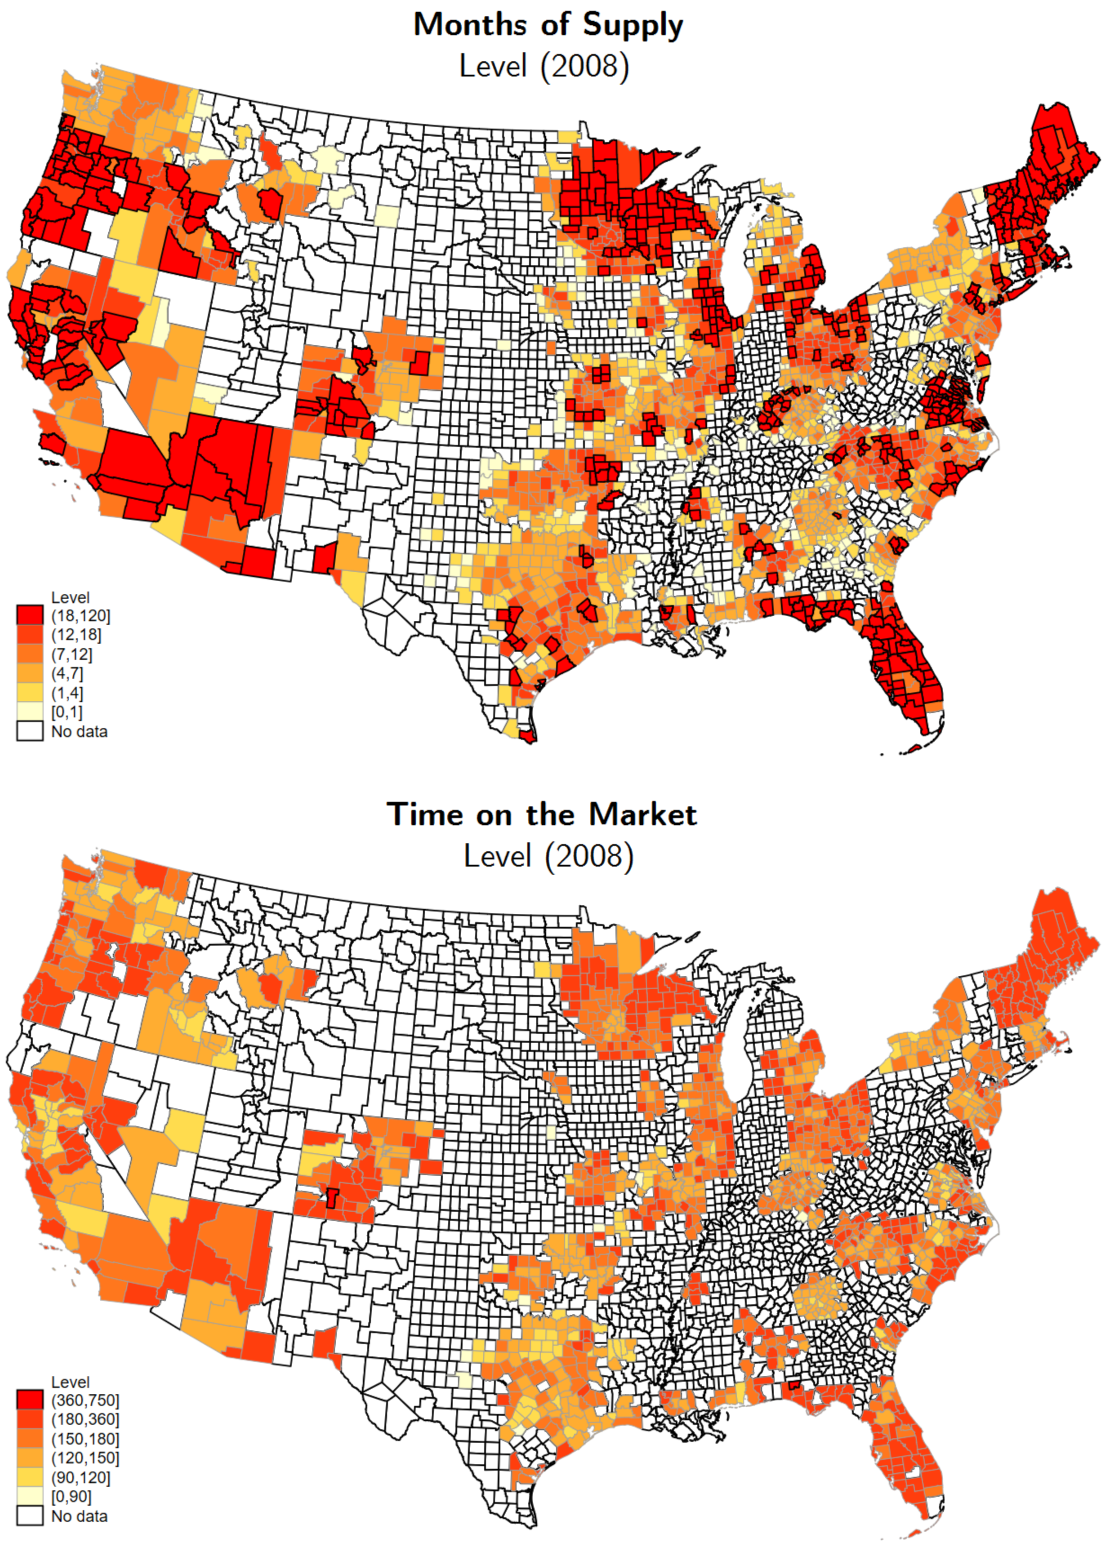

Figure 12: The *level* of months of supply and time on the market (measured in days) in 2008. Source: CoreLogic MLS data.

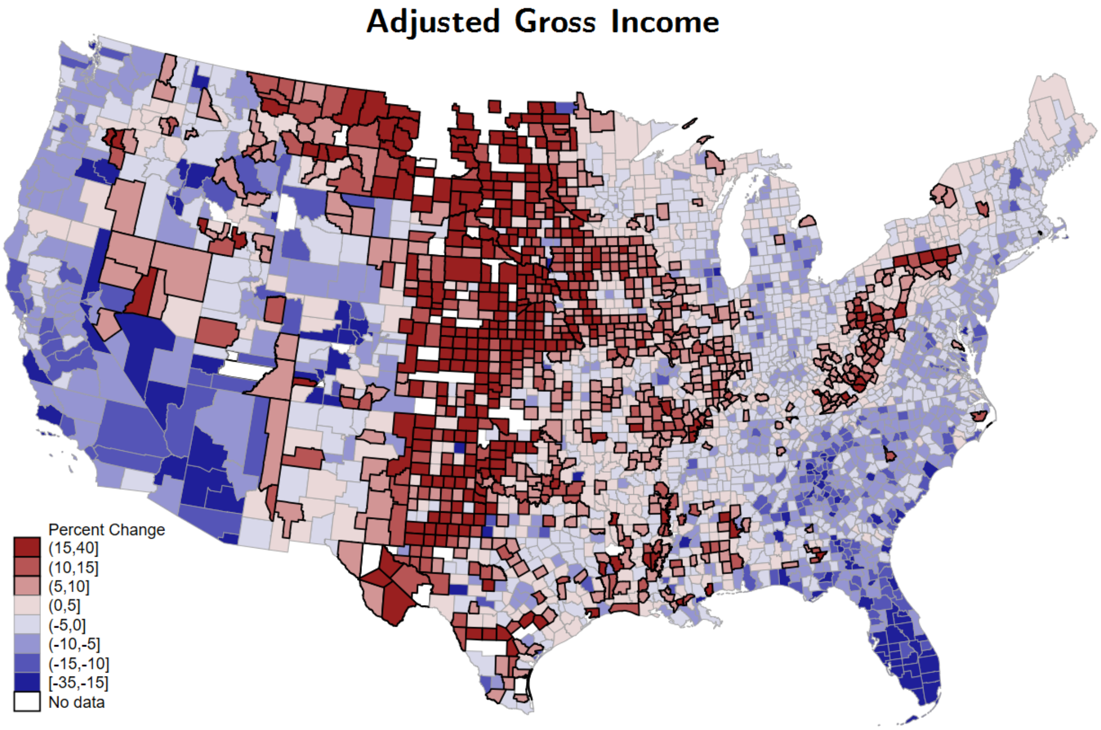

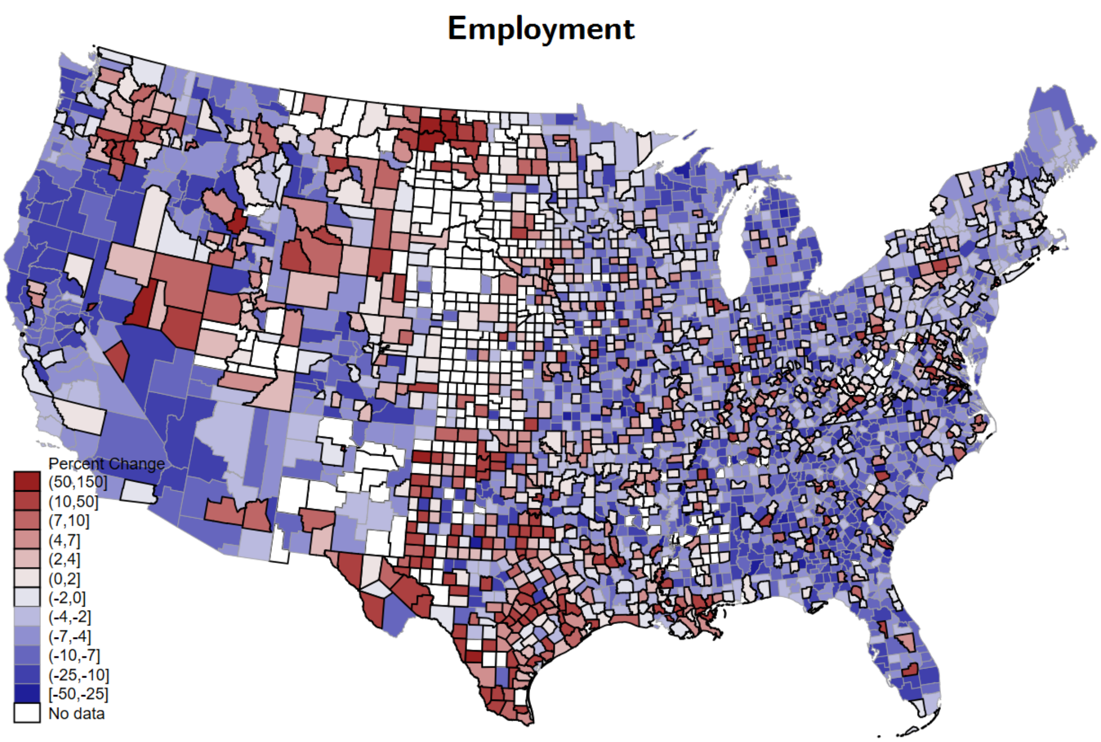

Figure 13: The percentage change in adjusted gross income and employment from 2006 to 2011. Sources: (Income) IRS Statistics of Income. (Employment) BLS Quarterly Census of Employment and Wages.

A.3 Regression Results

This section provides a detailed description of the regressions in section IV that analyze the transmission from housing to credit and the macroeconomy. Table 12 gives the full regression results corresponding to specification 10, which isolates the impact that changes to house prices and liquidity have on mortgage default at the county level. The positive and statistically significant coefficients for both measures of illiquidity indicate that selling delays are an important contributing factor to default. Given that average county-level months of supply rose by 10.57, the coefficient of 0.125 in table 12 implies an increase in mortgage default of 1.3 percentage points from illiquidity alone.

Table 12: Default and Liquidity

	(1) ΔDefault	(2) ΔDefault
$\%\Delta\text{Prices}\times\frac{H_{06}}{NW_{06}}$	-0.131^{***}	-0.140^{***}
	(0.005)	(0.005)
ΔMonths Supply	0.125^{***}	
	(0.009)	
ΔTime on Market		0.027^{***}
		(0.002)
Constant	0.891^{***}	0.872^{***}
	(0.105)	(0.110)
N	1021	935
R^2	0.540	0.545

Regressions are weighted by county population. Default and prices are from 2006 – 2010; months supply and time on the market are from 2005 – 2008. ***Significant at the 1% level.

Table 13 provides estimates for regression specification 13, which measures the macro spillovers—specifically, for adjusted gross income and nontradable employment at the county level—associated with declining house prices and rising illiquidity. Column 1 closely mirrors the main specification in Mian et al. (2013), except that this paper uses a wider time horizon from 2006 to 2011 that covers the entire housing bust instead of only the 2006–2009 period they use. Furthermore, their dependent variable is proprietary Mastercard consumption data instead of the publicly available AGI from the IRS. Even so, the price coefficient of 0.237 is remarkably similar to the 0.341 they obtain for the response of nondurable consumption to house price changes.

Revealing the novel importance of illiquidity, incorporating the lagged change in months of supply from 2005 to 2008 adds a similarly large but negative coefficient while causing the regression to explain a larger share of the AGI variation. For perspective, this coefficient yields a predicted 2 percentage point decline in AGI based on the observed average increase in months of supply during the crisis. Replacing months of supply with time on the market

Table 13: Elasticity to Changes in House Prices and Liquidity

	(1) %ΔAGI	(2) %ΔAGI	(3) %ΔAGI	(4) %ΔE_{NT}	(5) %ΔE_{NT}	(6) %ΔE_{NT}
%ΔPrices$\times\frac{H_{06}}{NW_{06}}$	0.237***	0.202***	0.229***	0.118***	0.091***	0.090***
	(0.011)	(0.012)	(0.012)	(0.023)	(0.025)	(0.026)
ΔMonths Supply		−0.188***			−0.143***	
		(0.022)			(0.047)	
ΔTime on Market			−0.029***			−0.035***
			(0.005)			(0.010)
Constant	−1.803***	−0.780***	−0.859***	−0.771	0.010	0.198
	(0.241)	(0.262)	(0.275)	(0.494)	(0.553)	(0.575)
N	1023	1023	934	1023	1023	934
R^2	0.304	0.350	0.348	0.025	0.034	0.036

AGI is adjusted gross income; E_{NT} is nontradable employment using the 4-digit NAICS industry classification from Mian and Sufi (2014). Regressions are weighted by county population. Dependent variables and prices are from 2006 – 2011; months supply and time on the market are from 2005 – 2008. ***Significant at the 1% level.

delivers a smaller coefficient, but critically, time on the market is measured in *days* and therefore experiences a much larger increase than does months of supply. A similar back of the envelope calculation based on the average rise in time on the market in table 11 predicts a 1.6 percentage point drop in AGI.

Switching from AGI to nontradable employment as the macro variable of interest as in Mian and Sufi (2014) changes the coefficients but not the underlying message. In particular, the price coefficient of 0.118 in the regression without illiquidity is close to the 0.190 estimate they obtain, even though the time horizon in this paper is 2006–2011 instead of 2007–2009. As in the case of AGI, introducing either months of supply or time on the market reveals an economically meaningful and statistically significant effect of illiquidity. Specifically, the regression predicts a 1.9 percentage point decline in nontradable employment based on the observed rise in time on the market.

Lastly, table 14 shows the results for specification 14 in section IV.C that measures the marginal impact of house price and liquidity changes on county-level income. If AGI is viewed as a proxy for consumption, the house price coefficient of 0.03 in the regression without illiquidity is in line with the 0.047 and 0.054 empirical estimates for the marginal propensity to consume out of house price changes in Aladangady (2017) and Mian et al. (2013), respectively. Put another way, the regression predicts a $30 fall in AGI in response to a $1,000 decline in house prices. Showing again the macroeconomic importance of illiquidity, column 2 reveals that each additional month of supply is associated with a $89 drop in AGI, and each one *day* increase in time on the market predicts a $31 fall in AGI, which corresponds to a $1,700 total effect based on the observed rise in time on the market in the data.

Table 14: Marginal Response of Income to Prices and Liquidity

	(1)	(2)	(3)
	ΔAGI	ΔAGI	ΔAGI
ΔPrices	0.030***	0.027***	0.033***
ΔMonths Supply		−89.264***	
		(17.412)	
ΔTime on Market			−31.282***
			(3.520)
Constant	−1262.632***	−766.795***	−3.910
	(165.868)	(190.267)	(202.868)
N	1023	1023	934
R^2	0.343	0.359	0.412

Regressions are weighted by county population. AGI and prices are from 2006 – 2011; months supply and time on the market are from 2005 – 2008. ***Significant at the 1% level.

B Supplementary Tables and Figures

This appendix provides companion figures and tables regarding model fit, transmission mechanisms, and alternative explanations for the housing crisis.

B.1 Additional Dimensions of Model Fit

Section IV discusses at length the importance of balance sheet depth during the crisis. In particular, illiquidity-induced debt overhang for households with high initial leverage causes a pronounced increase in foreclosure risk, a contraction in the supply of credit, and a larger drop in house prices and consumption. Thus, it is important that the pre-crisis parametrization properly capture the right tail of the mortgage leverage distribution that acts as the source for many of the model nonlinearities. Figure 14 shows that, indeed, the parametrized model does remarkably well at matching this important segment of the leverage distribution from the 2007 Survey of Consumer Finances while also replicating the targeted aggregate portfolio statistics in table 1.

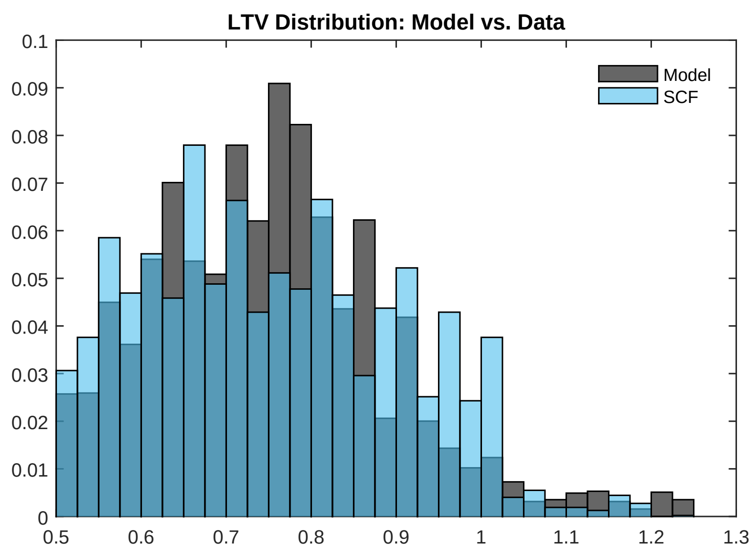

Figure 14: Pre-bust LTV distribution in the model and in the 2007 SCF.

B.1.1 Selected Shocks in the Model and Data

Figure 15 depicts the skewness and interest rate shocks from the model next to their empirical counterparts. Although households in the model inelastically supply *raw* labor, their individual effective labor supply $e_t \cdot z_t$ is stochastic. As explained in section III, the skewness shocks are implemented as temporary changes to the transition matrix π_z. First, downside risk in π_z is increased for 3 years to match the decline in aggregate employment from the BLS over that same horizon. Then, a reversal in downside risk is implemented to match the pace of the employment recovery. Finally, π_z is reset to its pre-crisis state. The left panel shows the behavior of aggregate labor in the model and data.

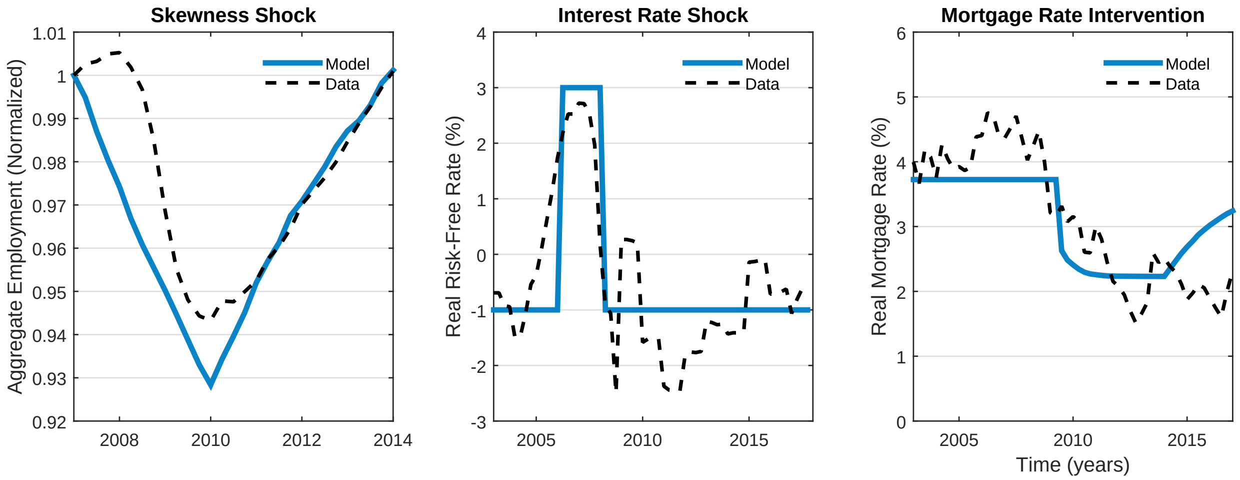

Figure 15: Select shocks in the model and data. (Left) Aggregate employment. (Middle) Real risk-free rate. (Right) Real mortgage rate. The decline in panel 3 corresponds to nominal mortgage rates falling from around 5.5% to 4%.

The middle panel shows the dynamics of the risk-free rate i_t. In the model, i_t is exogenously increased (because of the open economy assumption) for two years to approximate the tightening in monetary policy during 2006 and 2007 and then lowered again to reflect the Federal Reserve's reversal as the economy began to collapse. Despite this temporary increase in i_t, long-term mortgage rates r_t—which are set according to equation 30—do *not* rise, as seen in the right panel.[4] Instead, they initially remain flat and subsequently fall after policy interventions were enacted to reduce long-term borrowing costs. In the model, lower mortgage rates are instituted via reduced servicing costs ϕ_t.

B.1.2 Further Cross-Validation

The open economy assumption in the model implies that net financial flows—measured as the gap between total financial (i.e. liquid) assets and mortgage debt—are typically non-zero. Although the model parametrization targets several aggregate and cross-sectional portfolio statistics from the 2007 Survey of Consumer Finances, it does not target the pre-crisis size of this net financial position. Nevertheless, the initial gap between the two curves in figure 16 shows that the model closely mirrors the data.

During the crisis, net financial flows increase (become less negative) in the data and in the model, as seen both in the left panel for levels and in the right panel, which measures the *change* in these flows. These dynamics emerge from a combination of higher savings and lower debt. The temporary rise in the risk-free rate increases the return to saving, and heightened downside risk strengthens households' precautionary motive to save. The reduction in mortgage debt comes about because of higher mortgage outflows from foreclosures and reduced inflows from changes both on the demand and supply side. In particular, the increased riskiness of housing reduces the appeal of homeownership, and the contraction in credit—which occurs both because of the exogenous credit limit shocks and

[4] Recall that endogenous default risk is incorporated in mortgage *prices* q_t at origination.

the endogenous rise in default spreads—reduces the ability of buyers to finance purchases through borrowing. In addition to matching these net financial flows, the model generates aggregate foreclosure losses amounting to 12.6% of annual goods output, which is in line with the magnitude of losses reported by the IMF.[5]

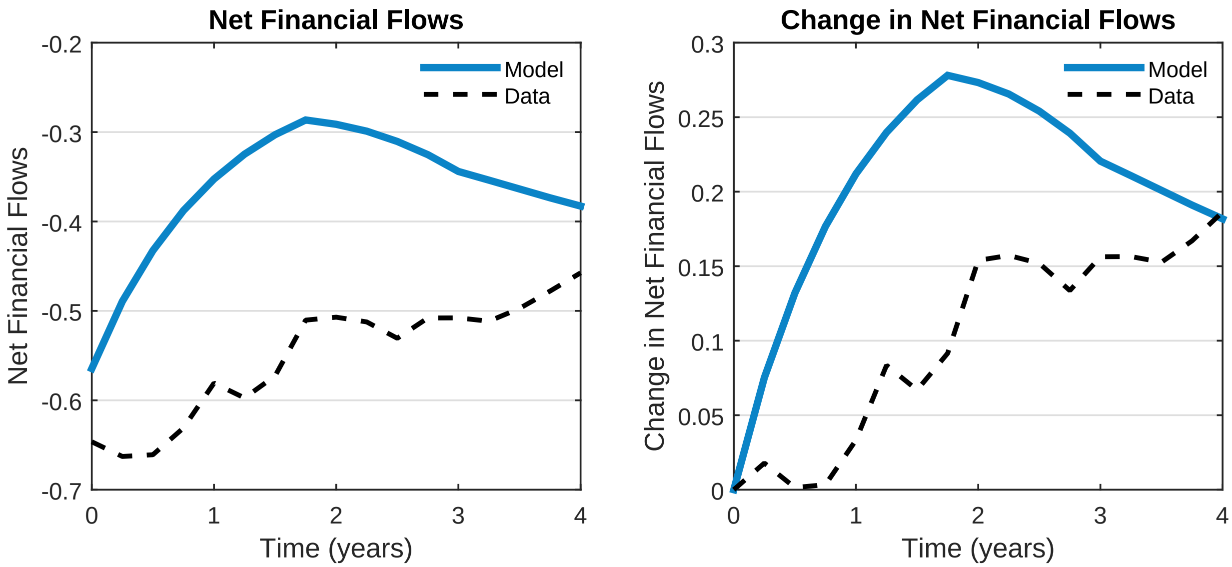

Figure 16: (Left) Net financial flows in the model and data. (Right) Dynamics after subtracting the pre-crisis gap between model and data. Net flows in the data are from the Flow of Funds and are defined as total financial deposits (FL154000025Q) minus the sum of mortgage debt (FL153165105Q) and consumer credit (FL153166000Q). This measure is deflated by the PCE and normalized by total annual earnings from the Bureau of Economic Analysis.

Figure 17 complements tables 3, 6, and 7 in section IV that show the model's success in matching the aggregate and cross-sectional behavior of consumption during the crisis. The top left panel shows that aggregate consumption in the model and data falls by approximately 10%, while the remaining three panels reveal the significant degree of heterogeneity in consumption dynamics by net worth, tenure status, and degree of leverage. In particular, heavily-indebted owners experience the largest decline in consumption followed by less leveraged owners and, finally, renters. These patterns reinforce the discussion in section IV.C about the importance of balance sheet depth.

[5]See https://www.nytimes.com/2009/04/22/business/global/22fund.html.

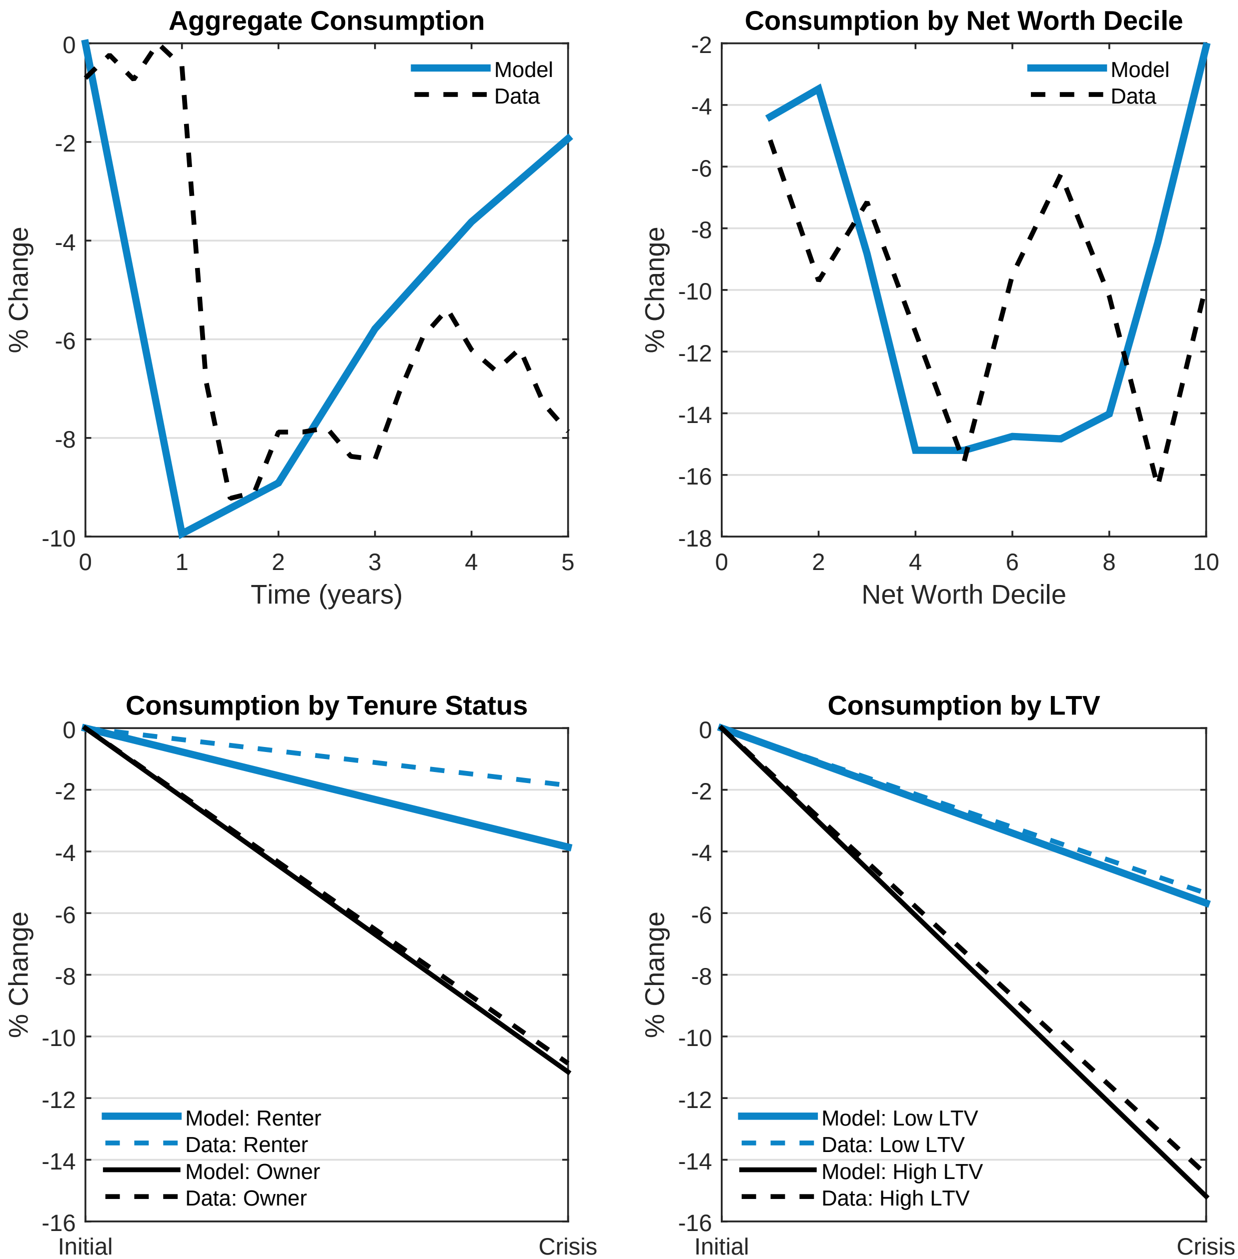

Figure 17: Sources: (Aggregate) BEA nondurable consumption deflated by the PCE (detrended from extrapolated 1991 – 2000 linear trend). (Disaggregated) PSID using the sampling criteria of Arellano et al. (2017). Low loan-to-value (LTV) is below 0.3; high LTV is above 0.8.

B.2 Housing Behavior in the Cross Section

This section adds to section IV.A by analyzing the cross-sectional behavior of foreclosures, homeownership, and the liquidity-adjusted double trigger.

B.2.1 Default, Tenure Flows, and the "New Narrative"

Two of the most salient features of the 2006–2011 housing crisis were the wave of mortgage defaults and the associated persistent homeownership decline. Challenging the view that these maladies were concentrated in the subprime market, several recent empirical papers have used administrative credit panel data to uncover evidence pointing to the broad-based nature of the foreclosure crisis.[6] The cross-sectional model results in the top row of figure 18 are consistent with this new narrative. In particular, foreclosures spike both in the bottom and middle segments of the market, with only high-income borrowers at the top end of the market emerging relatively unscathed. Turning to the top right panel, *relative* gross exits into renter status are initially most pronounced among owners of medium-sized houses whose deeper balance sheets and higher leverage make them more financially exposed to shocks. Over time, the collateral damage from tighter credit limits and increased downside earnings risk also takes its toll on owners of small houses as they transition out of owning and into renting. In *absolute* terms, the equilibrium house size distribution *conditional on ownership* remains relatively stable, although the staggered timing of exits from owning to renting—first for medium houses, and then for small houses—is evident in the bottom left panel. Lastly, the bottom right panel shows the decline in the aggregate housing stock as depressed construction fails to keep up with depreciation.

B.2.2 The Liquidity-Adjusted Double Trigger

As a supplement to section IV.B, figures 19 and 20 provide a visual depiction of mortgage default with and without endogenous housing liquidity. The height of the "mountain peak" at each point corresponds to the measure of homeowners with that combination of mortgage leverage and cash at hand at the time of the house price trough, while the shading indicates each household's foreclosure propensity, with brighter colors representing higher default risk. In the top panel of figure 19, low-income borrowers with negative equity (leverage above 1) default with near certainty, which is consistent with the standard double trigger. However, some borrowers with modest amounts of *positive* equity but little cash at hand *also* default with strictly positive probability because of selling delays from debt overhang, which is indicative of the **liquidity-adjusted double trigger**. By contrast, the bottom panel shows the stark "bang-bang" nature of default in the standard double trigger: for given cash at hand, the probability of foreclosure immediately jumps from 0 to 1 upon passing some leverage threshold.

The same differences between the standard and liquidity-adjusted double triggers appear for middle-income homeowners in figure 20, which also sheds light on the broad-based nature of the foreclosure crisis. Even though middle-income borrowers have lower default

[6] See, for example, Adelino, Schoar and Severino (2016), Foote, Loewenstein and Willen (2016), and Albanesi, DeGiorgi and Nosal (2017).

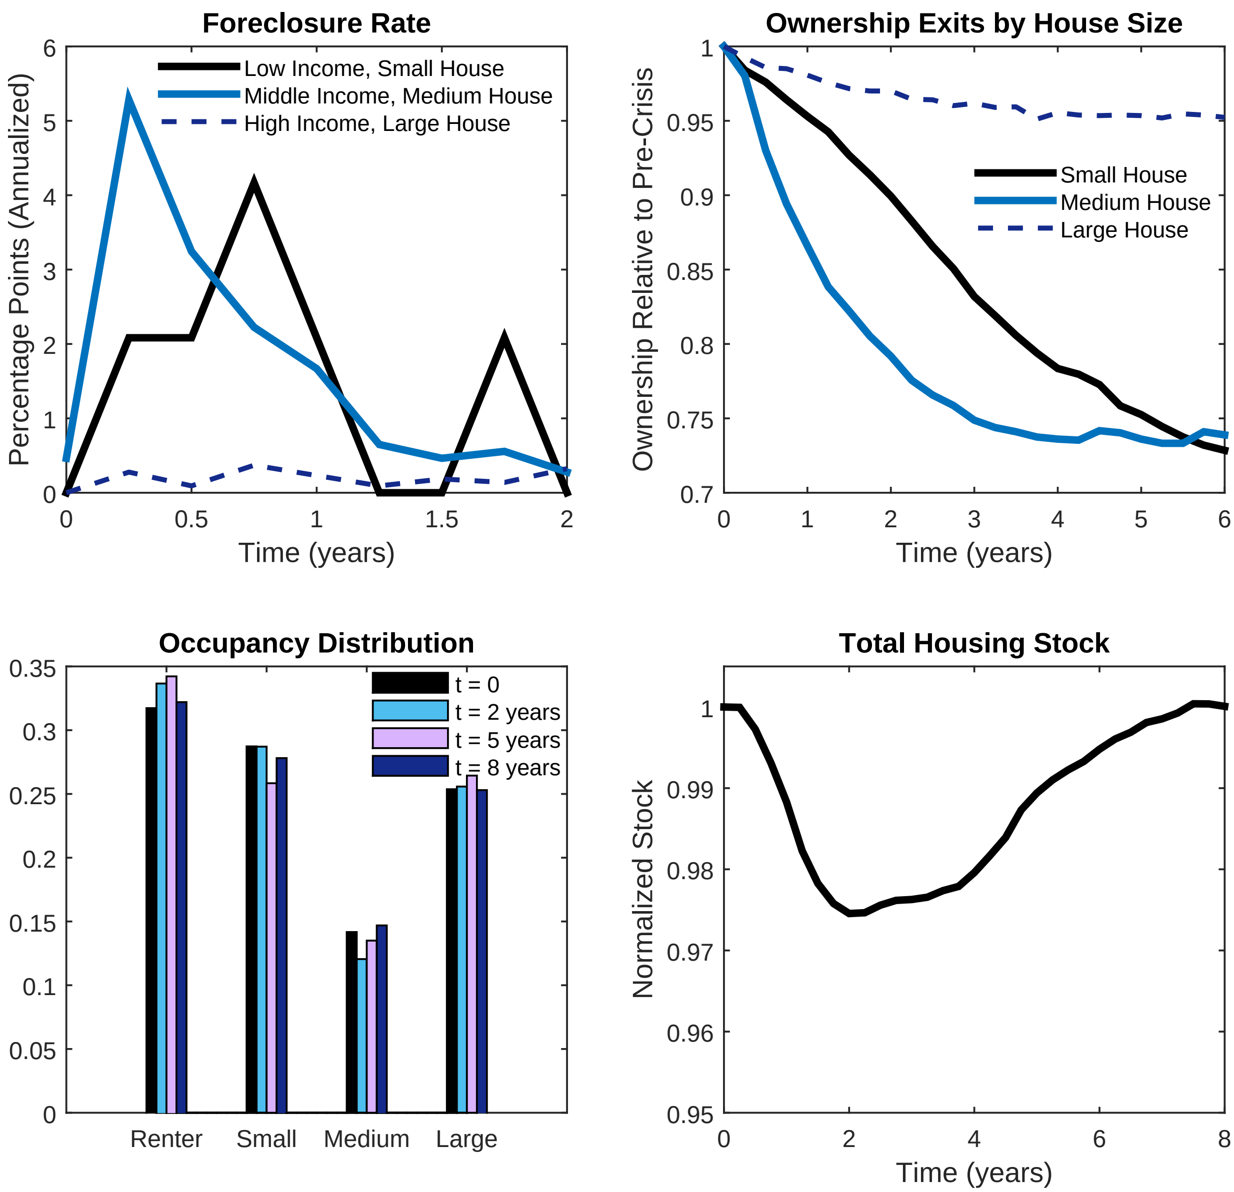

Figure 18: (Top Left) Foreclosure rate by pre-crisis income state and house size. (Top Right) Ownership rate by pre-crisis house size. (Bottom Left) Distribution of occupancy across renter status and house sizes. (Bottom Right) The dynamics of the aggregate housing stock, $H_t = (1 - \delta_h)H_{t-1} + Y_{ht}$.

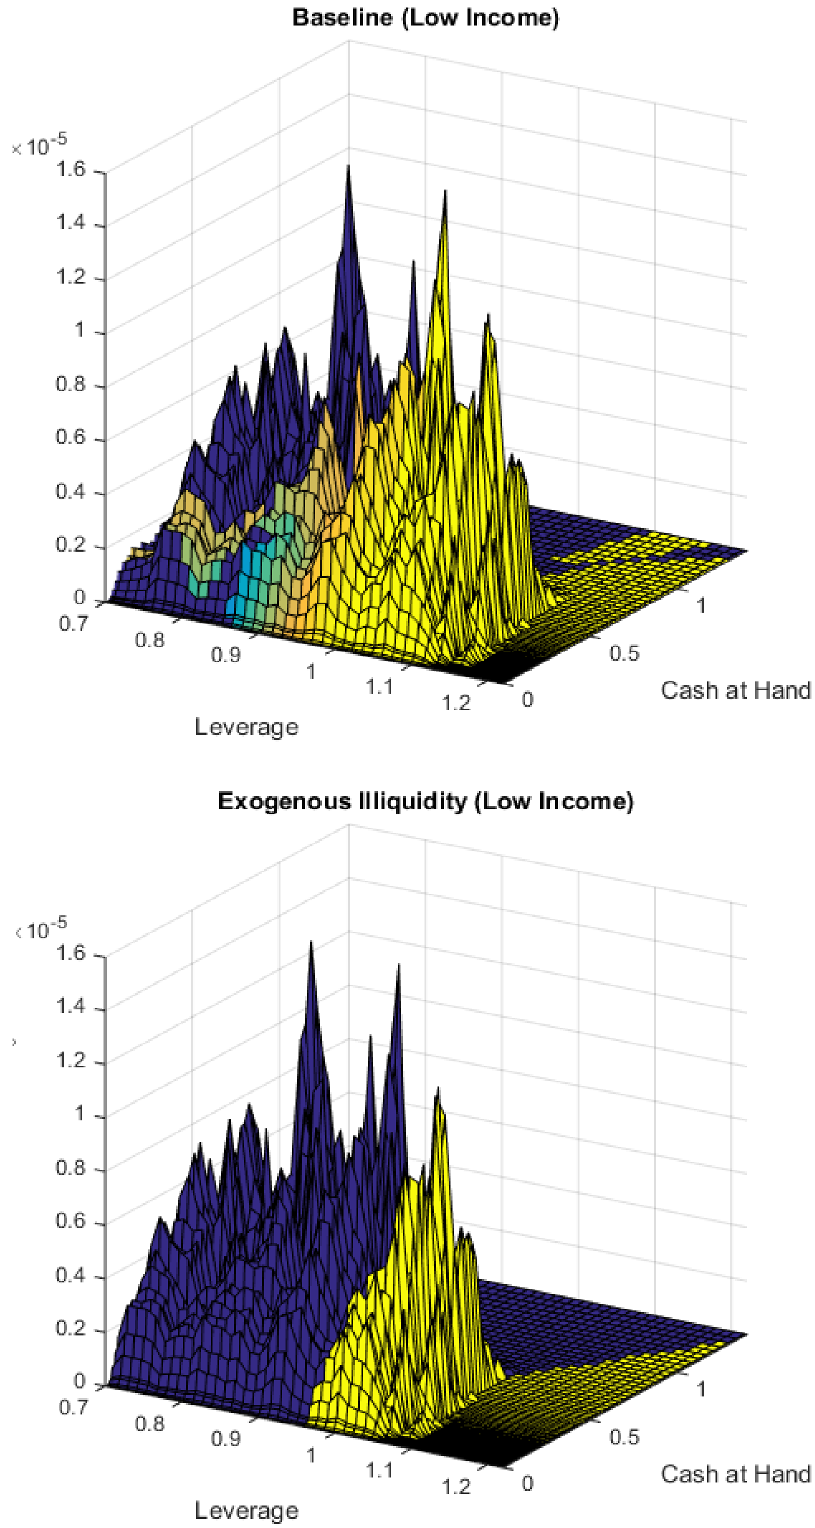

Figure 19: Distribution of low-income households in the bust with lighter shading indicating higher default probabilities. Foreclosures at lower LTV values in the baseline are driven by illiquidity-induced failures to sell.

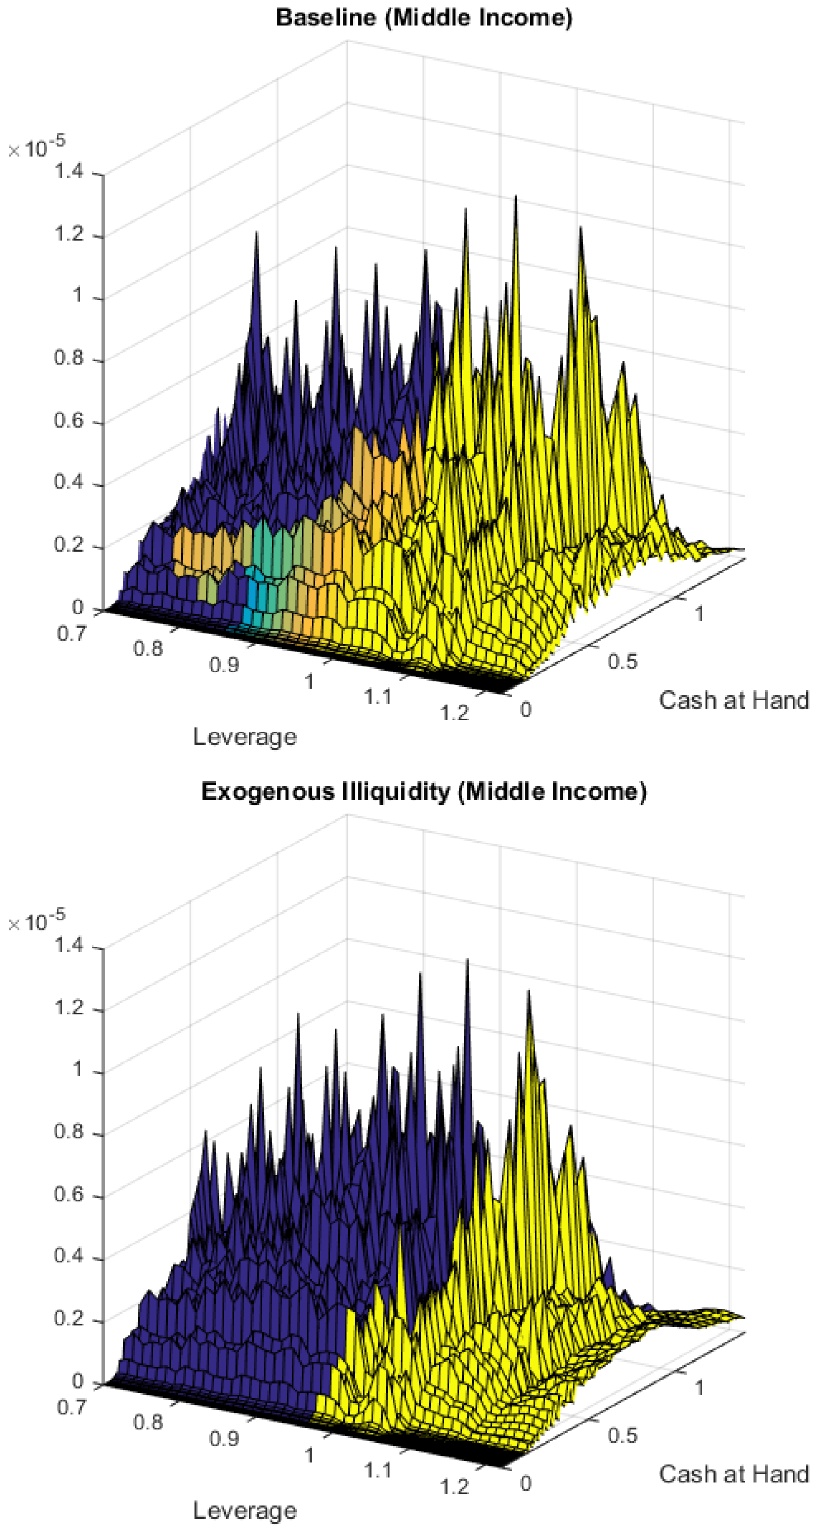

Figure 20: Distribution of middle-income households in the bust with lighter shading indicating higher default probabilities. Foreclosures at lower LTV values in the baseline are driven by illiquidity-induced failures to sell.

propensities than *otherwise identical* low-income borrowers with the same leverage and cash at hand, there is a *greater mass* of highly leveraged middle-income borrowers because their lower foreclosure risk grants them better access to credit at a reduced premium. Thus, middle-income borrowers contribute just as much to the total foreclosure rate during crises because they disproportionately inhabit a risky portion of the state space that exposes them to deteriorating housing market conditions.

B.2.3 Construction, Reshuffling, and the Occupancy Distribution

As in most workhorse macro-housing models, the construction of new housing is akin to the production of new capital–that is, housing is built in continuous and divisible units. It is only when households buy and sell houses that they are restricted to transacting indivisible house sizes from a discrete set. The conventional approach used in frictionless models assumes that the Walrasian auctioneer can costlessly reshuffle the housing stock across the discrete house sizes and clears the market at a uniform per-unit housing price. In this model, the real estate brokers provide a similar function, which leads to analogous (though not identical) equilibrium conditions. Proceeding in this manner gives rise to one equilibrium price (or price index) p_t as opposed to a vector of equilibrium prices $\{p_t(h)\}_h$ corresponding to each house size. To make the comparison between the Walrasian auctioneer and real estate brokers more evident, the Walrasian equilibrium condition (without foreclosures, for simplicity) is

$$\underbrace{\int h_t^* \mathbf{1}_{[buy_t^*]} d\Phi_t^{rent}}_{\mathcal{D}_t(p_t)} = \underbrace{Y_{ht}(p_t) + \int h \mathbf{1}_{[sell_t^*]} d\Phi_t^{own}}_{\mathcal{S}_t(p_t)},$$

where the indicator function on the left is the decision of whether to buy a house or not, and h_t^* is the choice of house size (with the dependence of these policy functions on p_t and state variables suppressed here). On the right-hand side, Y_{ht} is the construction of new housing given by the builder's standard first order conditions described in section C, and the second term is the total volume of housing sold by owners aggregated across house sizes, where the indicator function is the binary decision of whether to sell or not.

With directed search and brokers, the analogous equilibrium condition is

$$\int h_t^* \eta^{buy}(p_t^{bid*}, h_t^*; p_t) d\Phi_t^{rent} = Y_{ht}(p_t) + \int h \eta^{sell}(p_t^{list*}, h; p_t) d\Phi_t^{own}.$$

The main difference between these two equations is that, with search, only successful transactions (i.e. not failed searches) appear on either side. In the Walrasian model, all it takes is for a buyer or seller to decide they want to transact, which flips the indicator function from 0 to 1. By contrast, with search frictions, sellers (buyers) of house h choose a list price p_t^{list*} (bid price p_t^{bid*}) and succeed with probability $\eta^{sell}(p_t^{list*}, h; p_t)$ ($\eta^{buy}(p_t^{bid*}, h; p_t)$).

Although the equilibrium conditions with and without search frictions resemble each other, the baseline model actually does better at avoiding reshuffling between house sizes. Figure 21 provides a direct comparison of the amount of reshuffling in the baseline and Walrasian models. In the baseline model, there are only slight changes in the distribution of house sizes over time, whereas the Walrasian model shows larger amounts of reshuffling.

Specifically, in the Walrasian model, there is a significant decline in the number of small houses between $t = 2$ and $t = 5$ accompanied by a spike in the number of medium-sized houses during this period as buyers take advantage of cheap prices. By contrast, the increased difficulty of selling in the model with search-induced endogenous liquidity significantly attenuates this opportunistic upgrading during bad times.

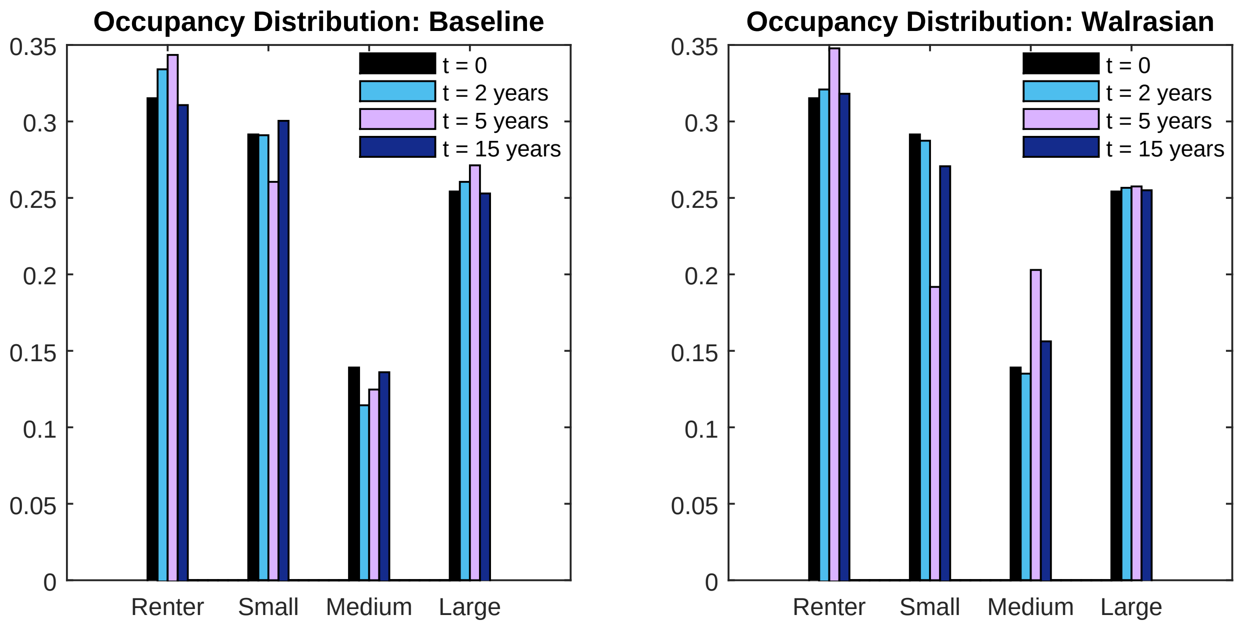

Figure 21: Dynamics of the occupancy distribution across house sizes in the baseline (left) and Walrasian (right) economies.

Table 15: Quantifying the Drivers of the Housing Bust

	Baseline	Exclude*	Alone**	Impact Bounds
Skewness Shock				
ΔHouse Prices	−23.4%	−14.8%	−11.6%	[−11.6%,−8.6%]
ΔOwnership	−2.8pp	+1.2pp	−3.1pp	[−4.0pp,−3.1pp]
ΔMonths Supply	+6.5m	+3.0m	+1.3m	[+1.3m,+3.5m]
ΔForeclosures	+5.1pp	+1.1pp	+0.2pp	[+0.2pp,+4.0pp]
ΔConsumption	−9.9%	−6.3%	−2.8%	[−3.6%,−2.8%]
Credit Shock				
ΔHouse Prices	−23.4%	−19.1%	−5.6%	[−5.6%,−4.3%]
ΔOwnership	−2.8pp	−3.0pp	+0.9pp	[+0.2pp,+0.9pp]
ΔMonths Supply	+6.5m	+3.5m	+0.3m	[+0.3m,+3.0m]
ΔForeclosures	+5.1pp	+2.3pp	−0.2pp	[−0.2pp,+2.8pp]
ΔConsumption	−9.9%	−7.0%	−2.2%	[−2.9%,−2.2%]
Productivity Shock				
ΔHouse Prices	−23.4%	−21.6%	−1.9%	[−1.9%,−1.8%]
ΔOwnership	−2.8pp	−2.9pp	+0.7pp	[+0.1pp,+0.7pp]
ΔMonths Supply	+6.5m	+5.5m	+0.5m	[+0.5m,+1.0m]
ΔForeclosures	+5.1pp	+3.6pp	−0.4pp	[−0.4pp,+1.5pp]
ΔConsumption	−9.9%	−8.0%	−1.0%	[−1.9%,−1.0%]
Interest Rate Shock				
ΔHouse Prices	−23.4%	−20.2%	−3.7%	[−3.7%,−3.2%]
ΔOwnership	−2.8pp	−2.9pp	+0.5pp	[+0.1pp,+0.5pp]
ΔMonths Supply	+6.5m	+4.8m	+0.5m	[+0.5m,+1.7m]
ΔForeclosures	+5.1pp	+4.4pp	−0.4pp	[−0.4pp,+0.7pp]
ΔConsumption	−9.9%	−8.7%	−2.0%	[−2.0%,−1.2%]

*The shock's effect in this case is the difference between the "baseline" and "exclude" columns. **The other shocks are removed.

B.3 Decomposing the Housing Bust

Section IV.A focuses on the role of worse earnings skewness and tighter credit limits as drivers of the housing crash. This section goes further by fully decomposing—visually and quantitatively—the contributions of all the productivity, interest rate, skewness, and credit limit shocks. Figure 22 shows the marginal impact of each shock. The curves prefaced with "no" remove the shock, and the "only" curves remove the other three shocks instead. The skewness and credit shocks have the largest impact on house prices, consumption, and foreclosures. Critically, the skewness shock is necessary to explain the decline in homeownership. Table 15 quantifies these effects.

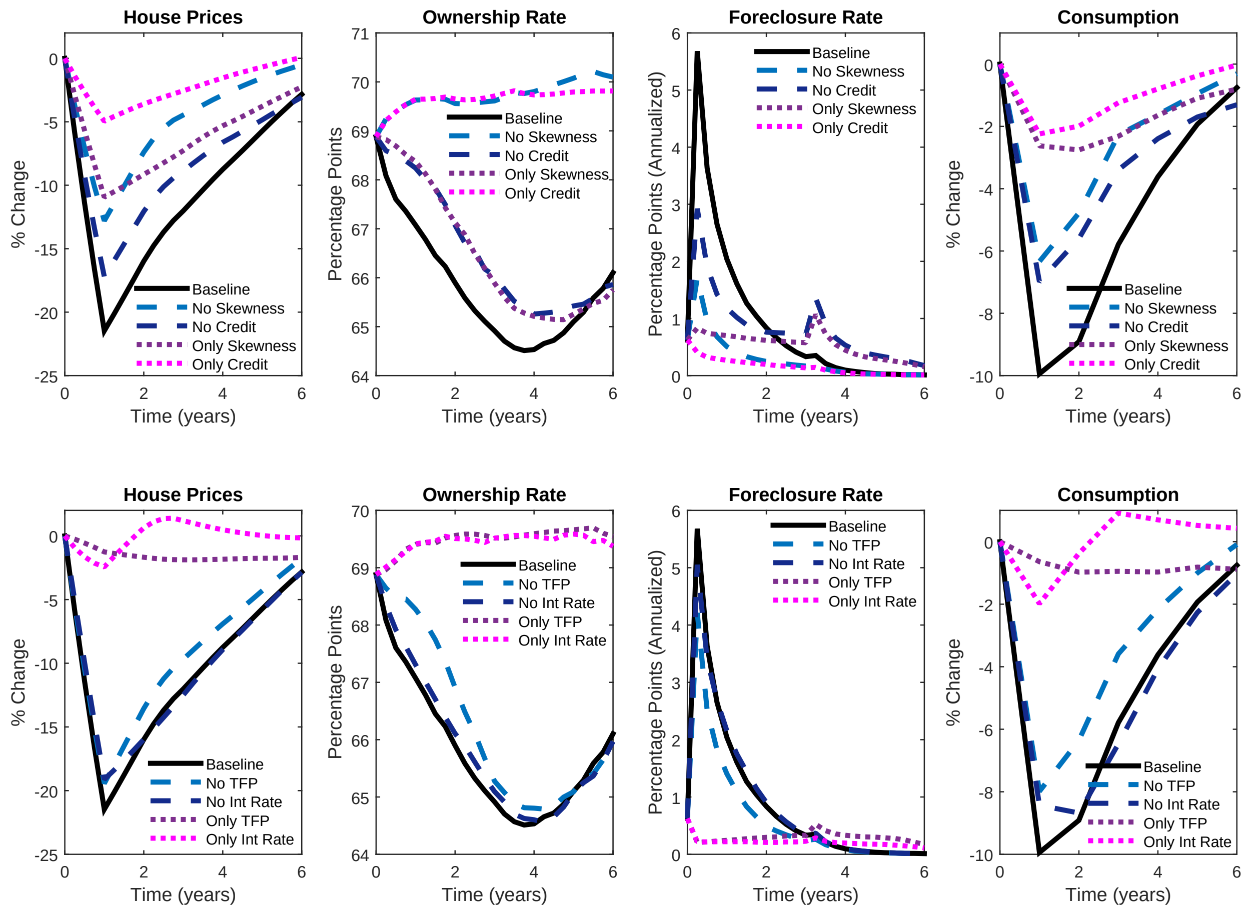

Figure 22: (Top) Decomposing the skewness and credit shocks. (Bottom) Decomposing the productivity and interest rate shocks. Plots with "No" remove the shock; plots with "Only" remove the *other* shocks.

B.3.1 Skewness Shocks: Realizations vs. Uncertainty

Figure 23 decomposes the effect of skewness shocks into earnings realizations vs. higher downside uncertainty. To isolate the effect of worse realizations, the model is simulated under the assumption that households do not perceive the change in their transition matrix π_z. The second simulation flips this scenario by removing the skewness shocks without the households' knowledge, leading them to false believe downside earnings risk is higher. Figure 23 also shows the baseline model and the version without skewness shocks entirely. Consistent with recent work by Berger, Dew-Becker and Giglio (2019), worse realizations have the largest negative impact, although uncertainty noticeably amplifies the response of sales, foreclosures, and consumption while also depressing homeownership.

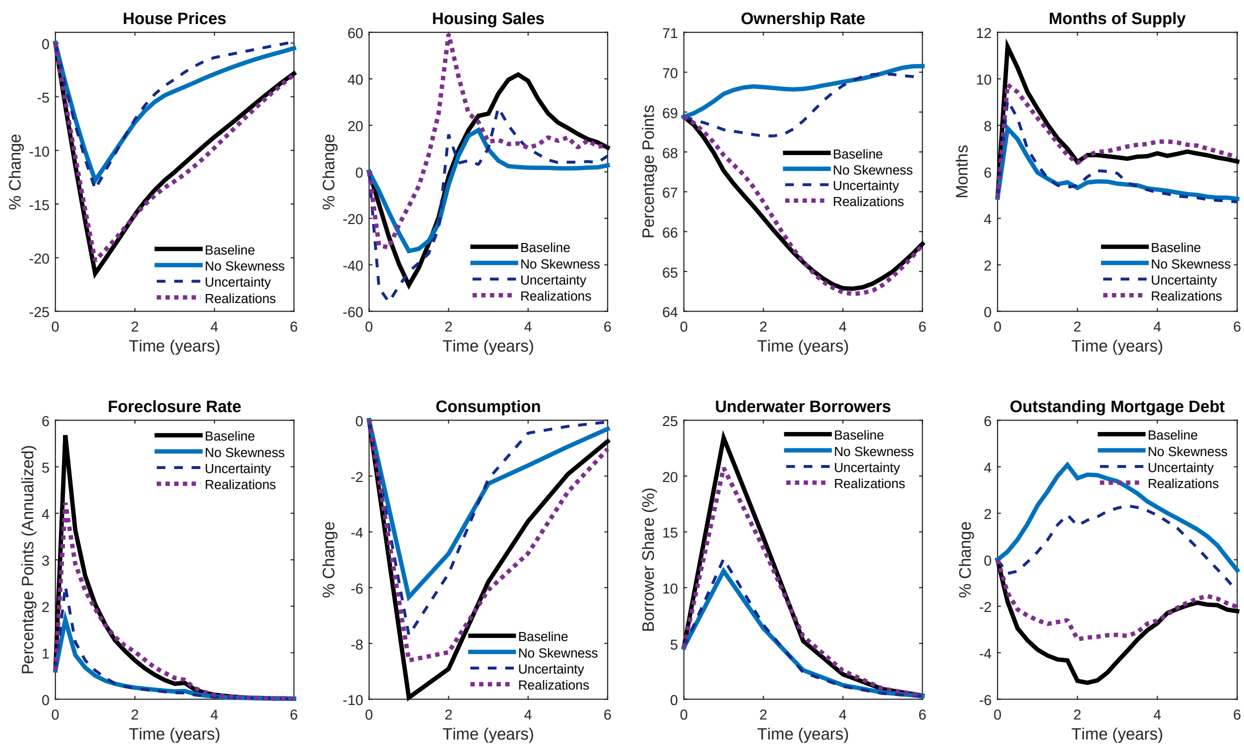

Figure 23: Decomposing the role of realizations vs. uncertainty from skewness shocks. "Uncertainty" captures just the effect of households' *belief* that skewness has worsened. The "Realizations" curve shows the impact of worse skewness when households are naively unaware of the change in earnings risk. The "No Skewness" plot shuts off skewness shocks entirely.

B.3.2 The Distributional Effects of Skewness and Credit Shocks

Complementing the discussion in section IV.A, figures 24 and 25 show the cross-sectional implications of skewness and credit limit shocks for consumption. The top row in each figure plots slices of average consumption over time by tenure, leverage, and financial distress status, and the bottom row shows the entire distribution of the peak-to-trough decline in consumption for each of these groups. In figure 24, removing the skewness shock provides significant relief to distressed owners, who no longer face the pressure to severely cut consumption in response to higher downside earnings risk. In figure 25, the removal of the credit limit shock relaxes homeowners' ability to extract equity through refinancing, which mitigates the consumption decline especially for highly leveraged and distressed owners.

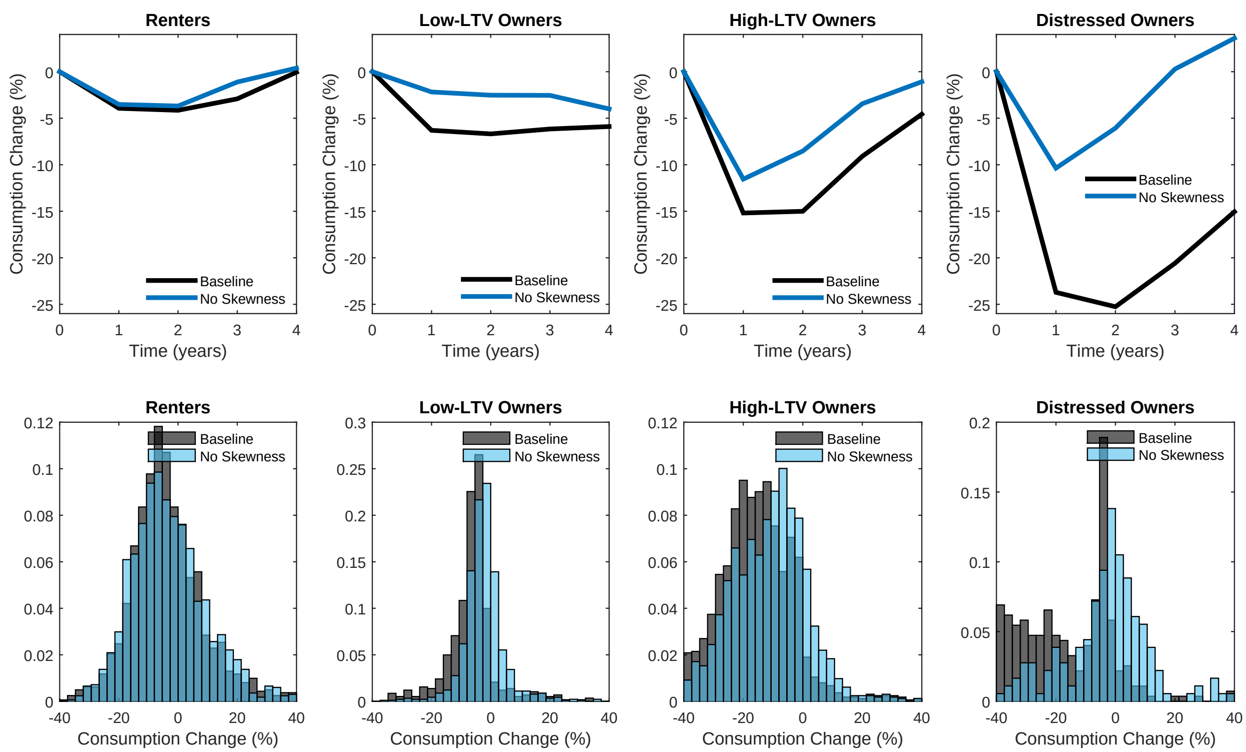

Figure 24: The consumption effects of removing skewness shocks. Low loan-to-value (LTV) is below 0.3; high LTV is above 0.8. Distressed owners are those who default 1 year *after* the beginning of the Great Recession.

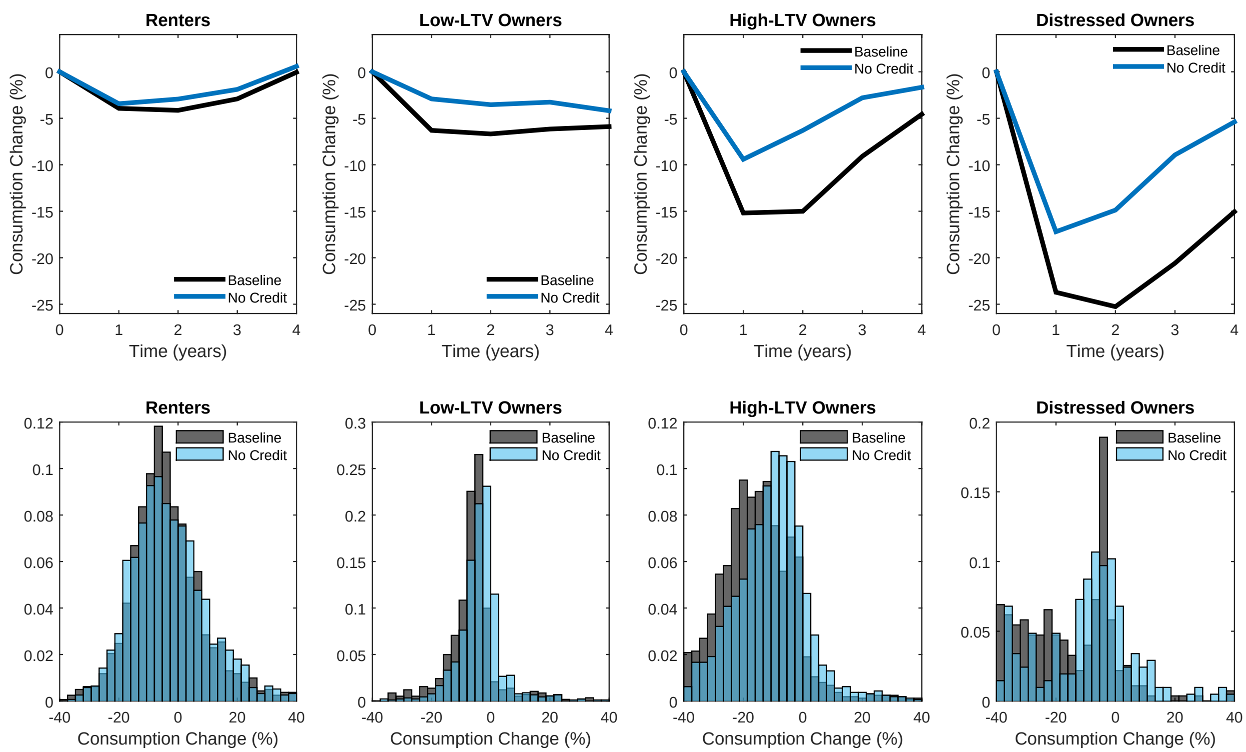

Figure 25: The consumption effects of removing credit shocks. Low loan-to-value (LTV) is below 0.3; high LTV is above 0.8. Distressed owners are those who default 1 year *after* the beginning of the Great Recession.

B.4 Robustness

This section explores the implications of altering either the timing or duration of the shocks that are summarized by the timeline in figure 26.

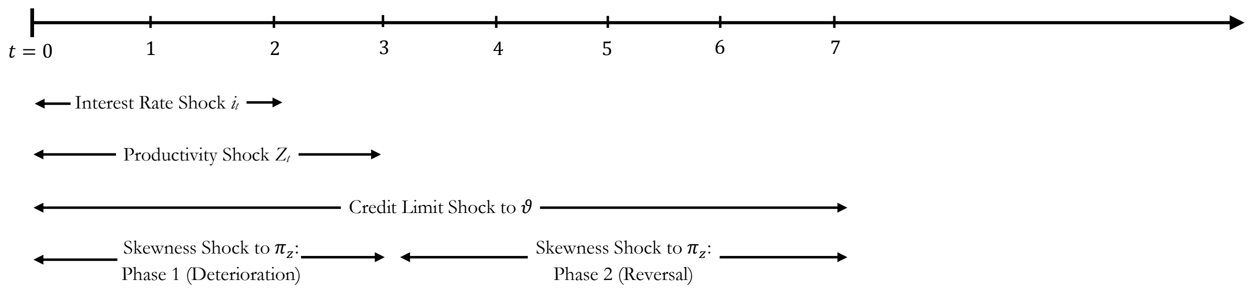

Figure 26: Shock timeline in the baseline.

B.4.1 Staggering the Arrival of Shocks

In the baseline implementation, all of the shocks arrive simultaneously. As a robustness test, this section subjects the economy to two waves of shocks. First, the unanticipated credit shocks arrive. Then, agents are surprised again one year later by the arrival of the skewness and productivity shocks. Figure 27 summarizes this staggered timeline.

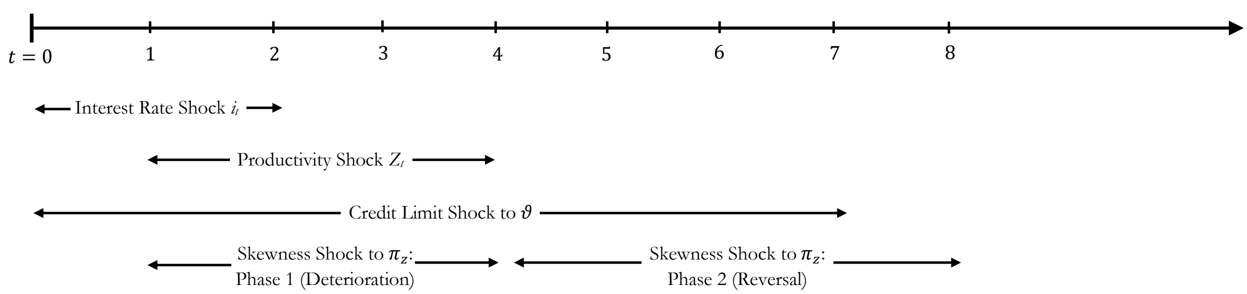

Figure 27: Staggered arrival of shocks.

As figure 28 makes evident, the dynamics of house prices and foreclosures are largely unchanged when the earnings and productivity shocks arrive with a delay. The staggered timing causes a slightly smaller deterioration in both variables, but it also produces more persistence. The delayed ownership decline in the middle panel reinforces the discussion in section IV.A that highlights the critical role of the skewness shocks in depressing ownership.

B.4.2 Forward-Looking Behavior and Terminal Conditions

Even though each of the following experiments only involves changing the *end date* of the shocks, terminal conditions affect economic dynamics during the entire transition path because of forward-looking behavior by households and lenders. Going from the most extreme to the least extreme case, extending the shocks to credit limits, interest rates,

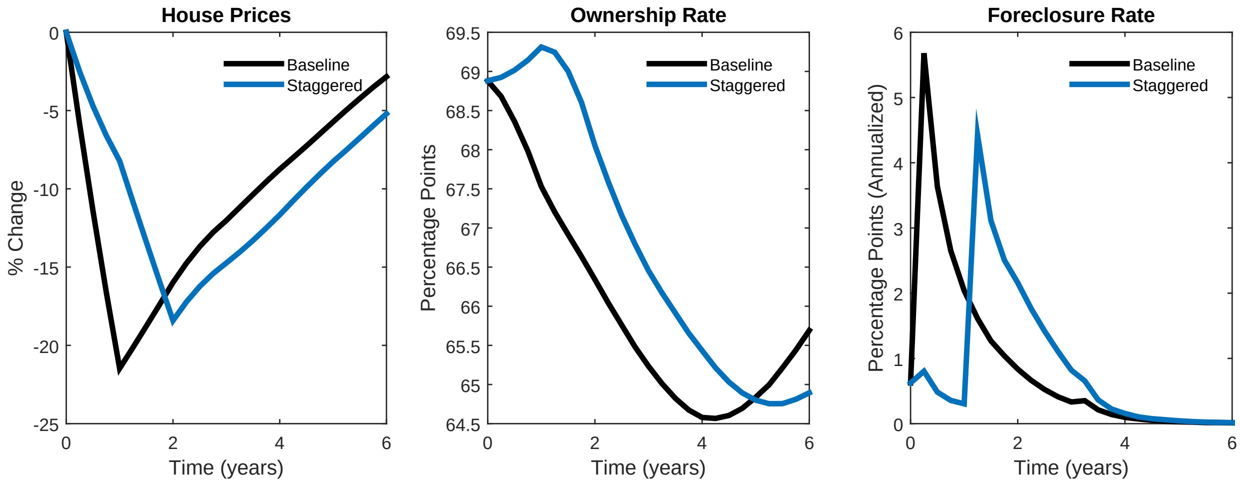

Figure 28: Results with a delayed arrival of skewness and productivity shocks.

and productivity indefinitely—which means a permanent 5% decline in wages, real risk-free rates of 3% instead of -1% (with higher mortgage rates as a result), and a permanent 10% minimum down payment requirement—causes a severe deepening of the housing crisis. House prices fall twice as far, the consumption decline is amplified by 60%, and the foreclosure rate peaks at 18% as nearly two-thirds of homeowners find themselves underwater (i.e. owing more on their mortgage than their house is worth). In addition, the stock of outstanding mortgage debt drops substantially. Needless to say, these results show the importance of future expectations about terminal conditions, but the simulated time series in this case are a dramatic exaggeration of what occurred in the data. Making only the credit limit and interest rate shocks permanent by allowing productivity to recover acts as as intermediate point between the baseline and the previous case, but the large, permanent decline in outstanding debt is at odds with the data. Lastly, making only the tighter down payment constraint permanent causes the trajectory of the economy to differ only modestly from the baseline path. However, a permanent 10% minimum down payment requirement is at odds with the return of low-LTV lending in recent years.

B.4.3 Alternative Drivers

Section IV.A affirms the role of credit as a driver of housing market behavior while also establishing the novel importance of downside uncertainty during the Great Recession. Although for this paper the elevated uncertainty in the model comes through higher left tail earnings risk from skewness shocks, a large body of recent empirical literature highlights the importance of uncertainty, broadly conceived, during the crisis.[7] For robustness, this

[7] See Shoag and Veuger (2016), Arellano, Bai and Kehoe (2019), Stock and Watson (2012), and Bloom, Floetotto, Jaimovich, Saporta-Eksten and Terry (2018) for an analysis of uncertainty in the Great Recession. In addition, Jurado, Ludvigson and Ng (2015) identify 2007–2009 as "the most striking period of heightened uncertainty since 1960," and Berger et al. (2019) emphasize the importance of left-skewed shocks for matching the relationship between uncertainty and the macroeconomy. Kozeniauskas, Orlik and Veldkamp (2018) provide a unified discussion of how to define and measure uncertainty.

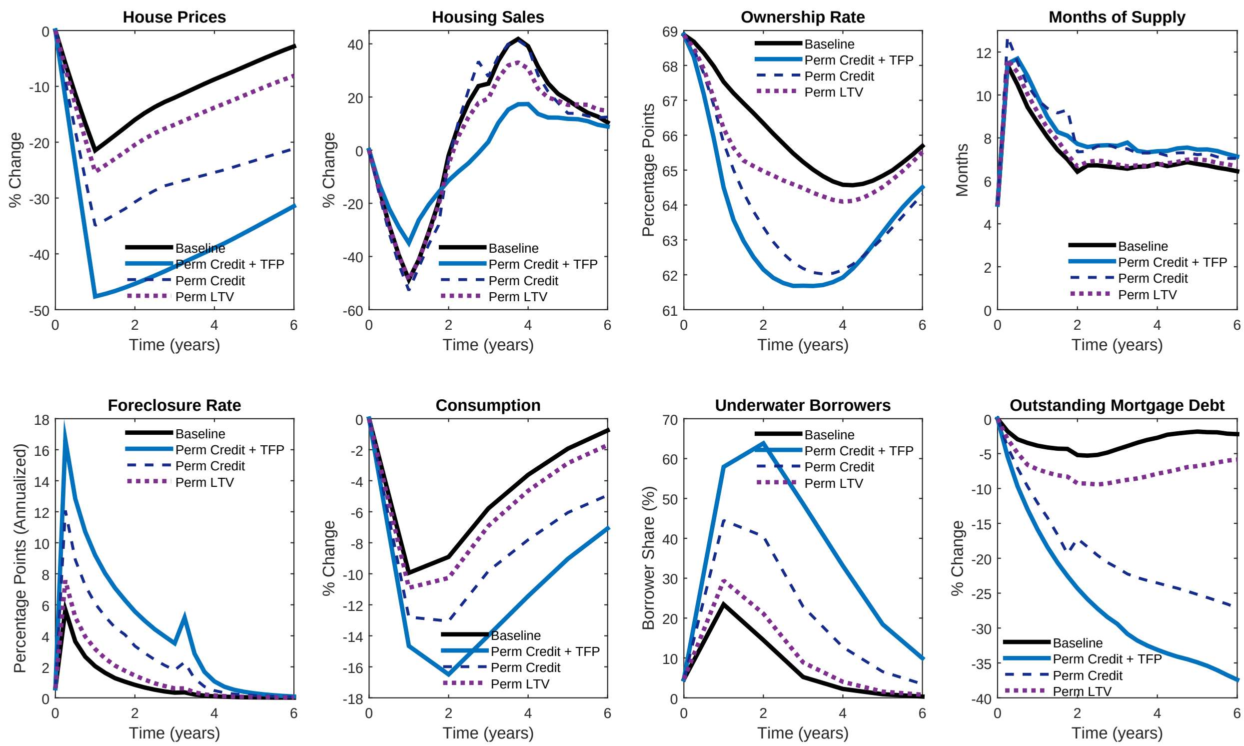

Figure 29: The effect of making shocks permanent. "Perm Credit + TFP" makes the credit, interest rate, and TFP shocks permanent, but the skewness shocks remain temporary. "Perm Credit" makes only the credit and interest rate shocks permanent (i.e. the real risk-free rate goes to 3% in the long run). "Perm LTV" makes only the tighter down payment constraint permanent.

section evaluates a selection of possible alternative drivers of the housing crash involving housing preference shocks, housing belief shocks, and unusually large productivity shocks.

Housing Preference Shocks Figure 30 shows the economic response to an *immediate, permanent* shock to preferences (specifically, a hike in the consumption weight ω) that reduces agents' taste for housing. The "pref shock only" curve shows the equilibrium path with *only* this shock, and the dashed line shows the combined effect of this preference shock together with the interest rate, credit limit, and productivity shocks from the baseline—that is, when the skewness shock is replaced with the preference shock. It is immediately evident from the counterfactual increase in consumption that the preference shock cannot be the only driving force behind the Great Recession. Furthermore, even though agents have a lower preference for housing, the decline in prices causes ownership and mortgage debt to counterfactually *rise*.

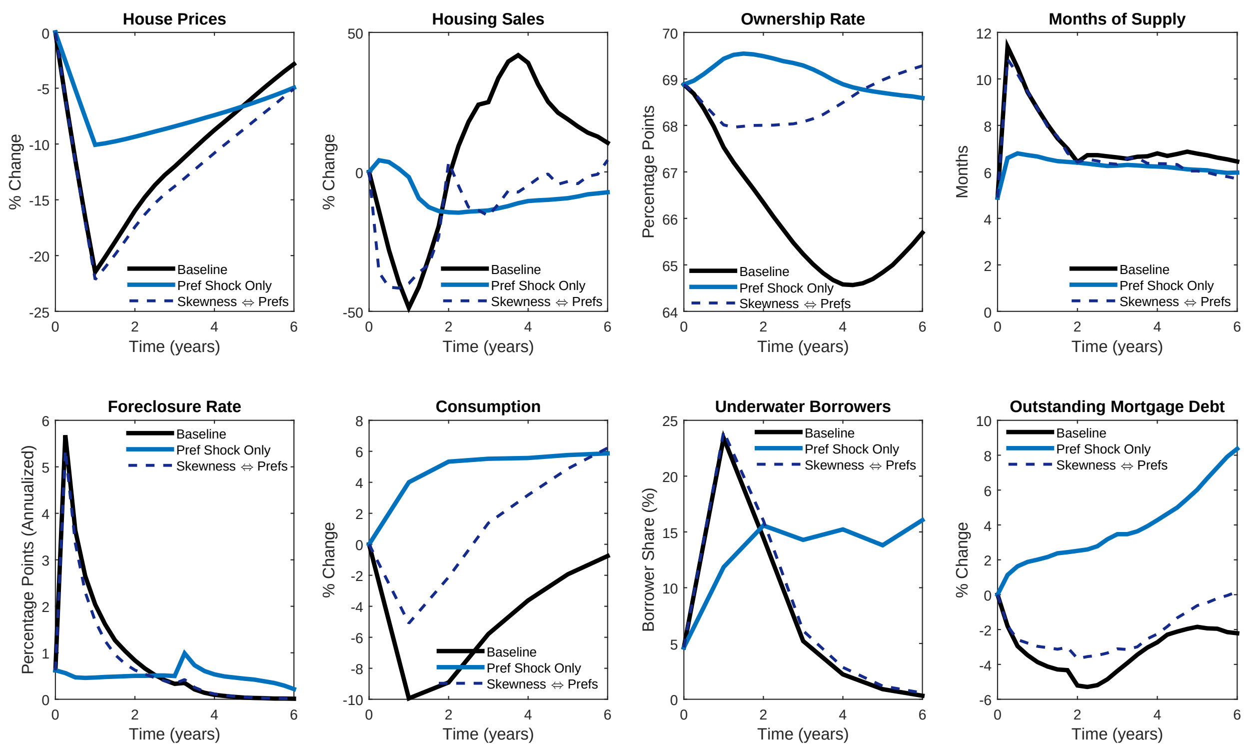

Figure 30: The effect of housing preference shocks. "Pref Shock Only" includes an *immediate, permanent* negative shock to the preference for housing as the only shock. "Skewness ⇔ Prefs" retains the credit, productivity, and interest rate shocks from the baseline but swaps out the skewness shock with the housing preference shock.

The addition of the interest rate, credit limit, and productivity shocks reverses this counterfactual consumption boom, but even so, the decline is too shallow and short-lived, and low-LTV owners still experience a counterfactual rise in consumption, as observed in figure 31. Moreover, homeownership barely budges instead of experiencing the deep decline observed in the data.

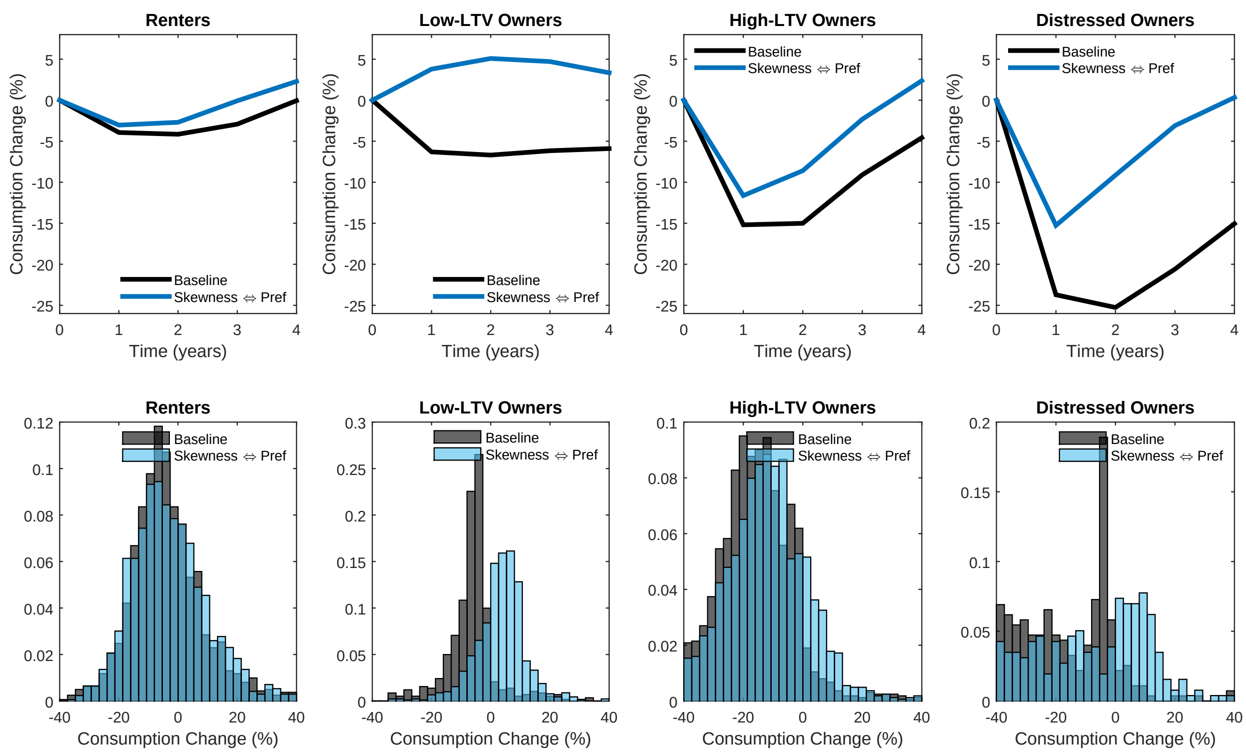

Figure 31: The distributional impact of swapping the skewness shock for a negative shock to housing preferences. Low loan-to-value (LTV) is below 0.3; high LTV is above 0.8. Distressed owners are those who default 1 year *after* the beginning of the Great Recession.

Pessimistic Housing Beliefs Kaplan, Mitman and Violante (2019) resolve the counterfactual negative co-movement between house prices and consumption induced by preference shocks during a simulated boom-bust episode by instead shocking households' *beliefs* about the likelihood of a *future* change in the taste for housing. Furthermore, to prevent these counterfactual dynamics from simply materializing later, Kaplan et al. (2019) then assume that the preference shock never actually materializes. Nevertheless, households still react immediately based on their (ex-ante rational, but ex-post false) beliefs about the future. Figure 32 shows the results of an analogous experiment in this model which announces to agents that there will be a *future, permanent* preference shock that decreases the taste for housing starting at $t = 3$. Regardless of whether the shock actually occurs (as in the figure) or unexpectedly never materializes (which would imply surprising households again at $t = 3$), the endogenous economic response between $t = 0$ and $t = 3$ to agents' pessimistic housing beliefs is the same.

Figure 32 shows that, if the anticipated future preference shock is large enough, it can induce a significant decline in house prices but still only a modest rise in foreclosures and drop in consumption. Moreover, the belief shock produces almost no change in homeownership, even if combined with the interest rate, credit limit, and productivity shocks, and yet still manages to generate excessive deleveraging as pessimistic homeowners unload mortgage debt from their balance sheets. This behavior conflicts with the empirical evidence in Bhutta (2015), which attributes most of the modest decline in outstanding mortgage debt to reduced inflows caused by a "dramatic falloff in first-time homebuying."

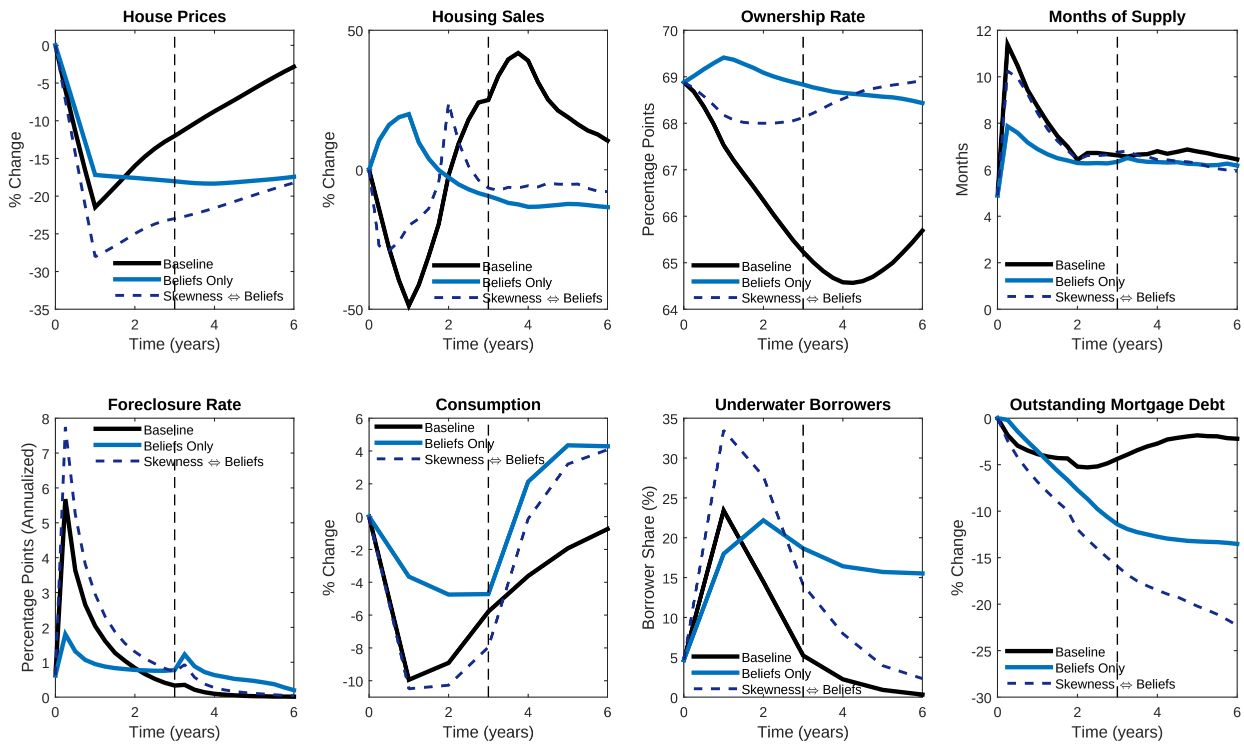

Figure 32: The effect of housing pessimism. "Beliefs Only" introduces a *future, permanent, and fully anticipated* negative shock to the preference for housing as the only shock. "Skewness ⇔ Beliefs" retains the credit, productivity, and interest rate shocks from the baseline but swaps out the skewness shock with the housing pessimism shock.

Productivity Disasters and Earnings Pessimism Section IV finds that the baseline 5% productivity shock contributes only modestly to the housing crash. By contrast, the skewness shock has large effects. On the surface, both types of shocks reduce aggregate earnings. However, they differ profoundly in other ways. Specifically, the productivity shock reduces earnings deterministically, immediately, and uniformly for all households. By contrast, the skewness shock causes earnings to fall *stochastically*, *gradually*, and *unevenly*. Thus, relative to productivity shocks, the deterioration in skewness concentrates bad earnings realizations among the lower and middle class while simultaneously creating a sense of *earnings pessimism* that induces precautionary behavior—most notably the desire to avoid the consumption commitment of homeownership until housing becomes more liquid and less risky.

To isolate the effect of earnings pessimism from the uneven incidence of lower earnings realizations—and to test for nonlinearities in the response of housing and consumption to aggregate earnings—the baseline skewness shocks to π_z are replaced with a sequence of across-the-board shocks to effective labor supply $e_t \cdot z_t$ that produce the same path of aggregate labor shown in figure 15. Notably, the combination of the original productivity shock with these labor shocks is isomorphic to a "productivity disaster" consisting of a gradual deterioration and recovery in aggregate productivity that bottoms out at just over a 10% decline, which is similar to Glover, Heathcote, Krueger and Ríos-Rull (2019). Thus, to simplify exposition, this experiment analyzes the impact of a productivity disaster working in conjunction with the interest rate and credit limit shocks. Importantly, because households perceive that productivity will continue to fall for multiple years, they have pessimism about future earnings. However, unlike with skewness shocks, households in this case know with certainty that they will all equally bear the burden of the aggregate earnings reduction.

Figure 33 shows that, even in this productivity disaster scenario, the deterioration in house prices, foreclosures, and consumption is noticeably smaller than in the baseline. Furthermore, the familiar counterfactual rise in homeownership and mortgage debt re-emerges. Thus, combining the insights of this experiment with the uncertainty vs. realizations decomposition in section B.3.1 reveals that the skewness shocks play such an important role as a driver of the housing crash because of the trifecta of higher uncertainty, worse pessimism, and an uneven incidence of the aggregate earnings decline across households felt most saliently at the bottom and middle of the income distribution.

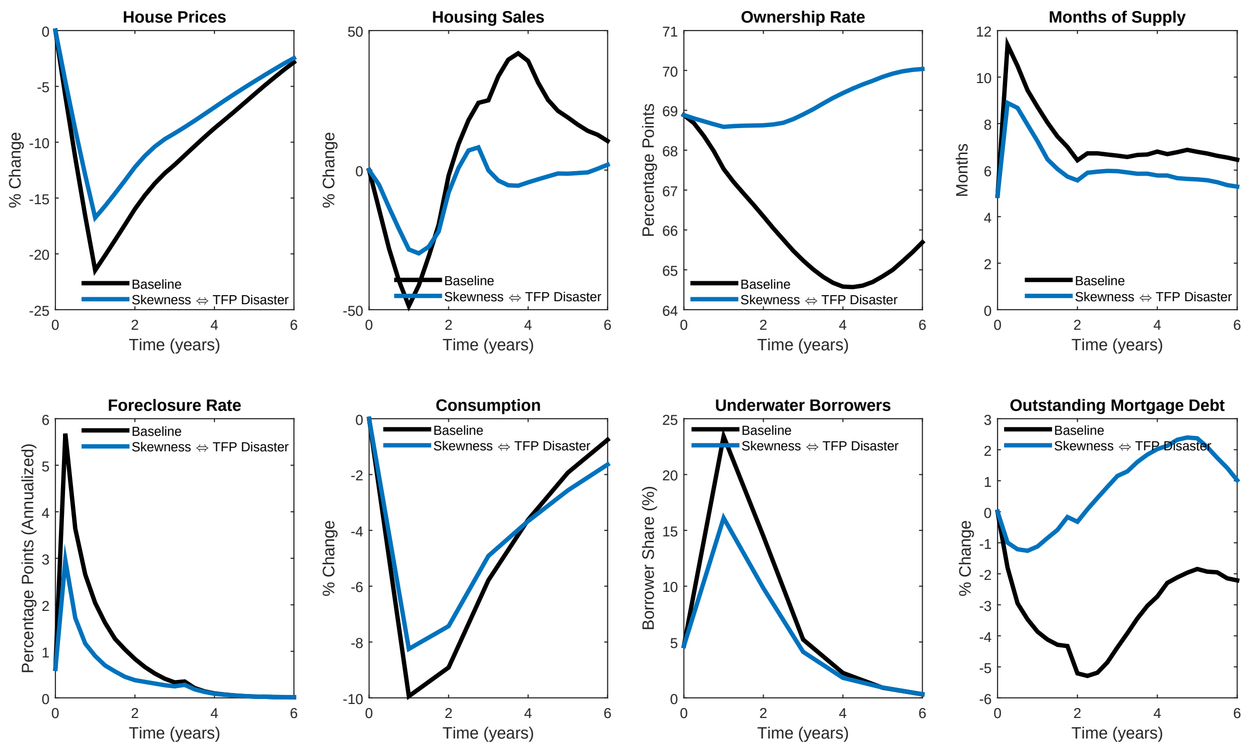

Figure 33: Replacing skewness shocks with a “productivity disaster” that combines the original productivity shock with a sequence of across-the-board shocks to effective labor supply that produces the same path of aggregate labor as does the skewness shocks.

B.5 Housing Spillovers to Consumption

This section provides supplemental information regarding the transmission of house prices to consumption through the balance sheet.

B.5.1 Aggregate Nonlinearities and Shock Dependence

Section IV.C mentions that the elasticity of consumption to house prices is not a single invariant number. Instead, the transmission from housing to consumption is dynamic and shock-dependent. Figures 34 and 35 provide a visual illustration of this point by taking turns either removing or introducing each shock one at a time, just as in the decompositions from section IV.A.

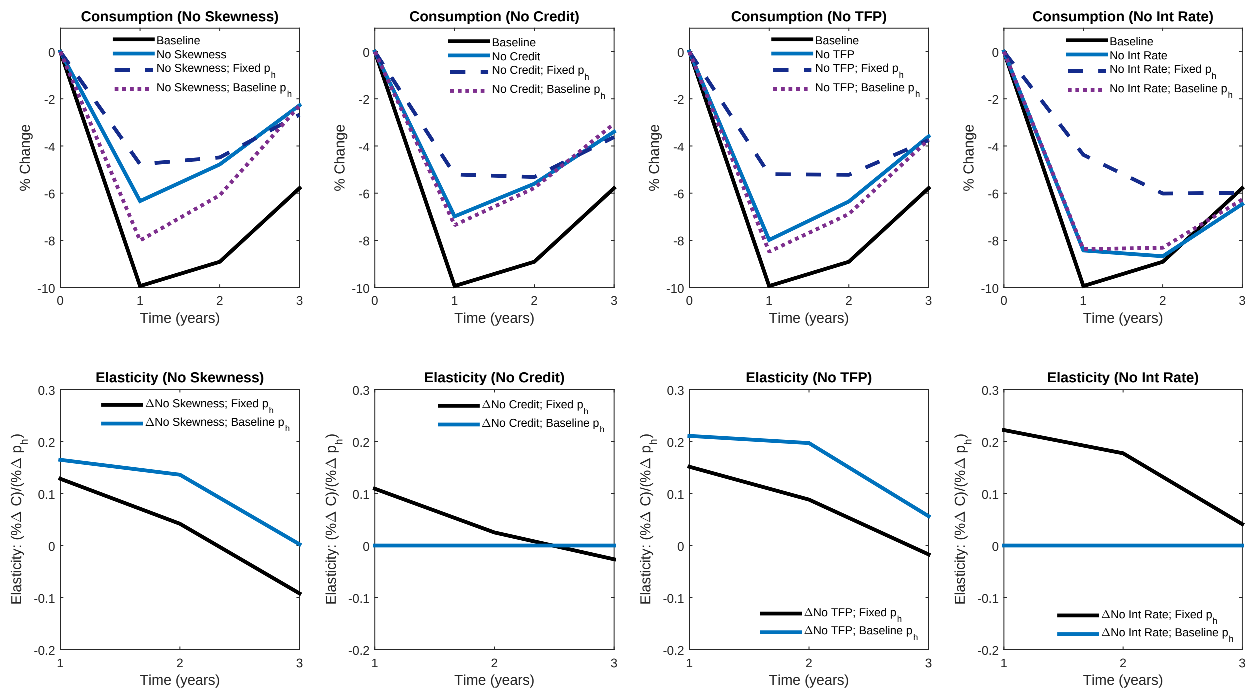

Figure 34: Consumption elasticity to house prices when one shock is removed. The elasticity is calculated by comparing equilibrium consumption to either: (fixed) consumption when house prices are fixed; (baseline) consumption when house prices follow their baseline trajectory.

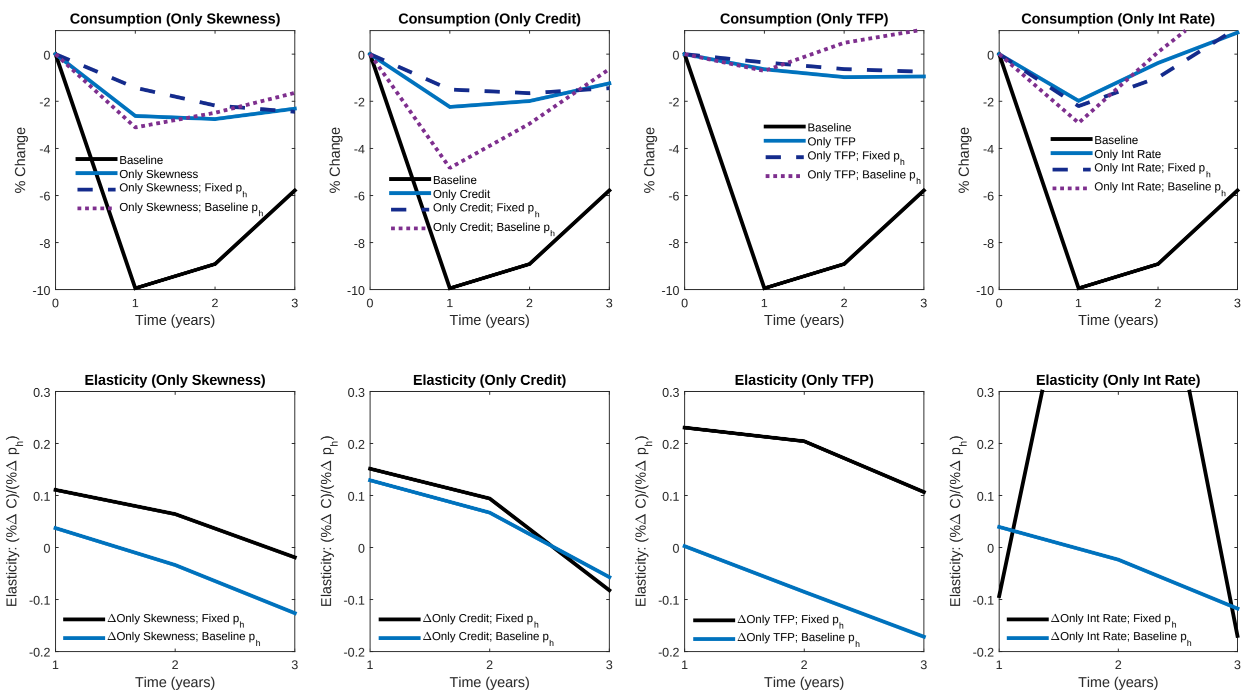

Figure 35: Consumption elasticity to house prices when *the other* shocks are removed. The elasticity is calculated by comparing equilibrium consumption to either: (fixed) consumption when house prices are fixed; (baseline) consumption when house prices follow their baseline trajectory.

B.5.2 Balance Sheet Depth and Consumption in the Cross Section

The main text points out that endogenous housing liquidity amplifies consumption differentially throughout the cross-section. In particular, selling delays increase the mass of the left tail of the consumption decline histogram for highly leveraged and distressed owners, as shown in figure 36. By contrast, renters are sheltered from the debt overhang caused by decreasing liquidity.

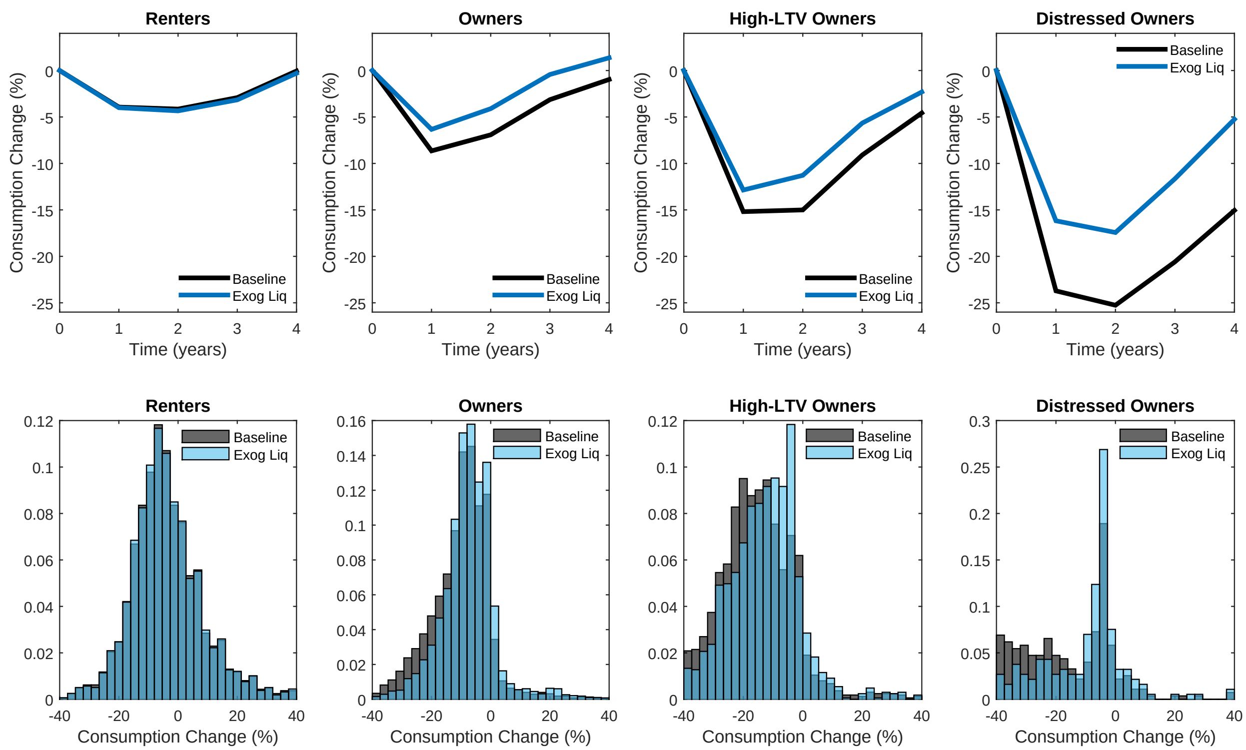

Figure 36: The amplification effects of endogenous liquidity on consumption. High loan-to-value (LTV) is above 0.8. Distressed owners are those who default 1 year *after* the beginning of the Great Recession. The histograms show the distribution of consumption changes during the bust.

Figure 37 corresponds to table 7 in section IV.C, which discusses the importance of balance sheet depth for the behavior of consumption during the crisis. The top row plots slices of average consumption over time by net worth bin, and the bottom row shows the entire distribution of the peak-to-trough decline in consumption for each of these groups.

Figure 38 refers to the discussion in section 38 about the higher order consumption implications of balance sheet depth. The bottom left panel shows that homeowners are better able than renters to insure against income shocks during normal times. However, the middle left panel provides a stark depiction of the deterioration in the risk properties of leveraged homeownership, as the left tail of consumption dynamics expands to the left for owners relative to renters. Similar patterns emerge when comparing homeowners with different amounts of leverage to each other, as shown in the middle column. Lastly, the column on the right reveals that distressed owners who are least able to extract equity

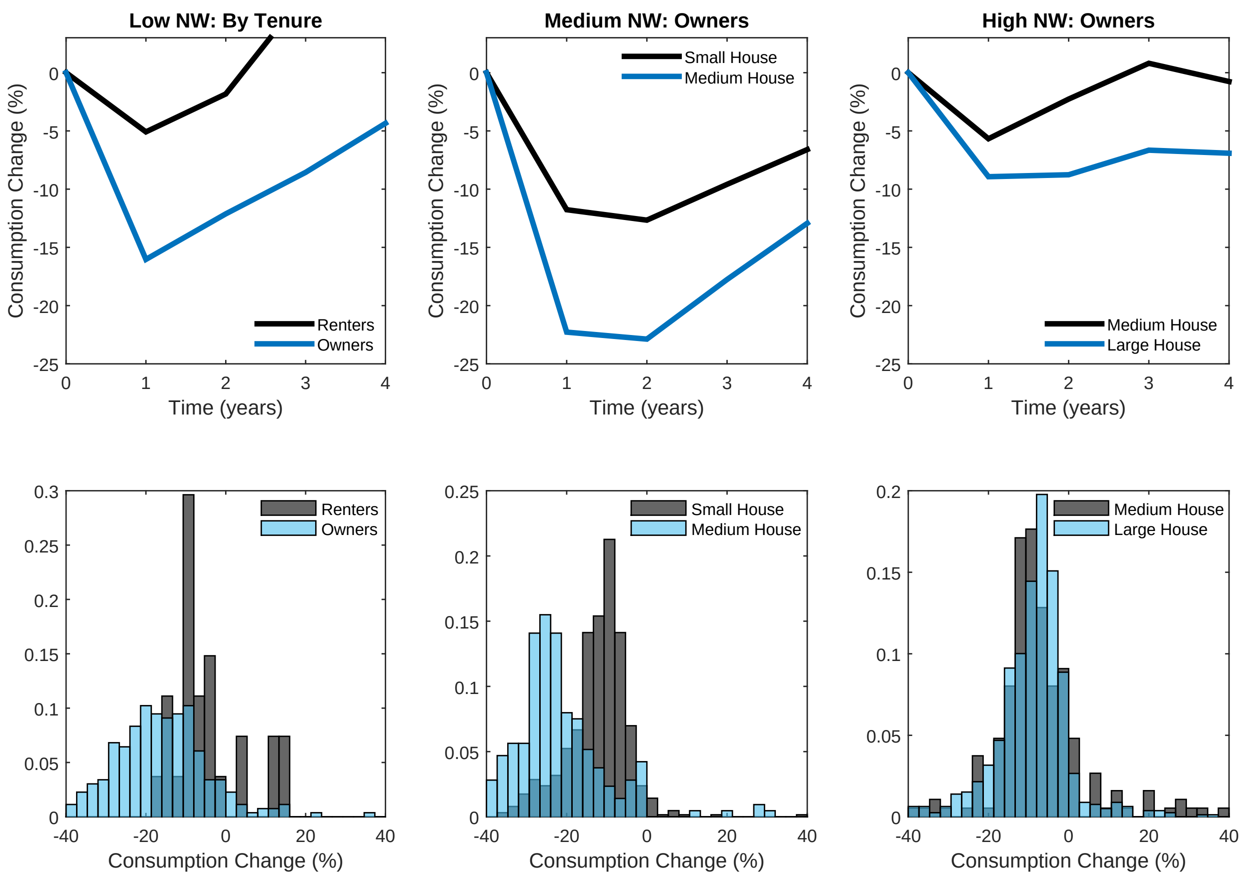

Figure 37: Consumption by net worth (NW) decile and balance sheet depth. Low NW is decile 4; medium NW is decile 6; high NW is decile 9.

either by selling (because of debt overhang) or refinancing (because of their default risk) suffer the largest consumption declines. Triggering the default option provides considerable consumption relief, though at the expense of eviction and several years of exclusion from the mortgage market.

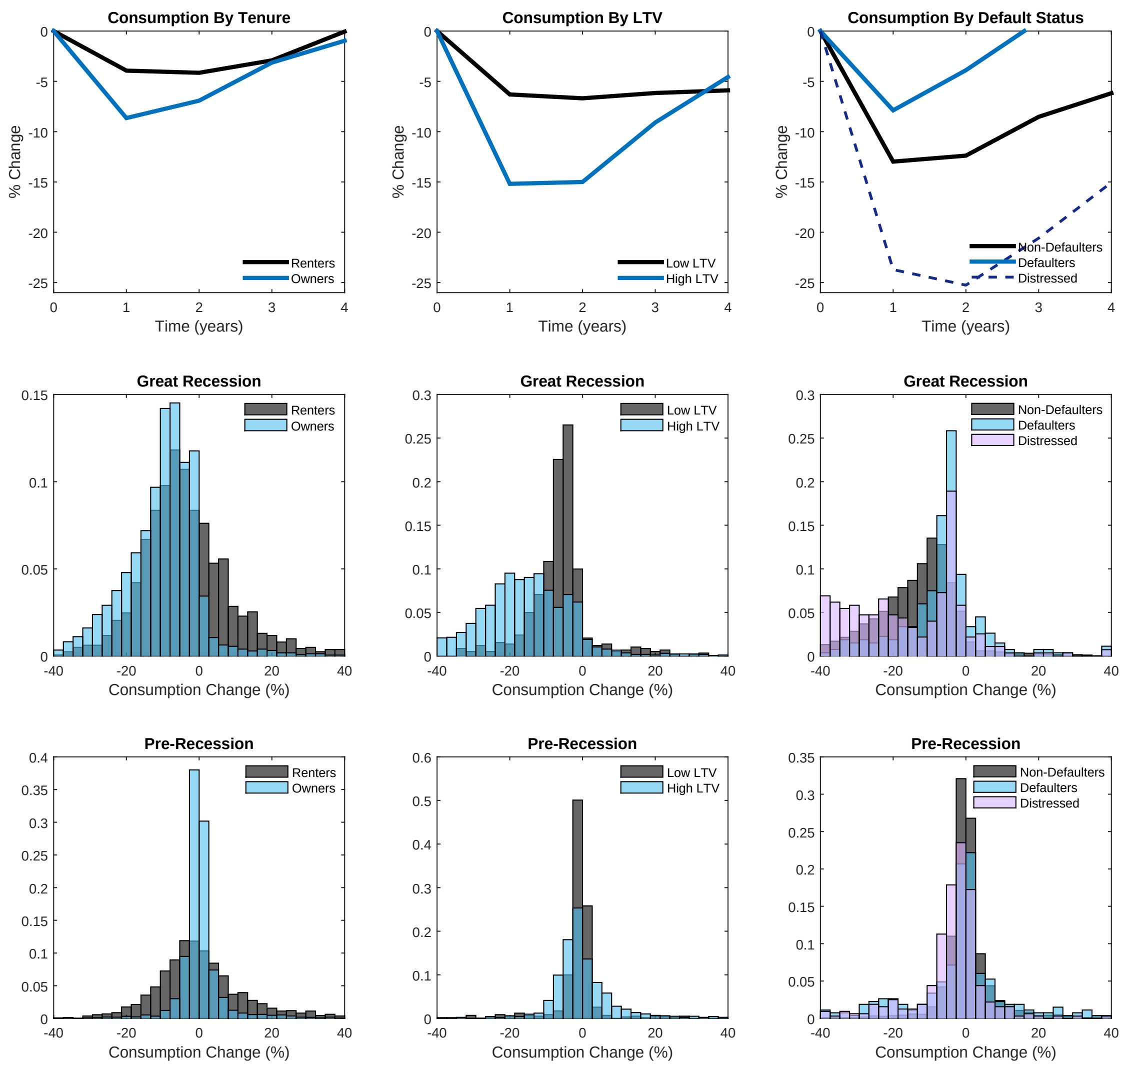

Figure 38: Consumption by tenure status, leverage, and financial distress. Low loan-to-value (LTV) is below 0.3; high LTV is above 0.8. Distressed owners are those who default 1 year *after* the beginning of the Great Recession. The middle row of histograms shows the distribution of consumption changes during the Great Recession. The bottom row displays the dispersion in pre-recession consumption dynamics.

B.6 Mortgage Rate Reductions and Heterogeneity

Section IV.D refers to figure 39 when it describes the ability of mortgage rate reductions to stimulate house prices during the crisis, which in turn accelerates the recovery of consumption by partially repairing household balance sheets. Figure 40 provides additional cross-sectional support for this mechanism by showing that the potency of the policy increases with leverage. Specifically, highly leveraged owners experience the largest consumption boost.

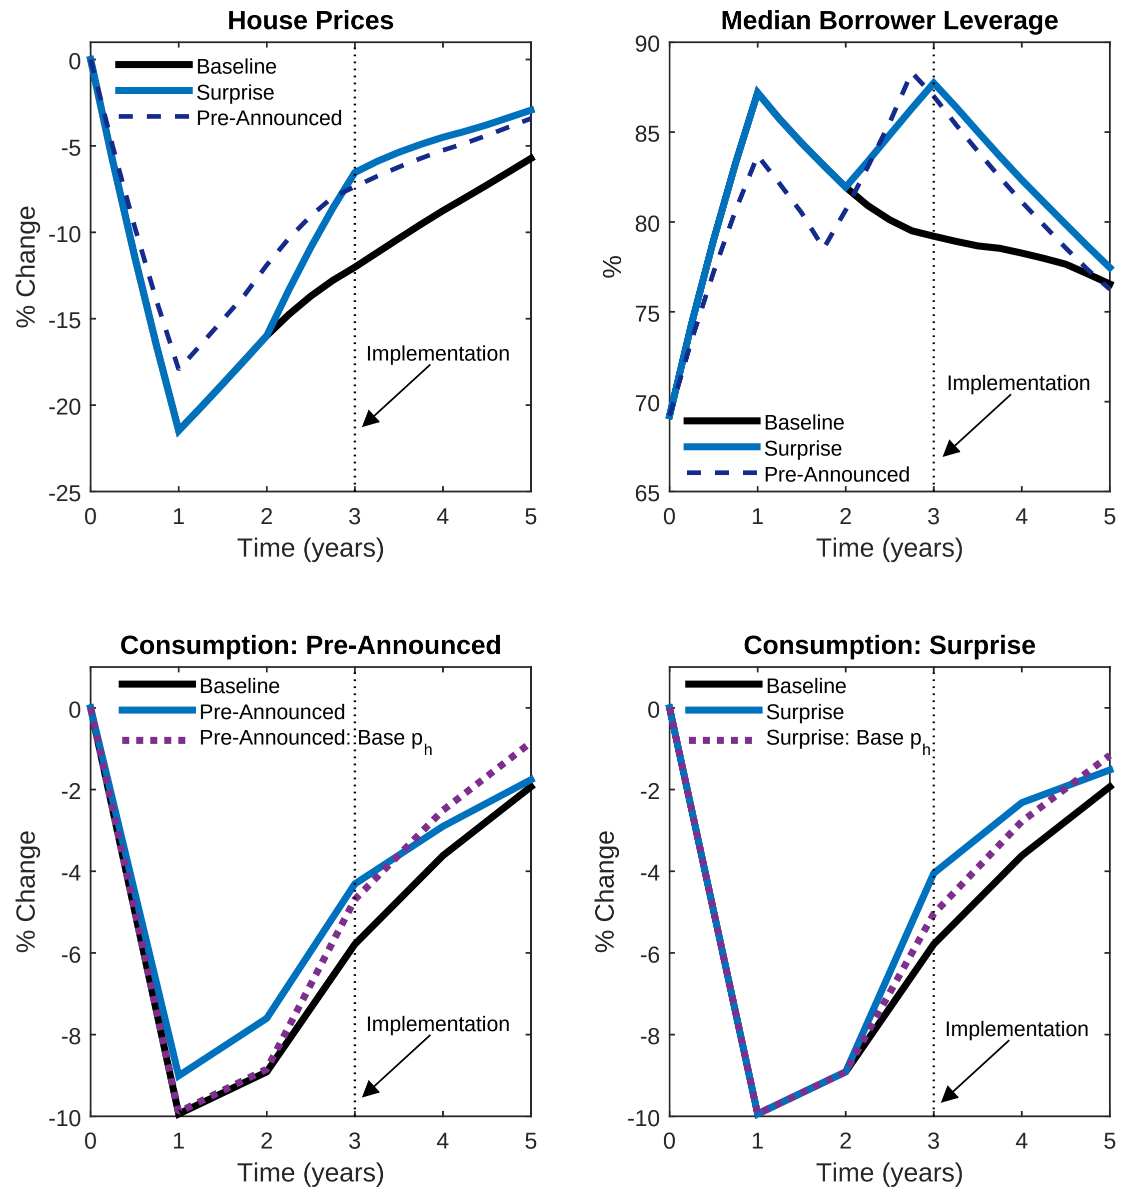

Figure 39: The effects of lowering mortgage rates at $t = 3$ when either pre-announced at the time of the shocks or else implemented by surprise.

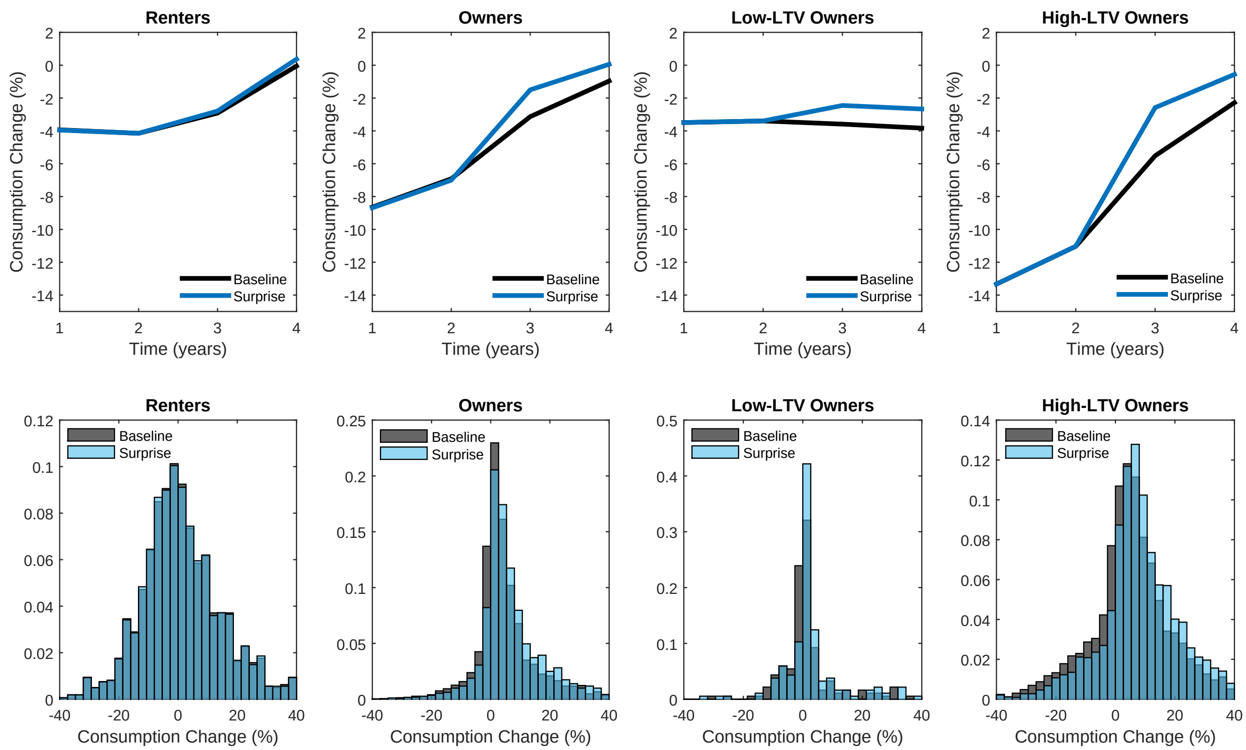

Figure 40: Consumption response to the mortgage rate reduction policy by tenure status and leverage. Low loan-to-value (LTV) is below 0.3; high LTV is above 0.8.

C Model Equations and Equilibrium

This section gives the complete definition of equilibrium from section II.F.

C.1 Household Value Functions

C.1.1 Consumption and Balance Sheet Decisions

Homeowners who take out a new mortgage:

$$V_t^{own,0}(y_t,h,z_t)=\max_{\substack{m_{t+1}\geq 0,\\ b_{t+1}\geq 0,\\ c_t\geq 0}} u(c_t,h)+\beta\mathbb{E}\left[\begin{array}{l}(1-\delta_h)[W_{t+1}^{own,0}(y_{t+1},(r_{t+1},m_{t+1}),h,z_{t+1})\\ +R_{t+1}^{sell,0}(y_{t+1},(r_{t+1},m_{t+1}),h,z_{t+1})]\\ +\delta_h[V_{t+1}^{rent,0}(y_{t+1},z_{t+1})+R_{t+1}^{buy,0}(y_{t+1},z_{t+1})]\end{array}\right]$$

$$\begin{gathered}\text{subject to}\\ c_t+\gamma p_t h+b_{t+1}/(1+i_{t+1})\leq y_t+q_t((r_{t+1},m_{t+1}),b_{t+1},h,z_t)m_{t+1}\\ q_t((r_{t+1},m_{t+1}),b_{t+1},h,z_t)m_{t+1}\leq\vartheta p_t h\\ y_{t+1}=w_{t+1}e_{t+1}z_{t+1}+b_{t+1}\end{gathered}\tag{15}$$

Homeowners who make a payment on their current mortgage:

$$V_t^{amort}(y_t,(\bar{r},m_t),h,z_t)=\max_{\substack{b_{t+1}\geq 0,\\ c_t\geq 0,\\ l_t}} u(c_t,h)+\beta\mathbb{E}\left[\begin{array}{l}(1-\delta_h)[W_{t+1}^{own,0}(y_{t+1},(\bar{r},m_{t+1}),h,z_{t+1})\\ +R_{t+1}^{sell,0}(y_{t+1},(\bar{r},m_{t+1}),h,z_{t+1})]\\ +\delta_h[V_{t+1}^{rent,0}(y_{t+1},z_{t+1})+R_{t+1}^{buy,0}(y_{t+1},z_{t+1})]\end{array}\right]$$

$$\begin{gathered}\text{subject to}\\ c_t+\gamma p_t h+b_{t+1}/(1+i_{t+1})+l_t\leq y_t\\ \frac{\bar{r}}{1+\bar{r}}m_t\leq l_t\leq m_t\\ m_{t+1}=(m_t-l_t)(1+\bar{r})\\ y_{t+1}=w_{t+1}e_{t+1}z_{t+1}+b_{t+1}\end{gathered}\tag{16}$$

Homeowners who default but are not foreclosed on:

$$V_t^{delinq}(y_t,(\bar{r},m_t),h,z_t)=\max_{\substack{b_{t+1}\geq 0,\\ c_t\geq 0}} u(c_t,h)+\beta\mathbb{E}\left[\begin{array}{l}(1-\delta_h)[W_{t+1}^{own,0}(y_{t+1},(\bar{r},m_t),h,z_{t+1})\\ +R_{t+1}^{sell,0}(y_{t+1},(\bar{r},m_t),h,z_{t+1})]\\ +\delta_h[V_{t+1}^{rent,0}(y_{t+1},z_{t+1})+R_{t+1}^{buy,0}(y_{t+1},z_{t+1})]\end{array}\right]$$

$$\begin{gathered}\text{subject to}\\ c_t+\gamma p_t h+b_{t+1}/(1+i_{t+1})\leq y_t\\ y_{t+1}=w_{t+1}e_{t+1}z_{t+1}+b_{t+1}\end{gathered}\tag{17}$$

Homeowners with a default flag:

$$V_t^{own,1}(y_t,h,z_t)=\max_{b_{t+1},c_t\geq 0} u(c_t,h)+\beta\mathbb{E}\left[\begin{array}{l}(1-\delta_h)\{(1-\lambda_f)[W_{t+1}^{own,0}(y_{t+1},0,h,z_{t+1})\\ +R_{t+1}^{sell,0}(y_{t+1},0,h,z_{t+1})]\\ +\lambda_f[V_{t+1}^{own,1}(y_{t+1},h,z_{t+1})+R_{t+1}^{sell,1}(y_{t+1},h,z_{t+1})]\}\\ +\delta_h\{(1-\lambda_f)[V_{t+1}^{rent,0}(y_{t+1},z_{t+1})+R_{t+1}^{buy,0}(y_{t+1},z_{t+1})]\\ +\lambda_f[V_{t+1}^{rent,1}(y_{t+1},z_{t+1})+R_{t+1}^{buy,0}(y_{t+1},z_{t+1})]\}\end{array}\right]$$
$$\text{subject to}$$
$$c_t+\gamma p_t h+b_{t+1}/(1+i_{t+1})\leq y_t$$
$$y_{t+1}=w_{t+1}e_{t+1}z_{t+1}+b_{t+1} \tag{18}$$

Renters with good credit:

$$V^{rent,0}(y_t,z_t)=\max_{\substack{b_{t+1},c_t\geq 0,\\ 0\leq a_t\leq \overline{a}}} u(c_t,a_t)+\beta\mathbb{E}\left[V_{t+1}^{rent,0}(y_{t+1},z_{t+1})+R_{t+1}^{buy,0}(y_{t+1},z_{t+1})\right]$$
$$\text{subject to} \tag{19}$$
$$c_t+b_{t+1}/(1+i_{t+1})+r_a a_t\leq y_t$$
$$y_{t+1}=w_{t+1}e_{t+1}z_{t+1}+b_{t+1}$$

Renters with a default flag:

$$V^{rent,1}(y_t,z_t)=\max_{\substack{b_{t+1},c_t\geq 0,\\ 0\leq a_t\leq \overline{a}}} u(c_t,a_t)+\beta\mathbb{E}\left[\begin{array}{l}(1-\lambda_f)[V_t^{rent,0}(y_t,z_t)+R^{buy,0}(y_t,z_t)]\\ +\lambda_f[V_t^{rent,1}(y_t,z_t)+R^{buy,1}(y_t,z_t)]\end{array}\right]$$
$$\tag{20}$$
$$\text{subject to}$$
$$c_t+b_{t+1}/(1+i_{t+1})+r_a a_t\leq y_t$$
$$y_{t+1}=w_{t+1}e_{t+1}z_{t+1}+b_{t+1}$$

C.1.2 House Buying Decisions

The value of searching to buy a house:

$$R_t^{buy,0}(y_t,z_t)=\max\{0,\max_{\substack{h_t\in H,\\ p_t^{bid}\leq y_t-\underline{y}}}\eta_t^{buy}(p_t^{bid},h_t)[V_t^{own,0}(y_t-p_t^{bid},h_t,z_t)-V_t^{rent,0}(y_t,z_t)]\} \tag{21}$$
$$R_t^{buy,1}(y_t,z_t)=\max\{0,\max_{\substack{h_t\in H,\\ p_t^{bid}\leq y_t}}\eta_t^{buy}(p_t^{bid},h_t)[V_t^{own,1}(y_t-p_t^{bid},h_t,z_t)-V_t^{rent,1}(y_t,z_t)]\} \tag{22}$$

C.1.3 Mortgage Default, Amortization, and Refinancing Decisions

The value function for the decision to default, refinance, or make a payment:

$$\begin{aligned} W_t^{own,0}(y_t,(\bar{r},m_t),h,z_t) = \max \Big\{ \varphi [V_t^{rent,1}(y_t,z_t) + R_t^{buy,1}(y_t,z_t)] \\ +(1-\varphi)V_t^{delinq}(y_t,(\bar{r},m_t),h,z_t), V_t^{amort}(y_t,(\bar{r},m_t),h,z_t), V_t^{own,0}(y_t-m_t,h,z_t) \Big\} \end{aligned} \tag{23}$$

C.1.4 House Selling Decisions

The option value of selling for an owner with good credit:

$$\begin{aligned} R_t^{sell,0}(y_t,(\bar{r},m_t),h,z_t) = \max\{0, \max_{p_t^{list}\geq 0} \eta_t^{sell}(p_t^{list},h) \left[V_t^{rent,0}\left(y_t + p_t^{list} - m_t, z_t\right) \right. \\ \left. +R_t^{buy,0}\left(y_t + p_t^{list} - m_t, z_t\right) - W_t^{own,0}(y_t,(\bar{r},m_t),h,z_t)\right] + \left[1 - \eta_t^{sell}(p_t^{list},h)\right](-\xi)\} \\ \text{subject to} \\ p_t^{list} \geq m_t - y_t \end{aligned} \tag{24}$$

The option value of selling for an owner with a default flag:

$$\begin{aligned} R_t^{sell,1}(y_t,0,h,z_t) = \max\{0, \max_{p_t^{list}\geq 0} \eta_t^{sell}(p_t^{list},h) \left[V_t^{rent,1}\left(y_t + p_t^{list}, z_t\right) \right. \\ \left. +R_t^{buy,1}\left(y_t + p_t^{list}, z_t\right) - V_t^{own,1}(y_t,h,z_t)\right] + \left[1 - \eta_t^{sell}(p_t^{list},h)\right](-\xi)\} \end{aligned} \tag{25}$$

C.2 Production

C.2.1 Composite Consumption

The optimality condition for numeraire good production is

$$w_t = Z_t \tag{26}$$

C.2.2 Apartment Space

The optimality condition for apartment space production is

$$r_a = \frac{1}{A} \tag{27}$$

C.2.3 Housing Construction

The optimality conditions for new housing construction are

$$1 = p_t \frac{\partial F_h(\overline{L}, S_{ht}, N_{ht})}{\partial S_{ht}} \tag{28}$$

$$w_t = p_t \frac{\partial F_h(\overline{L}, S_{ht}, N_{ht})}{\partial N_{ht}} \tag{29}$$

C.3 Financial Sector

The interest rate for new mortgages satisfies

$$1 + r_{t+1} = \frac{(1+\phi)(1+i_{t+1}^{LR})}{1-\delta_h} \Rightarrow r_{t+1} \approx i_{t+1}^{LR} + \phi + \delta_h \tag{30}$$

where i_{t+1}^{LR} is the *long-run* cost of external financing (e.g. the 10-year treasury rate) that can deviate from the short-run rate i_{t+1} outside of steady state.
Mortgage prices satisfy the recursive relationship

$$\begin{aligned}
(1+\zeta)q_t((\overline{r}, m_{t+1}), b_{t+1}, h, z_t) = \frac{1}{1+r_{t+1}}\mathbb{E}\Bigg\{ & \overbrace{\eta_{t+1}^{sell}}^{\text{sell, repay}} + \overbrace{(1-\eta_{t+1}^{sell})}^{\text{no house sale}}\Bigg[\overbrace{d_{t+1}^*}^{\text{default}} \overbrace{\varphi \min\left\{1, \frac{J_{t+1}^{REO}(h)}{m_{t+1}}\right\}}^{\text{foreclosure recovery ratio}} \\
& + \underbrace{d_{t+1}^*(1-\varphi)(1+\zeta)q_{t+1}^{delinq}}_{\text{delinquency continuation value}} + (1-d_{t+1}^*)\left\{ \underbrace{\mathbf{1}_{[\text{Refi,t+1}]}}_{\text{repay in full}} + \mathbf{1}_{[\text{No Refi,t+1}]} \underbrace{\left(\frac{l_{t+1}^* + (1+\zeta)q_{t+1}^{cont} m_{t+2}^*}{m_{t+1}}\right)}_{\text{payment + continuation value}}\right\}\Bigg]\Bigg\}
\end{aligned}$$

such that

$$\begin{aligned}
\eta_{t+1}^{sell} &\equiv \eta_s(\theta_s(p_{t+1}^{list*}, h; p_{t+1})) \ \textbf{(probability of house sale)} \\
q_{t+1}^{delinq} &\equiv q_{t+1}((\overline{r}, m_{t+1}), b_{t+2}^{delinq*}, h, z_{t+1}) \ \textbf{(mark-to-market price for delinquent } m_{t+1}\textbf{)} \\
q_{t+1}^{cont} &\equiv q_{t+1}((\overline{r}, m_{t+2}^*), b_{t+2}^*, h, z_{t+1}) \ \textbf{(mark-to-market price for updated } m_{t+2}^*\textbf{)} \\
m_{t+2}^* &= (m_{t+1} - l_{t+1}^*)(1+\overline{r}) \ \textbf{(endogenous amortization)}
\end{aligned} \tag{31}$$

The value of repossessing a house h is

$$J_t^{REO}(h) = R_t^{REO}(h) - \gamma p_t h + \frac{1-\delta_h}{1+i_{t+1}} J_{t+1}^{REO}(h)$$

$$R_t^{REO}(h) = \max\left\{0, \max_{p_t^{REO} \geq 0} \eta_t^{sell}(p_t^{REO}, h)\left[(1-\chi)p_t^{REO} - \left(-\gamma p_t h + \frac{1-\delta_h}{1+i_{t+1}} J_{t+1}^{REO}(h)\right)\right]\right\} \tag{32}$$

C.4 Housing Market Equilibrium

C.4.1 Market Tightnesses

Market tightnesses satisfy

$$\kappa_b h_t \geq \overbrace{\alpha_{bt}(\theta_{bt}(p_t^{bid}, h_t))}^{\text{prob of match}} \overbrace{(p_t^{bid} - p_t h_t)}^{\text{broker surplus}} \tag{33}$$

$$\kappa_s h_t \geq \underbrace{\alpha_{st}(\theta_{st}(p_t^{list}, h_t))}_{\text{prob of match}} \underbrace{(p_t h_t - p_t^{list})}_{\text{broker surplus}} \tag{34}$$

with $\theta_{bt}(p_t^{bid}, h_t) \geq 0$, $\theta_{st}(p_t^{list}, h_t) \geq 0$, and complementary slackness. Recall that the index p_t is a sufficient statistic for the dependence of tightnesses on the household distribution Φ_t, i.e. $\theta_{bt}(\cdot) \equiv \theta_b(\cdot; p_t(\Phi_t))$ and $\theta_{st}(\cdot) \equiv \theta_s(\cdot; p_t(\Phi_t))$.

C.4.2 Determining the House Price Index

Housing supply $\mathcal{S}_t(p_t)$ includes sales of new construction, owner-occupied houses, and REO inventories:

$$\mathcal{S}_t(p_t) = \overbrace{Y_{ht}(p_t)}^{\text{new housing}} + \overbrace{S_t^{REO}(p_t)}^{\text{REO housing}} + \overbrace{\int h\eta_s(\theta_s(p_t^{list*}, h; p_t)) d\Phi_t^{own}}^{\text{sold by owner}} \tag{35}$$

The supply of REO housing is given by

$$S_t^{REO}(p_t) = \sum_{h \in H} h\eta_s(\theta_s(p_t^{REO}, h; p_t)) \left[\underbrace{H_t^{REO}(h)}_{\text{existing REOs}} + \underbrace{\int d_t^*[1 - \eta_s(\theta_s(p_t^{list*}, h; p_t))] d\Phi_t^{own}(\cdot; h)}_{\text{new foreclosures after a failed listing}} \right] \tag{36}$$

Housing demand $\mathcal{D}_t(p_t)$ equals housing purchased by matched buyers,

$$\mathcal{D}_t(p_t) = \int h_t^* \eta_b(\theta_b(p_t^{bid*}, h_t^*; p_t)) d\Phi_t^{rent} \tag{37}$$

The house price index p_t sets supply equal to demand,

$$\mathcal{D}_t(p_t) = \mathcal{S}_t(p_t) \tag{38}$$

which is reminiscent of Walrasian market clearing. The key difference is that, in this frictional setting, individual buyers and sellers may fail to transact, in which case they do not immediately appear on either side of equation (38). They may, however, influence future demand and supply if they make another attempt to trade in subsequent periods.

C.5 Equilibrium Definition

In a stationary equilibrium, all prices and quantities are constant, and i and r are exogenous because of the open economy assumption. Also, the supply of new permits $\overline{L}$ is exogenous, which removes the need to solve for their price.

Definition 1 *A stationary recursive equilibrium is*

1. *Household value functions $R^{sell,f}$, $W^{own,f}$, $R^{buy,f}$, $V^{rent,f}$, $V^{own,f}$, V^{amort}, and their associated policy functions*
2. *REO value function J^{REO} and its associated policy function p^{REO}*
3. *Mortgage pricing function q*
4. *Market tightness functions θ_b and θ_s*
5. *Prices w, r_a, and p*
6. *Quantities S_h, N_c, and N_h*
7. *Stationary distributions H^{REO} of REO properties and Φ of households*

such that

1. ***Household Optimality:*** *The value/policy functions solve (1) – (25).*
2. ***Production Optimality:*** *Conditions (26) – (29) are satisfied.*
3. ***Lender Optimality:*** *Conditions (30) – (32) are satisfied.*
4. ***Market Tightnesses:*** *$\theta_b(\cdot;p)$ and $\theta_s(\cdot;p)$ satisfy (33) – (34).*
5. ***Labor Market Clearing:*** *$N_c + N_h = \int \int_E e \cdot zF(de)d\Phi$.*
6. ***House Price Index:*** *$\mathcal{D}(p) = \mathcal{S}(p)$.*
7. ***Stationarity:*** *H^{REO} and Φ are invariant with respect to the Markov process induced by the exogenous processes and relevant policy functions.*

D Calibration and Computation

D.1 Income Dynamics

As explained in section III, quarterly income processes cannot be estimated directly from the PSID because it is annual data. Instead, a labor process is specified akin to that in Storesletten, Telmer and Yaron (2004) except without the life cycle or permanent shock at birth. Their values for the autocorrelation of the persistent shock and the variances of the persistent and transitory shocks are transformed into quarterly values in the manner described below.

D.1.1 Persistent Shocks

It is assumed that in each period households play a lottery in which, with probability 3/4, they receive the same persistent shock as they did in the previous period, and with probability 1/4, they draw a new shock from a transition matrix calibrated to the persistent process in Storesletten et al. (2004) (in which case they still might receive the same persistent labor shock). This is equivalent to choosing transition probabilities that match the expected amount of time that households expect to keep their current shock. Storesletten et al. (2004) report an annual autocorrelation coefficient of 0.952 and a frequency-weighted average standard deviation over expansions and recessions of 0.17. The Rouwenhorst method is used to calibrate this process, which gives the following transition matrix:

$$\widetilde{\pi}_z(\cdot,\cdot) = \begin{pmatrix} 0.9526 & 0.0234 & 0.0006 \\ 0.0469 & 0.9532 & 0.0469 \\ 0.0006 & 0.0234 & 0.9526 \end{pmatrix}$$

As a result, the transition matrix *prior to adding the fourth state corresponding to the top 1%* is

$$\pi_z(\cdot,\cdot) = 0.75I_3 + 0.25\tilde{g}_z(\cdot,\cdot) = \begin{pmatrix} 0.9881 & 0.0059 & 0.0001 \\ 0.0171 & 0.9883 & 0.0171 \\ 0.0001 & 0.0059 & 0.9881 \end{pmatrix}$$

D.1.2 Transitory Shocks

Storesletten et al. (2004) report a standard deviation of the transitory shock of 0.255. To replicate this, it is assumed that the annual transitory shock is actually the sum of four, independent quarterly transitory shocks. The same identifying assumption as in Storesletten et al. (2004) is used, namely, that all households receive the same initial persistent shock. Any variance in initial labor income is then due to different draws of the transitory shock. Recall that the labor productivity process is given by

$$\ln(e \cdot z) = \ln(e) + \ln(z)$$

Therefore, total labor productivity (which, when multiplied by the wage w, is total wage income) over a year in which s stays constant is

$$(e \cdot z)_{\text{year 1}} = \exp(z_0)[\exp(e_1) + \exp(e_2) + \exp(e_3) + \exp(e_4)]$$

For different variances of the transitory shock, total annual labor productivity is simulated for many individuals, logs are taken, and the variance of the annual transitory shock is computed. It turns out that quarterly transitory shocks with a standard deviation of 0.49 give the desired standard deviation of annual transitory shocks of 0.255.

D.2 Computation

The household problem is solved using value function iteration. The state space $(y, (\overline{r}, m), h, z)$ for homeowners with good credit is discretized using 275 values for y, 2 values for $\overline{r}$ (the long-run mortgage rate and also its value during quantitative easing), 131 values for m, 3 values for h, and 4 values for z. Homeowners with a default flag, $f = 1$, have state (y, h, z), and renters have state (y, z, f). To compute the equilibrium transition path, the algorithm starts with an initial guess for the path of the house price index, $\{p_t\}_{t=1}^{T}$. The algorithm then does backward induction on the recursive mortgage pricing equation and the household Bellman equations before forward iterating on the distribution of households and REO properties. New equilibrium house prices $\{\widehat{p}_t\}_{t=1}^{T}$ are calculated period-by-period during the forward iteration. If the guessed and solved price paths differ, a convex combination of these two sequences is used for the next guess. This process continues until convergence.

References

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, "Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class," *Review of Financial Studies*, 2016, *29* (7), 1635–1670.

Aladangady, Aditya, "Housing Wealth and Consumption: Evidence from Geographically Linked Microdata," *American Economic Review*, 2017, *107* (11), 3415–3446.

Albanesi, Stefania, Giacomo DeGiorgi, and Jaromir Nosal, "Credit Growth and the Financial Crisis: A New Narrative," 2017. Working Paper.

Arellano, Cristina, Yan Bai, and Patrick Kehoe, "Financial Frictions and Fluctuations in Volatility," *Journal of Political Economy*, Forthcoming 2019.

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme, "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework," *Econometrica*, May 2017, *85* (3), 693–734.

Berger, David, Ian Dew-Becker, and Stefano Giglio, "Uncertainty Shocks as Second-Moment News Shocks," *Review of Economic Studies*, Forthcoming 2019.

Bhutta, Neil, "The Ins and Outs of Mortgage Debt During the Housing Boom," *Journal of Monetary Economics*, 2015, *76*, 284–298.

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry, "Really Uncertain Business Cycles," *Econometrica*, 2018, *86* (3), 1031–1065.

Foote, Christopher L., Lara Loewenstein, and Paul S. Willen, "Cross-Sectional Patterns of Mortgage Debt during the Housing Boom: Evidence and Implications," 2016. Working Paper.

Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Víctor Ríos-Rull, "Intergenerational Redistribution in the Great Recession," 2019. Working Paper.

Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng, "Measuring Uncertainty," *American Economic Review*, 2015, *105* (3), 1177–1216.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante, "The Housing Boom and Bust: Model Meets Evidence," 2019. Working Paper.

Kozeniauskas, Nicholas, Anna Orlik, and Laura Veldkamp, "What Are Uncertainty Shocks?," 2018. Working Paper.

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

_ , Kamalesh Rao, and Amir Sufi, "Household Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics*, 2013, *128* (4), 1687–1726.

Shoag, Daniel and Stan Veuger, "Uncertainty and the Geography of the Great Recession," *Journal of Monetary Economics*, 2016, *84*, 84–93.

Stock, James H. and Mark W. Watson, "Disentangling the Channels of the 2007 – 2009 Recession," *Brookings Papers on Economic Activity*, Spring 2012.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, June 2004, *112* (3), 695–717.